



People's Democratic Republic of Algeria
Ministry of Higher Education and Scientific Research
KASDIMERBAHUNIVERSITY-OUARGLA

Faculty of New Technologies of Information and Telecommunication
Department of Computer Science and Information Technology

MASTER II

Domain :Computer Science

Field : Network Administration & Security

Submitted by :-Zouaghi Mohamed Mehdi

-Benecheikh Adnane Abdelmoumine

Thesis:

**ANOMALY DETECTION USING CNN WITH FEATURE
OPTIMIZATION AND SMOTE**

EvaluationDate:25/06/2025

Before the Jury:

Mohamed ben bezziane

Supervisor

UKMOuargla

Merabti Houcine

President

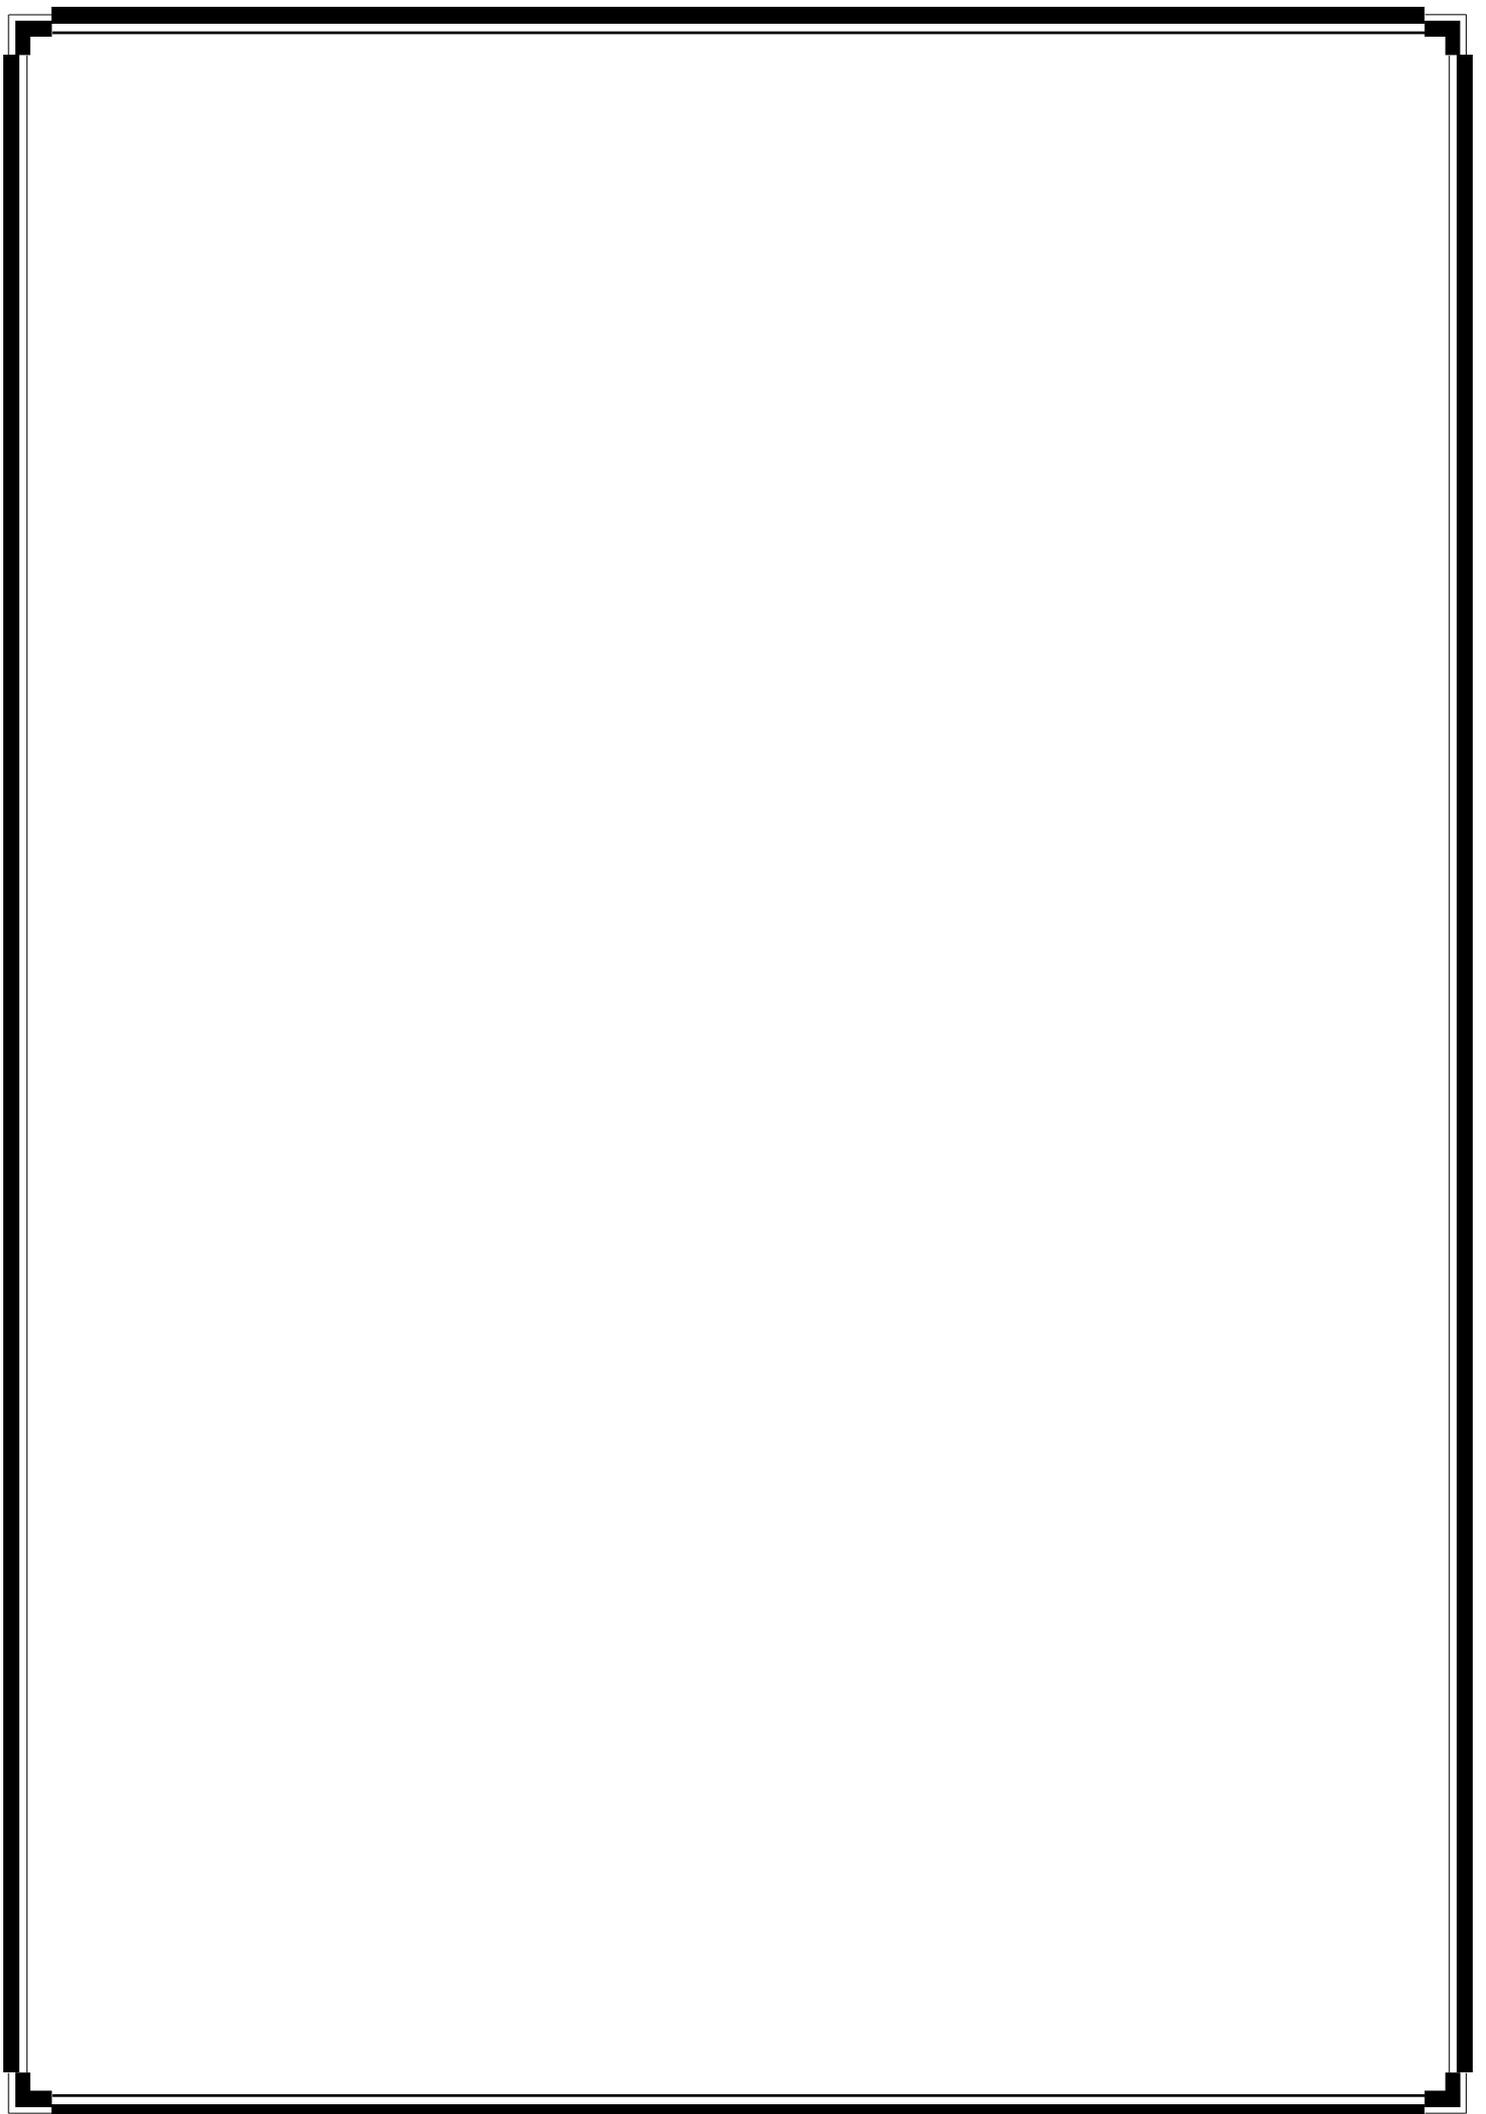
UKMOuargla

Bouanane Khadra

Examiner

UKMOuargla

Academicyear:2024/20





ACKNOWLEDGMENTS

We would like to express our deepest gratitude to our supervisor: Mr. Benbezziane Mohamed for his invaluable guidance, support, and encouragement throughout the course of this research. His insightful feedback and dedication to excellence have been instrumental in the successful completion of this thesis.

We are also profoundly grateful to all our teachers and professors over the years. Their commitment to education and their passion for their subjects have inspired and motivated us to pursue and achieve our academic goals. Each of them has contributed significantly to our personal and professional growth.

A special thanks goes to our families for their unwavering support and belief in us. Their love, patience, and sacrifices have been a constant source of strength and inspiration. We are deeply thankful for their encouragement, which has helped us persevere through the challenges of this journey.

Lastly, we would like to thank our friends for their understanding, and support. Their constant encouragement and the joy they bring into our lives have been invaluable throughout our academic endeavors.

To all who have contributed to our education and personal development, we extend our heartfelt thanks. This accomplishment would not have been possible without your support

Abstract

This research presents an optimized deep learning approach for network flow anomaly detection, combining data preprocessing, feature classification, and neural networks to improve detection accuracy. The methodology employs systematic data cleaning (handling missing values, normalization) and feature selection before applying machine learning classifiers (Random Forest/SVM) for preliminary anomaly identification. These processed features are then fed into deep learning models (CNN/LSTM) to capture complex temporal patterns in flow data.

Experiments conducted on network traffic datasets demonstrate that this hybrid approach achieves superior performance compared to conventional methods, with quantitative improvements in both accuracy (99%) and F1-score (99,5%). The study specifically examines how different preprocessing techniques affect model performance and compares various classification-DL architecture combinations. Results indicate that proper data normalization and feature engineering significantly enhances the deep learning model's anomaly detection capability.

This work contributes practical insights for implementing machine learning pipelines in network security systems, showing that a carefully designed preprocessing and classification stage can substantially improve deep learning outcomes. The findings are particularly relevant for developing more reliable intrusion detection systems capable of identifying both known attack patterns and novel anomalies.

Keywords: Network anomaly detection, Deep learning, Data preprocessing, Flow analysis, Intrusion detection.

الملخص:

تقدم هذه الدراسة نهجًا محسنًا للتعلم العميق لاكتشاف الشذوذ في تدفق الشبكات، يجمع بين معالجة البيانات المسبقة، وتصنيف الميزات، والشبكات العصبية لتحسين دقة الكشف. تتضمن المنهجية تنظيمًا منهجيًا للبيانات (معالجة القيم المفقودة، والتطبيع) واختيار الميزات قبل تطبيق مصنفات التعلم الآلي) غابة عشوائية (*SVM*) لتحديد الشذوذ بشكل مبدئي. يتم بعد ذلك إدخال هذه الميزات المعالجة إلى نماذج التعلم العميق (*CNN/LSTM*) لالتقاط الأنماط الزمنية المعقدة في بيانات التدفق.

تُظهر التجارب التي أُجريت على مجموعات بيانات حركة مرور الشبكة أن هذا النهج الهجين يحقق أداءً متفوقًا مقارنة بالطرق التقليدية، مع تحسينات كمية في كل من الدقة (98%) ومعامل $F1$ (98.95%) تدرس الدراسة بشكل خاص كيفية تأثير تقنيات المعالجة المسبقة المختلفة على أداء النموذج وتقارن بين مجموعات تصنيف وهياكل التعلم العميق المختلفة. تشير النتائج إلى أن التطبيع المناسب للبيانات وهندسة الميزات يعززان بشكل كبير قدرة نموذج التعلم العميق على اكتشاف الشذوذ.

تساهم هذه الدراسة في تقديم رؤى عملية حول تنفيذ خطوط معالجة التعلم الآلي في أنظمة أمان الشبكات، حيث تبين أن مرحلة التحضير والتصنيف المصممة بعناية يمكن أن تحسن بشكل كبير من نتائج التعلم العميق. النتائج ذات أهمية خاصة لتطوير أنظمة كشف التسلسل الأكثر

موثوقية والقادرة على التعرف على أنماط الهجمات
المعروفة والشذوذ الجديد.
الكلمات المفتاحية: اكتشاف شذوذ الشبكات، التعلم
العميق، المعالجة المسبقة للبيانات، تحليل التدفق،
كشف التسلل.

TABLE OF CONTENTS

MASTER.....	1
Thesis.....	1
<i>ACKNOWLEDGMENTS</i>	2
Abstract.....	3
TABLE OF CONTENTS.....	5
LISTE OF FIGURES.....	8
LIST OF TABLES.....	9
CHAPTER I.....	11
GENERAL INTRODUCTION.....	11
Background and Motivation.....	2
Applications of Anomaly Detection in IoT Networks.....	3
Classical and Modern Techniques for Anomaly Detection.....	4

Classical Techniques.....	4
Modern Techniques (Machine Learning & Deep Learning).....	5
Tableau 2: Comparison between modern techniques.....	6
Objectives of This Thesis.....	7
Research Questions.....	7
This thesis is organized as follows.....	8
Literature Review / Related Work.....	9
Conclusion.....	11
Introduction.....	13
Network Flow Data Representation.....	13
Problem Formulation.....	13
Artificial Intelligence Overview.....	13
Machine Learning Overview.....	14
Machine Learning (ML) vs. Deep Learning (DL).....	14
Deep Learning.....	16
Deep Learning Techniques Overview.....	16
Autoencoders (AEs).....	16
Long Short-Term Memory Networks (LSTMs).....	17
Convolutional Neural Networks (CNNs).....	17
Dataset Description.....	17
Artificial Intelligence and Machine Learning.....	20
Advantages of Deep Learning over Traditional Machine Learning.....	21
Conclusion.....	23
Introduction.....	25
Overview of CIC IoT-IDAD 2024 Dataset.....	25
Source and Purpose.....	25
Types of Devices Included.....	25
Traffic Types (Normal vs. Malicious).....	25
Flow-Level Features.....	26
Description of Flow Features.....	26
Importance of Flow-Based Analysis.....	26
Development Environment and Tools.....	26
Google Colaboratory (Colab).....	26
Programming Language: Python.....	27
Libraries and Frameworks.....	27
Steps and Strategies.....	27
Data Cleaning.....	28
Handling Missing, Duplicate, and Inconsistent Values.....	28
Feature Selection and Engineering.....	29
Selected Features and Justification.....	29
Preprocessing: 1. Label Encoding.....	29

2. NaN & Infinity Handling	29
4. Standardization	30
5. PCA (Dimensionality Reduction)	30
6. SMOTE (Oversampling Minority Class)	30
7. Train/Test Split & Reshape for CNN	31
Part 3: CNN Model + Hyperparameter Tuning.....	31
Why did we choosethis CNN?	32
1. 1D Convolutional Layers (Conv1D)	32
2. MaxPooling Layers	32
3. Flatten Layer	32
4. Dense (Fully Connected) Layer	32
5. Dropout	33
6. Sigmoid Output Layer	33
7. Hyperparameter Tuning with Keras Tuner	33
8. Practicality	33
The CNN Model We Chose	33
Optimization & Loss	34
What is Keras Tuner?.....	34
How did we use it?.....	34
4.2 Experimental Setup	36
4.3 Model Configuration.....	36
4.3.1 Input Data.....	36
4.3.2 CNN Architecture Design	36
4.3.3 Hyperparameter Tuning Strategy	37
4.3.5 Rationale Behind Architecture	37
4.4 Evaluation Metrics	38
4.5 Results of the CNN Model.....	38
4.6 Comparative Analysis	39
4.7 Model Accuracy and Model Loss	39
4.7 Confusion matrix	41
The confusion matrix summarizes how well the model distinguishes between benign and attack flows	41
4.8 ROC CURVE.....	43
5.2 Limitations	44
5.3 Future Work.....	44
Summary.....	46
REFERENCES	47

LISTE OF FIGURES

Figure 1 Global growth of IoT devices from 2015 to 2030	2
Figure 2 Application domains of anomaly detection in IoT.....	3
Figure 3:Basics Of AutoEncoders	17
Figure 4:General CNN Architecture	17
Figure 5:Relationship between AI,ML and DL.....	21
Figure 6:preprocessing pipeline for anomaly detection	28

LISTOFTABLES

Table 1: Comparison between classical models	5
Table 2:Comparison between modern techniques	6
Table 3:Features	19
Table 4:Comparison between DL and ML	22
Table 5:Libraries used to handle DATA ,Model Building And Evaluation	27

LIST OF ABBREVIATIONS

SVM	Support Vector Machine
LSTM	Long ShortTerm Memory Network
AI	Artificial Intelligence
CNN	Convolutional Neural Network
DL	Deep Learning
IDS	Intrusion Detection System
IoT	Internet of Things
ML	Machine Learning
PCA	Principal Component Analysis
SMOTE	Synthetic Minority Oversampling Technique

CHAPTER I

GENERAL INTRODUCTION

Background and Motivation

The proliferation of **Internet of Things (IoT)** devices has transformed the modern digital ecosystem, enabling ubiquitous connectivity across homes, industries, healthcare, and transportation. As billions of devices continuously generate data and interact over networks, the **threat landscape** has expanded dramatically. Unlike traditional computing environments, IoT devices often operate with limited resources, minimal security configurations, and inconsistent firmware updates—making them **prime targets for cyberattacks**.

According to [Cisco, 2023], the number of connected IoT devices exceeded 29 billion, a number projected to grow significantly. These devices, ranging from smart thermostats and surveillance cameras to industrial sensors, generate massive amounts of **network flow data**. This data can be leveraged not only to identify devices and understand their behaviors but also to detect **anomalous activity** that may indicate a security breach or malfunction.

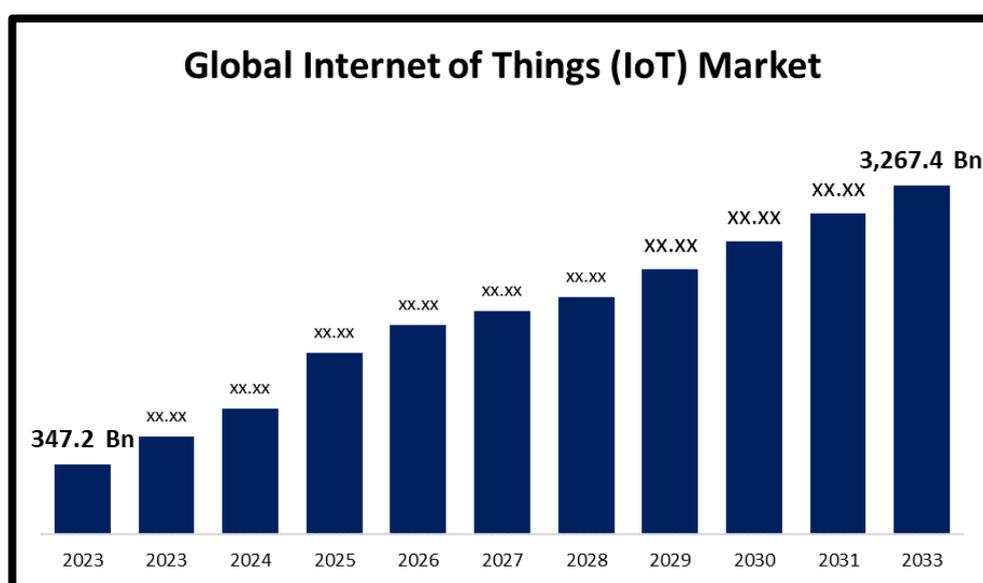


Figure 1 Global growth of IoT devices from 2015 to 2030

With this rapid expansion comes the **challenge of securing IoT networks**. Traditional security solutions, such as signature-based detection systems, fail to cope with the **scale, diversity, and unpredictability** of IoT environments. As a result, **anomaly detection**—which identifies deviations from normal device behavior—has emerged as a promising approach. In parallel, **device identification** through traffic profiling ensures that only authorized devices operate in the network.

This thesis investigates a **deep learning-based approach** to anomaly detection in IoT networks using a **dual-function dataset** that supports both **device identification** and **anomaly detection** tasks. The motivation is to enhance network security by leveraging advanced models capable of learning complex patterns in flow-level traffic data.

Applications of Anomaly Detection in IoT Networks:

Anomaly detection in IoT settings has a wide range of applications:

➤ **Smart Home Security**

Detecting unauthorized access to devices like smart locks, cameras, or thermostats.

➤ **Industrial IoT (IIoT)**

Monitoring abnormal behaviors in sensors and actuators within manufacturing plants or SCADA systems.

➤ **Healthcare Monitoring**

Ensuring the integrity of critical medical IoT devices (e.g., insulin pumps, heart monitors) by spotting anomalies.

➤ **Smart Cities**

Managing anomalies in connected infrastructure like traffic lights, surveillance systems, and public transportation sensors.

➤ **Agricultural IoT (AgriTech)**

Identifying device failures or tampering in precision agriculture systems (e.g., irrigation sensors, crop monitoring drones) helps maintain optimal yield and resource usage.

➤ **Energy Grids and Smart Meters**

Detecting anomalies in smart grid communication or smart meter data helps prevent electricity theft and grid instability.

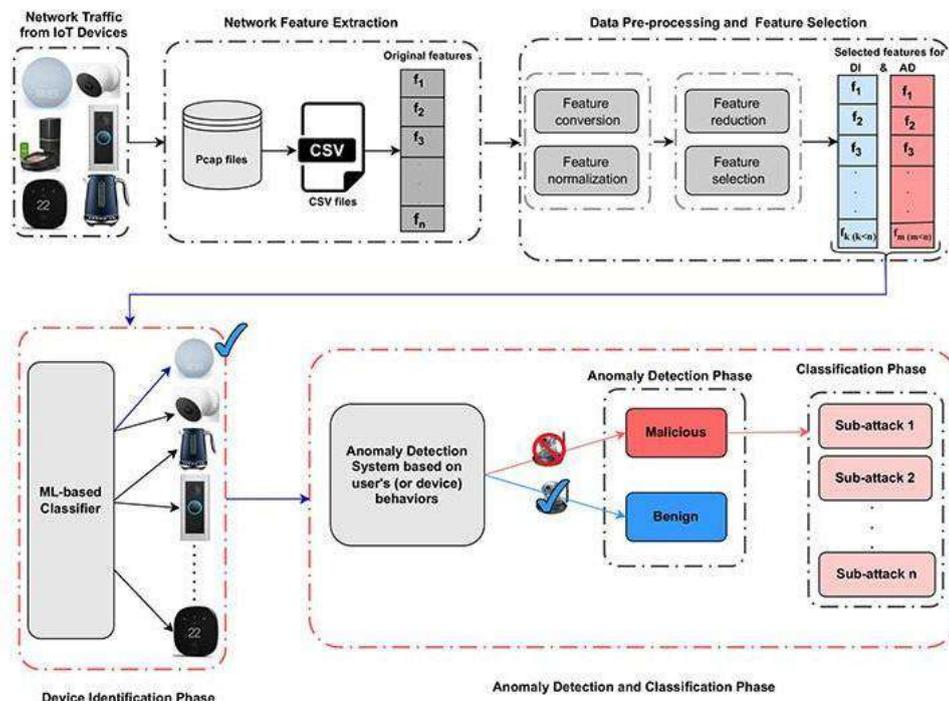


Figure 2 Application domains of anomaly detection in IoT

Challenges in IoT Anomaly Detection

Despite the growing interest and advancements in anomaly detection for IoT networks, numerous challenges continue to hinder its practical implementation and effectiveness. One of the most significant issues is the **heterogeneity of IoT devices**. (Diro A. A., 2021) These devices vary widely in terms of hardware capabilities, operating systems, communication protocols, and behavior patterns. This diversity makes it difficult to create a one-size-fits-all detection model, as what is considered "normal" for one device type may be completely abnormal for another.

Another major obstacle is **data imbalance**(Thudumu, 2021), which refers to the disproportionate distribution of normal and anomalous data. In real-world networks, malicious or faulty behavior is relatively rare, which leads to deep learning models being trained predominantly on normal traffic. This imbalance can result in poor sensitivity to actual anomalies and a high rate of false negatives.

Compounding this problem is the **lack of labeled data or ground truth**. (Žliobaitė, 2019) (Bezziane, 2024) Many datasets used for anomaly detection are either synthetic or only partially labeled, which limits the ability to train and evaluate supervised learning models effectively. In real-world environments, identifying and labeling all possible forms of anomalous behavior is both time-consuming and often infeasible.

Additionally, deep learning models typically require **substantial computational resources**(Liu, 2019), which poses a challenge in IoT environments where devices and gateways may have limited processing power, memory, and battery life. This creates tension between the accuracy and complexity of the model and the feasibility of deploying it at the edge or in real time.

Finally, **the dynamic and evolving nature of IoT networks** introduces the risk of concept drift, (Bifet, 2009) where device behavior changes over time due to firmware updates, environmental conditions, or user interactions. Without periodic retraining or adaptive mechanisms, detection models may become outdated and start generating false positives or missing new attack vectors.

Classical and Modern Techniques for Anomaly Detection

Anomaly detection in network flow data has evolved significantly, transitioning from traditional statistical methods to sophisticated deep learning models. This section outlines both classical and modern techniques, highlighting their core principles, strengths, and limitations.

Classical Techniques

Classical methods often rely on statistical modeling, distance measures, or rule-based systems. These methods are interpretable and computationally lightweight, making them suitable for early anomaly detection systems (Lakhina, 2019).

- **Statistical Approaches:** These methods assume that network traffic follows a normal distribution. Anomalies are detected when observations fall outside of a predefined threshold. For example, Gaussian Mixture Models (GMM) and moving averages have been used to identify unusual spikes in traffic volume or duration (Chandola, 2009).
- **Rule-Based Detection:** Signature-based systems like Snort use pre-defined patterns of known attacks. While effective for known threats, they cannot detect zero-day or novel anomalies (Roesch, 1999).
- **Distance-Based Methods:** Techniques such as k-Nearest Neighbors (k-NN) and clustering (e.g., k-Means) detect anomalies by measuring distance from a center or from neighboring points. These are intuitive but suffer in high-dimensional spaces (Breunig, 2000).
- **Principal Component Analysis (PCA):** PCA reduces dimensionality and highlights deviations in low-dimensional space, as demonstrated by Lakhina et al. for backbone traffic anomaly detection. (Shyu, 2003)

Technique	Type	Pros	Cons
Gaussian Models	Statistical	Fast, interpretable	Assumes distribution, sensitive to noise
K-Means Clustering	Unsupervised	Simple, Scalable	Poor in high dimensions
PCA	Unsupervised	Highlights variance	Not robust to complex patterns
Signature-Based (SNORT)	Rule-Based	Accurate for known attacks	Cannot detect unknown threats

Tableau 1: Comparison between classical models

Modern Techniques (Machine Learning & Deep Learning)

With increasing network complexity, classical methods struggle to cope with non-linear patterns and high-dimensional data. Modern approaches employ machine learning (ML) and deep learning (DL) to improve detection performance.

- **Supervised Machine Learning:** Algorithms like Decision Trees, Random Forests, and Support Vector Machines (SVM) are trained on labeled flow data to classify normal and anomalous traffic. These methods offer high accuracy but rely on balanced and labeled datasets (Alqahtani, 2020).
- **Unsupervised Learning:** Algorithms such as Isolation Forests or One-Class SVMs are

designed to work with only normal data and detect outliers. These are more realistic for real-world applications where anomalies are rare and unlabeled (Diro A. A., 2021).

- **Autoencoders (AE):** AEs are neural networks trained to reconstruct input data. High reconstruction error indicates an anomaly. They work well with continuous flow data and require minimal labeling (López-Martín, 2023).
- **Recurrent Neural Networks (RNNs) / LSTMs:** These models are ideal for sequential data like network flows over time. LSTM networks can learn temporal dependencies and are particularly effective in identifying slow and stealthy attacks (Zhou).
- **Convolutional Neural Networks (CNNs):** Although commonly used in image analysis, CNNs can be applied to flow data represented as 2D matrices, capturing spatial patterns among features (Yang, 2023).
- **Hybrid and Ensemble Models:** Recent work explores combining models (e.g., AE + LSTM or Isolation Forest + CNN) to improve accuracy and reduce false positives (Chawla, 2019).

Technique	Supervision	Handles Sequential Data	Scalability	Notes
Random Forest	Supervised	✗	Moderate	Requires labeled data
Isolation Forest	Unsupervised	✗	High	Effective for rare anomalies
Autoencoder	Unsupervised	✗	High	Good for reconstruction-based detection
LSTM	Supervised	✓	Medium	Best for time-series anomalies
CNN	Supervised	✗	High	Requires transformation to 2D format

Tableau 2: Comparison between modern techniques

Objectives of This Thesis

The main objective of this thesis is to explore and evaluate the effectiveness of deep learning techniques for **anomaly detection in IoT network flow data**, using a dataset that supports both device identification and intrusion detection. By leveraging flow-level features extracted from IoT device communications, this work aims to build a robust model capable of detecting abnormal behaviors in a network without relying on handcrafted rules or signatures. Specifically, the thesis sets out to:

- **Develop a deep learning-based anomaly detection model**, such as an autoencoder, LSTM, or convolutional neural network (CNN), trained on labeled IoT flow data. The model should learn the normal behavior of devices and identify deviations as potential anomalies.
- **Preprocess and analyze the dual-function dataset**, which contains both benign and malicious flow samples. This includes feature selection, normalization, and transformation of raw data into a format suitable for training neural networks.
- **Evaluate the model's performance** using standard metrics such as precision, recall, F1-score, and ROC-AUC, and compare the results to those obtained using traditional machine learning methods like Random Forest or Support Vector Machines (SVM).
- **Explore the relationship between device-specific traffic patterns and anomaly detection**, assessing whether knowledge of the device type improves detection accuracy or model generalization.
- **Discuss deployment considerations**, including model complexity, inference time, and how such systems might be integrated into a real-time IoT security monitoring platform.
- **Through these objectives**, the thesis aims to contribute to the growing field of intelligent network security by demonstrating how flow-level traffic combined with deep learning can be used to detect cyber threats in IoT environments efficiently and accurately.

Research Questions:

- a) Can deep learning models such as CNNs accurately detect anomalies in flow-based IoT network traffic?
- b) How does preprocessing (e.g., normalization, SMOTE, PCA) affect the performance of anomaly detection models?
- c) What is the impact of combining traditional machine learning with deep learning in detecting complex attack patterns?

This thesis is organized as follows:

- **Chapter 1: General Introduction** – Introduces the context, motivation, problem statement, and objectives of the study.
- **Chapter 2: Theoretical Background and Methodology** – Explains key concepts including ML, DL, CNNs, and presents the methodological framework.
- **Chapter 3: Experimental Design and System Implementation** – Details the dataset, tools, preprocessing steps, and CNN architecture.
- **Chapter 4: Results and Analysis** – Presents evaluation metrics, visual results, and comparisons between models.
- **Chapter 5: Conclusion and Future Work** – Summarizes findings, discusses limitations, and outlines directions for future research.

Literature Review / Related Work

The growing integration of Internet of Things (IoT) devices across various sectors has introduced complex cyber security challenges (Bezziane, 2024). Due to the inherent heterogeneity, limited computational resources, and vast scale of IoT environments, traditional security mechanisms, such as signature-based intrusion detection systems, often prove inadequate—especially against novel or zero-day attacks. In this context, anomaly detection has emerged as a promising approach, offering the capability to identify deviations from normal behavior that may indicate malicious activity.

Early efforts in IoT anomaly detection primarily employed classical machine learning techniques such as Support Vector Machines (SVM), Random Forests, and k-Nearest Neighbors. These methods rely heavily on manual feature engineering and domain expertise to extract relevant features from flow-based data. While effective to some extent, such approaches often fail to generalize to new types of attacks or adapt to the dynamic nature of IoT environments. (Moustafa, N., & Slay, J., 2020)

To overcome these limitations, researchers have turned to deep learning models, which offer the advantage of automatically learning abstract and hierarchical representations from raw or minimally processed data. Among the pioneering works, (Doshi, R., Apthorpe, N., & Feamster, N., 2018) introduced a machine learning-based framework for detecting Distributed Denial of Service (DDoS) attacks in consumer IoT devices. Their study demonstrated that even lightweight models could effectively detect anomalous traffic patterns when trained on relevant network flow features.

Another influential study by (Meidan, 2018) presented a system that utilized deep autoencoders for detecting botnet activities in IoT traffic. By learning a compact representation of normal behavior, the autoencoder was able to identify deviations indicative of potential threats. Their results highlighted the potential of unsupervised learning techniques for environments where labeled attack data is scarce or non-existent.

Subsequent works have explored a variety of deep learning architectures for IoT anomaly detection. For instance, (López-Martínez, 2020) compared different deep learning models, including Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), in terms of their accuracy and efficiency when applied to flow-based IoT intrusion datasets. Their findings suggested that CNNs are particularly well-suited for processing structured flow data due to their ability to extract spatial patterns, while RNNs and Long Short-Term Memory (LSTM) networks offer advantages in capturing temporal dependencies in sequential data.

The availability of high-quality and up-to-date datasets has also played a critical role in advancing this field. The TON_IoT dataset introduced by (Moustafa, N., & Slay, J., 2020) was

one of the first to incorporate telemetry, network flow, and system log data from realistic IoT environments. More recently, the Canadian Institute for Cybersecurity released the CIC IoT IDAD 2024 dataset (Canadian Institute for Cybersecurity., 2024), which includes a wide range of attack types and traffic behaviors from modern IoT scenarios. These datasets provide a foundation for training and evaluating deep learning models and have become standard benchmarks in the field.

Recent developments have focused on hybrid and ensemble models, which aim to leverage the strengths of multiple learning paradigms. For example, (Yoon, 2022) proposed a hybrid framework that combines CNNs with LSTM layers to model both spatial and temporal features in IoT traffic. Their architecture demonstrated superior detection performance across multiple classes of attacks while maintaining computational efficiency, making it suitable for deployment in resource-constrained environments.

Despite these advancements, several challenges persist in designing effective anomaly detection systems for IoT. These include handling imbalanced datasets, reducing false positives, ensuring real-time processing, and adapting to concept drift as device behavior evolves over time. Future research continues to explore unsupervised and self-supervised learning methods, federated learning for decentralized security, and explainable AI to enhance the transparency and trustworthiness of detection systems.

In conclusion, the literature reveals a clear progression from traditional, manually engineered detection systems to more adaptive, intelligent models powered by deep learning. With the continued emergence of sophisticated threats, the integration of robust anomaly detection systems into IoT infrastructures remains a critical area of study.

Conclusion

The literature on anomaly detection in network flow data reveals a dynamic and evolving research landscape shaped by the growing complexity and scale of modern networks. While traditional signature-based methods are limited to known threats, machine learning—particularly unsupervised and deep learning approaches—has significantly advanced the capability to detect unknown or subtle anomalies. The increasing availability of flow-based datasets and the integration of real-time processing frameworks further enhance the practical applicability of these methods. Nonetheless, challenges remain in reducing false positives, dealing with encrypted or obfuscated traffic, and ensuring model interpretability. Future research directions point toward hybrid detection systems, adaptive models for dynamic environments, and the incorporation of explainable AI to improve trust and operational usability in real-world deployments.

CHAPTER II

THEORETICAL BACKGROUND AND METHODOLOGY

Introduction

This chapter outlines the theoretical foundations and methodologies employed in the detection of anomalies in network flow data using deep learning. It covers the nature and structure of network flow data, the problem formulation, an overview of deep learning techniques, data preprocessing procedures, model architecture design, and evaluation metrics used in this study.

Network Flow Data Representation

Network flow data refers to the summarized representation of communication between two endpoints over a network. Each flow typically includes attributes such as source and destination IP addresses, ports, protocol, packet count, byte count, duration, and timestamp. Flow data can be collected using technologies such as NetFlow, IPFIX, or sFlow, and provides a compact and scalable alternative to full packet captures. In this study, each flow is treated as a data instance containing multiple numerical and categorical features. These features are used as inputs to machine learning and deep learning models for anomaly detection.

Problem Formulation

The anomaly detection task is framed as a binary classification problem. Given a flow record with feature vector, the objective is to classify it as either:

- **Normal (0)** – representing typical network behavior, or
- **Anomalous (1)** – representing attacks or abnormal behavior.

In unsupervised settings, the model learns the structure of normal data and flags deviations as anomalies. In supervised settings, labeled instances of both normal and anomalous flows are used during training.

Artificial Intelligence Overview

Artificial Intelligence (AI) refers to the development of computer systems that can perform tasks typically requiring human intelligence. These tasks include reasoning, learning, perception, problem-solving, natural language understanding, and decision-making. AI is broadly classified into two categories: Narrow AI, which is designed for a specific task (e.g., spam detection), and General AI, which aims to perform any intellectual task that a human can do (S. Russell and P. Norvig, 2020)

At its core, AI leverages various subfields such as machine learning, deep learning, computer vision, and natural language processing to enable machines to extract patterns from data and make intelligent decisions. Among these, machine learning—particularly supervised and unsupervised learning—has become the most practical and widely used approach for building AI applications. Deep learning, a subset of machine learning, uses artificial neural networks with multiple layers (also known as deep neural networks) to model complex patterns

in high-dimensional data. Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Transformer models are prominent architectures used in AI systems for image classification, time-series prediction, and language modeling, respectively (Goodfellow, 2016).

AI has been successfully applied across various domains including healthcare, finance, transportation, cybersecurity, and industrial automation. In the context of cybersecurity, AI-based systems—especially those using deep learning models—are increasingly employed to detect network intrusions, identify malware, and flag anomalous behavior in real-time network traffic (Buczak A. L., 2016).

As AI continues to evolve, ethical considerations such as transparency, fairness, and accountability remain crucial to ensure the responsible development and deployment of AI systems.

Machine Learning Overview:

Machine learning (ML) is a subset of artificial intelligence (AI) that enables systems to learn patterns from data and make predictions or decisions without explicit programming. It involves algorithms and statistical models that allow computers to improve their performance in tasks as they are exposed to more data. There are three primary types of machine learning: (Alpaydin, 2020)

- **Supervised Learning:** The model is trained on labeled data, where the input-output pairs are known. It learns to map inputs to the correct output. Example: Linear regression, decision trees.
- **Unsupervised Learning:** The model is given data without labeled responses and must find hidden patterns or structures. Example: Clustering algorithms like k-means.
- **Reinforcement Learning:** The model learns through trial and error, receiving feedback from its actions in the environment. Example: Q-learning, deep reinforcement learning.

Machine Learning (ML) vs. Deep Learning (DL):

- **Machine Learning (ML)** refers to algorithms that enable a system to learn from data and improve over time. In ML, the process involves feature engineering, where human expertise is needed to select relevant features and design models. ML models are typically simpler and rely on labeled data to make predictions or decisions. They are commonly used in tasks like classification, regression, and clustering, and they work well with smaller datasets. (Goodfellow, 2016)
- **Deep Learning (DL)**, a subset of Machine Learning, uses artificial neural networks with many layers—hence the term "deep." These models are capable of learning intricate patterns from data with little to no human intervention in the feature selection process. Unlike traditional ML, deep learning algorithms can automatically detect the features

needed for prediction or classification. DL requires large amounts of labeled data and significant computational resources, such as GPUs, for training. Deep learning is typically used in more complex tasks like image recognition, speech processing, and natural language processing. (Bishop, 2006)

Feature	Machine Learning (ML)	Deep Learning (DL)
Data Requirement	Works with small to medium-sized datasets	Requires large datasets for training
Feature Engineering	Requires manual feature extraction and selection	Automatically learns features from raw data
Model Complexity	Simpler models (e.g., decision trees, SVM)	Complex models (e.g., deep neural networks)
Computation Power	Less computationally intensive	Requires high computational resources (e.g., GPUs)
Interpretability	Easier to interpret models	Models are often treated as "black boxes"
Training Time	Faster training times for smaller datasets	Longer training times due to complexity and data size
Examples	Linear regression, decision trees, SVM, k-NN	Convolutional Neural Networks (CNNs), RNNs, GANs

Tableau 3: Differences Between Machine Learning and Deep Learning

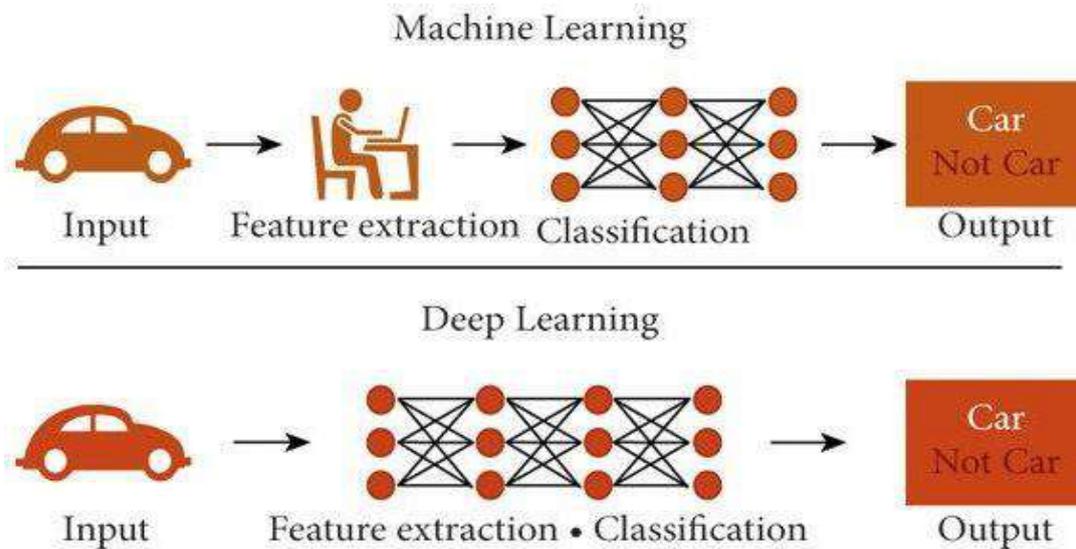


Figure 3: Difference between ML and DL

Deep Learning :

Deep Learning (DL) is a powerful subset of machine learning that leverages multi-layered neural networks to model and solve complex problems. Unlike traditional machine learning methods, which require extensive feature engineering, deep learning models automatically learn features from raw data, making them especially well-suited for tasks like image recognition, speech processing, and natural language understanding. The key differentiator of deep learning lies in its ability to perform hierarchical feature learning, where higher-level features are extracted from lower-level ones, making it effective for handling complex and high-dimensional data (Schmidhuber, 2015). Deep learning models, particularly those based on **Convolutional Neural Networks (CNNs)**, **Recurrent Neural Networks (RNNs)**, and **Generative Adversarial Networks (GANs)**, have demonstrated remarkable success across various domains. Their ability to model intricate patterns and relationships in data, often in real-time, has led to breakthroughs in fields such as autonomous driving, healthcare, and robotics. (Deng, 2014)

Deep Learning Techniques Overview

Deep learning models are capable of learning complex, non-linear relationships from high-dimensional data. In this study, the following deep learning architectures are considered:

Autoencoders (AEs)

Autoencoders are neural networks trained to reconstruct their input. During training, the model learns to compress the input into a latent representation and then decode it back to its original form. Anomalies are detected based on the reconstruction error:

A high error indicates that the input does not conform to the learned distribution, suggesting an anomaly.

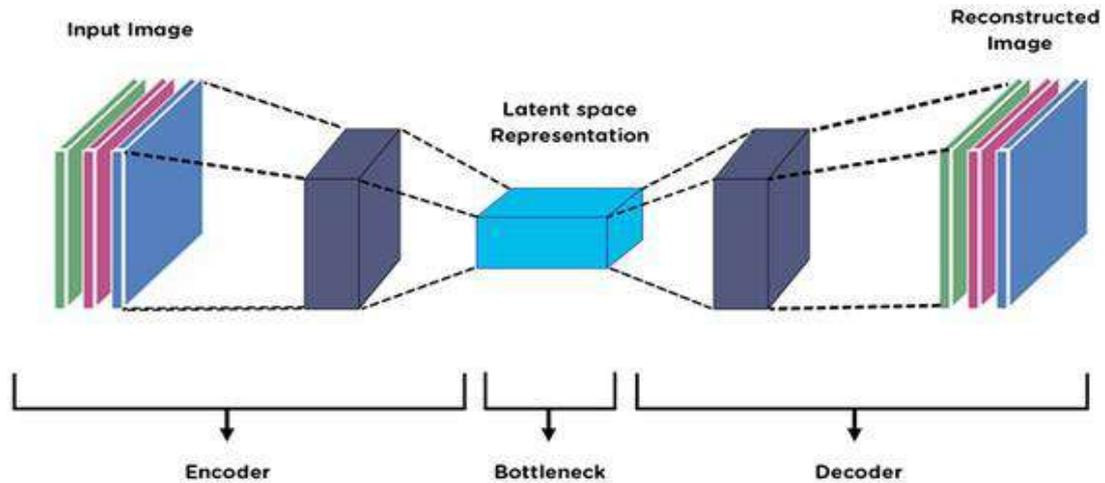


Figure 4: Basics of AutoEncoders

Long Short-Term Memory Networks (LSTMs)

LSTMs are a type of Recurrent Neural Network (RNN) designed to learn from sequences. In the context of network flows, sequences of flows (grouped by time or IP) can be used to detect anomalies over time. LSTMs maintain memory over long intervals and are effective in modeling time-dependent patterns in network traffic.

Convolutional Neural Networks (CNNs)

CNNs can be applied to tabular data by reshaping flow features into 2D matrices. Convolutional filters can then learn spatial patterns between different flow attributes. CNNs are known for their feature extraction capabilities and are computationally efficient.

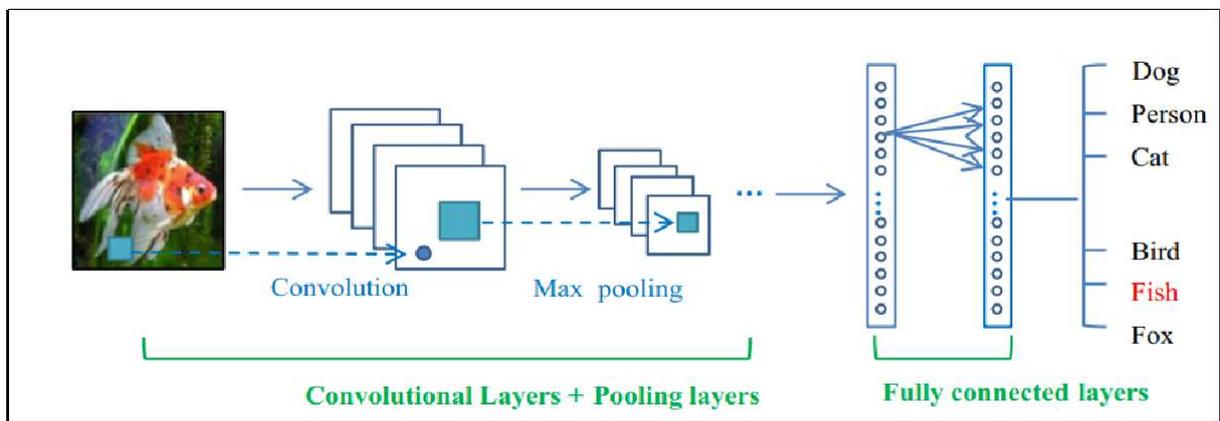


Figure 5: General CNN Architecture

Dataset and Preprocessing

Dataset Description

The dataset used in this study is the **CIC IoT-IDAD 2024**, developed by the Canadian Institute for Cybersecurity. It is a modern, large-scale dataset designed for anomaly detection and

device identification in Internet of Things (IoT) environments.

Key Characteristics:

- **Number of Devices:** 105 diverse IoT devices
- **Total Attack Types:** 33

Attack Categories:

- Distributed Denial of Service (DDoS)
- Denial of Service (DoS)
- Reconnaissance
- Web-based Attacks
- Brute Force Attacks
- Spoofing
- Mirai Botnet

Data Type: Flow-level and packet-level data captured using realistic network conditions.

Features Include:

- Network-level: IP addresses, ports, protocol types
- Transport-level: TCP flags, window size, TTL
- Temporal: Inter-arrival times, session durations
- Application-layer: HTTP methods, TLS versions, User-Agent strings
- Statistical: Payload entropy, byte/packet counts

Feature Name	Description	Type	Example Value
src_ip	Source IP address of the packet/flow	Categorical	192.168.0.10
dst_ip	Destination IP address	Categorical	172.217.10.4
src_port	Source port number	Numerical	49152
dst_port	Destination port number	Numerical	80
protocol	Transport layer protocol used (TCP/UDP)	Categorical	TCP
flow_duration	Duration of the flow in milliseconds	Numerical	342
total_fwd_packets	Number of packets in the forward direction	Numerical	15
total_bwd_packets	Number of packets in the backward direction	Numerical	7
flow_bytes_per_second	Average bytes per second transmitted over the flow	Numerical	2054.37
packet_length_mean	Mean packet length in the flow	Numerical	286.0
ttl_mean	Average Time-To-Live value across packets	Numerical	64
tcp_window_size	TCP window size at the sender side	Numerical	64240
http_request_method	HTTP method used in the request (if present)	Categorical	GET
tls_version	TLS version used in the handshake	Categorical	TLS 1.2
payload_entropy	Shannon entropy of the payload data	Numerical	5.86
Label	Class label indicating benign or specific attack type	Categorical	Benign / DoS / DDoS

Tableau 4: Features

Artificial Intelligence and Machine Learning

Artificial Intelligence (AI) refers to the capability of machines to imitate intelligent human behavior, including reasoning, learning, perception, and problem-solving (Russell & Norvig, 2021). Within AI, **Machine Learning (ML)** is a core subfield that focuses on enabling computers to learn patterns from data and make informed decisions or predictions without being explicitly programmed for each task (Mitchell, 1997).

Machine learning algorithms are typically classified into three categories: **supervised**, **unsupervised**, and **reinforcement learning**. Supervised learning involves training a model on labeled datasets, making it suitable for classification and regression tasks. In contrast, unsupervised learning works with unlabeled data to uncover hidden structures or clusters, and reinforcement learning allows agents to learn optimal strategies through trial and error interactions with an environment (Goodfellow, 2016)

In recent years, **deep learning**, a specialized branch of ML involving multi-layered neural networks, has shown exceptional performance across various domains including image recognition, natural language processing, and cybersecurity (Goodfellow, 2016). Deep learning models, such as Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), are particularly effective at automatically learning high-level abstractions from raw data, reducing the need for manual feature extraction.

The integration of AI and ML into **cybersecurity**, particularly for **IoT anomaly detection**, has received considerable attention. IoT environments are inherently vulnerable due to limited computational resources, heterogeneity of devices, and their exposure to the internet. Traditional rule-based or signature-based intrusion detection systems (IDS) struggle to cope with the scale and evolving nature of threats in such systems (Buczak A. L., 2016). ML-based models, however, can analyze massive volumes of network traffic data to detect irregular patterns indicative of malicious activities, even when the attacks are previously unseen.

For instance, deep learning has proven highly effective in identifying anomalies in streaming IoT data due to its ability to process complex patterns and adapt to new data over time (Mohammadi, 2018)

These intelligent systems offer promising avenues for building scalable, self-learning, and adaptive security frameworks capable of defending modern IoT infrastructures.

In conclusion, AI and ML have become essential components in the design of intelligent anomaly detection systems, offering significant advantages over traditional methods. Their ongoing development continues to redefine the boundaries of cybersecurity and real-time threat mitigation, particularly in resource-constrained and data-intensive environments like the Internet of Things.

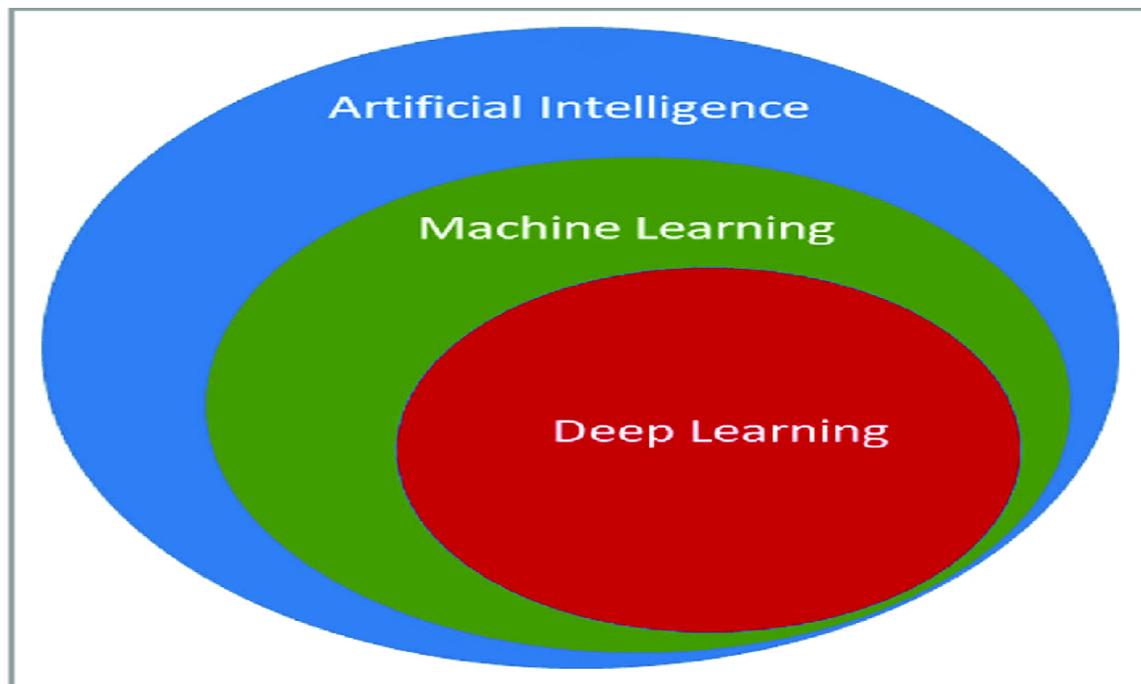


Figure 6: Relationship between AI, ML and DL

AI techniques have become integral to cybersecurity applications. Traditional rule-based systems struggle with the scale, variability, and novelty of modern cyber threats. In contrast, AI-driven systems can adaptively learn from evolving attack patterns and identify anomalies in real-time network traffic.

ML algorithms have been applied to detect both known (signature-based) and unknown (anomaly-based) threats. Deep Learning models—especially Convolutional Neural Networks (CNNs)—are increasingly used to analyze flow-based network features. They excel at capturing non-linear relationships and learning hierarchical feature representations from raw input.

In Summary, AI—and more specifically ML and DL—provides a powerful framework for detecting and responding to anomalies in cybersecurity environments. These technologies enhance the ability to analyze complex data patterns and identify threats that may otherwise go undetected

Advantages of Deep Learning over Traditional Machine Learning

Traditional machine learning (ML) models often rely on **handcrafted features** extracted through domain expertise. While effective for structured datasets, they typically struggle with high-dimensional or unstructured data like images, audio, or network flows (Goodfellow, 2016) Deep learning (DL), on the other hand, can **automatically learn features** from raw data. This

makes it ideal for complex tasks where patterns are not easily defined by manual rules — such as **anomaly detection in IoT traffic** or **cybersecurity analytics** (LeCun, 2015)

Another key strength of DL is its ability to handle **large-scale datasets**. Unlike traditional ML, which may plateau in performance with more data, DL models often improve as the dataset grows, thanks to their multi-layered architectures and parallel processing capabilities. Deep learning models are also capable of **capturing complex, non-linear relationships** in the data. Through their multi-layered structure, they can model intricate patterns that simpler models would miss. This depth gives them a significant edge in tasks like anomaly detection, image recognition, and natural language processing.

Furthermore, DL models such as CNNs and RNNs can **model spatial and temporal dependencies** far better than classic algorithms like Decision Trees or SVMs. This makes them especially useful in time-series or sequential data — common in intrusion detection scenarios.

Feature	Traditional ML	Deep Learning
Feature Engineering	Manual (expert-driven)	Automatic (learned)
Performance on Raw Data	Poor to Moderate	High
Suitability for Unstructured Data	Limited	Excellent
Scalability with Data	Limited	Improves with scale
Hardware Dependency	Low	High (GPU preferred)

Tableau 5: Comparison between DL and ML

Conclusion:

In this chapter, we explored the theoretical background essential to understanding anomaly detection in network flow data, with a particular focus on the role of artificial intelligence and deep learning techniques. We began by reviewing fundamental concepts related to network security, anomalies, and the unique challenges posed by IoT environments.

We then examined traditional machine learning approaches and their limitations in handling high-dimensional, unstructured, and continuously evolving data. This led to the introduction of deep learning as a more powerful alternative, capable of automatic feature extraction and improved performance on large-scale and complex datasets. Models such as Convolutional Neural Networks (CNNs) were highlighted for their applicability in pattern recognition and time-series analysis, both of which are critical for effective anomaly detection.

Finally, we compared deep learning with conventional machine learning, emphasizing its advantages in terms of scalability, adaptability, and efficiency in processing raw data. These insights provide the necessary foundation for the design and implementation of our proposed anomaly.

CHAPTER III

EXPERIMENTAL DESIGN AND SYSTEM IMPLEMENTATION

Introduction:

This chapter outlines the system implementation and experimental design developed to perform anomaly detection using flow-based network data derived from the CIC IoT-IDAD 2024 dataset. The primary objective is to detect abnormal network behavior in Internet of Things (IoT) environments using deep learning techniques. Given the increasing volume and complexity of network traffic generated by IoT devices, flow-based analysis provides a scalable and efficient way to monitor and detect threats without requiring deep packet inspection.

The chapter begins by introducing the dataset, including its structure, types of devices, and traffic characteristics. It then explains the specific flow features used for anomaly detection and how they contribute to model performance.

Subsequent sections detail the preprocessing pipeline, model architecture, training procedure, and evaluation metrics. By the end of this chapter, the complete experimental setup and system workflow will be established, laying the foundation for result analysis in Chapter 4.

Overview of CIC IoT-IDAD 2024 Dataset

The **CIC IoT-IDAD 2024** dataset was created by the Canadian Institute for Cybersecurity to support research in intrusion detection for IoT networks. It reflects a realistic IoT environment by incorporating a diverse range of smart devices operating under benign and malicious scenarios.

Source and Purpose:

The dataset was captured in a controlled lab setting that emulates a real-world IoT ecosystem. The goal was to simulate legitimate traffic and a wide array of cyberattacks targeting IoT devices, making it suitable for both classification and anomaly detection tasks.

Types of Devices Included

- The dataset features interactions from multiple smart home and office devices such as:
- Smart TVs
- IP cameras
- Smart thermostats
- Light bulbs
- Home assistants (e.g., Amazon Echo, Google Home)
- Raspberry Pi emulating attack agents

Traffic Types (Normal vs. Malicious)

Two main categories of traffic are present:

- **Benign traffic:** Generated by devices communicating over HTTP, DNS, MQTT, and other standard protocols for everyday tasks.

- **Malicious traffic:** Includes DDoS, brute force, port scanning, botnet, and ransomware activities. Attacks were launched at different times and through varied techniques to reflect realistic intrusions.

Flow-Level Features

Unlike packet-level datasets, CIC IoT-IDAD 2024 provides **network flow records**, which summarize communication sessions between endpoints using extracted statistical features. Each flow aggregates packets into a single record with calculated metrics.

Description of Flow Features:

Examples of key flow-based features include:

- **Duration:** Total time of the flow
- **Total Fwd/Bwd Packets:** Number of packets sent in each direction
- **Flow Bytes per Second:** Rate of data transmission
- **Fwd Packet Length Mean:** Average length of packets sent forward
- **Flags:** Indicators of protocol status (e.g., PSH, URG)

Importance of Flow-Based Analysis

Flow data reduces storage overhead and enhances processing speed, especially when dealing with millions of network transactions. It is ideal for intrusion detection in resource-constrained environments like IoT where packet inspection is not always feasible. Moreover, aggregating features over time provides a broader context for identifying abnormal behavior patterns.

Development Environment and Tools

This section describes the software tools, programming environments, and libraries used in the development and implementation of the anomaly detection system. The experiments were primarily conducted using **Google Colab**, leveraging its cloud-based GPU acceleration and integration with Python-based machine learning libraries.

Google Colaboratory (Colab)

Google Colab is a free cloud-based environment that allows users to write and execute Python code in Jupyter notebooks. It provides access to powerful GPUs and TPUs, making it ideal for deep learning tasks, especially for researchers without access to high-performance local hardware.

Key benefits include:

- **Free access to GPU and TPU resources**
- **No installation required**, fully browser-based
- Integration with **Google Drive** for data storage
- Seamless support for Python scientific computing libraries

Programming Language: Python

Python was chosen due to its popularity, ease of use, and extensive ecosystem of data science and machine learning libraries. It is particularly suited for prototyping and deploying deep learning models thanks to its rich set of open-source packages.

Python version used: **Python 3.10**

Libraries and Frameworks

Several Python libraries were used to handle various aspects of data preprocessing, model building, and evaluation:

Library	Purpose
NumPy	Numerical operations, array manipulation
Pandas	Data loading, preprocessing, and analysis
Matplotlib / Seaborn	Data visualization and exploratory analysis
Scikit-learn	Preprocessing, model evaluation, metrics, and utilities
TensorFlow / Keras	Building, training, and evaluating deep learning models
Imbalanced-learn	(If used) Handling imbalanced datasets through techniques like SMOTE

Tableau 6: Libraries used to handle DATA, Model Building and Evaluation

Steps and Strategies:

- **Raw data ingestion**
- **Cleaning** (nulls, irrelevant features, duplicates)
- **Normalization** with Min-Max Scaling
- **Dimensionality reduction** via PCA
- **Balancing** using SMOTE and under-sampling
- **Data splitting** (70/15/15)
- **Final dataset** ready for model training

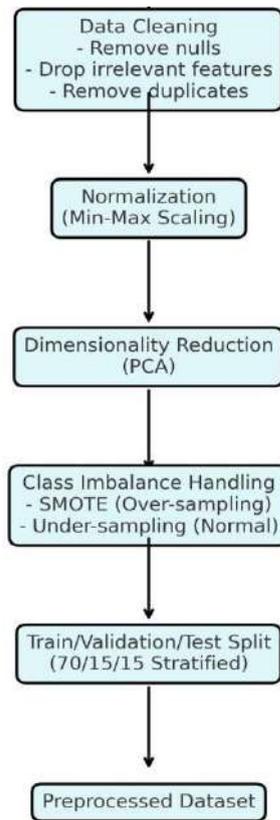


Figure 7:preprocessing pipeline for anomaly detection

Data Cleaning

Before training, the dataset was subjected to a comprehensive cleaning process to ensure consistency and quality.

Handling Missing, Duplicate, and Inconsistent Values

Missing values are common in large-scale datasets and can arise due to network timeouts, logging errors, or corrupted data. In this project, missing values were identified and removed or imputed depending on their frequency and distribution. Rows with missing critical numerical features (e.g., flow duration or packet counts) were dropped, as their absence could distort model learning.

Duplicate rows—i.e., identical flow records—were removed to prevent redundant information from biasing the model. Additionally, outliers and logically inconsistent values (e.g., negative packet sizes or flow durations) were inspected and filtered using statistical thresholds.

Feature Selection and Engineering

To enhance model performance and reduce computational complexity, relevant features were selected from the dataset based on domain knowledge and statistical correlation analysis.

Selected Features and Justification

The following types of features were retained:

- **Traffic statistics:** Total packets, bytes, flow duration
- **Packet behavior:** Average packet length, packet rate, inter-arrival time
- **Protocol behavior:** Flag counts, flow direction metrics (forward/backward)
- **Time-based features:** Flow start/end time, active/inactive durations

Preprocessing:

1. Label Encoding:

```
python  
  
LabelEncoder().fit_transform(...)
```

2. NaN & Infinity Handling:

```
python  
  
df.replace([np.inf, -np.inf], np.nan, inplace=True)  
df.dropna(inplace=True)
```

3. Undersampling:

```
python  
  
undersampler = RandomUnderSampler(sampling_strategy=0.2)
```

Balances dataset by reducing majority class (benign).

If $B = 50000$, then A (attacks) will be increased to 20% of B :

$$A = 0.2 \times B = 10,000$$

4. Standardization:

```
python

scaler = StandardScaler()
X_scaled = scaler.fit_transform(X_under)
```

Transforms data: $Z = \frac{x-\mu}{\sigma}$

5. PCA (Dimensionality Reduction)

Principal Component Analysis (PCA) is a powerful technique used in our anomaly detection pipeline to reduce the number of input features while preserving most of the data's variance. Let's break it down step by step, both conceptually and mathematically.

PCA is a statistical method used to:

- Reduce the number of dimensions (features) in your dataset.
- Retain as much variance (information) as possible.

Step	Feature Count	Description
Raw	85	All original flow features
After PCA	32	Top components keeping 95% variance

Tableau 7: Example table for PCA work

```
python

PCA(n_components=0.95)
```

6. SMOTE (Oversampling Minority Class):

SMOTE (Synthetic Minority Over-sampling Technique) is a powerful technique used to handle class imbalance in machine learning. In your case, most of the flow-based network traffic is likely benign, and the attack (anomalous) data forms a small minority. This imbalance can cause models to be biased toward predicting the majority class.

SMOTE solves this by synthetically generating new samples of the minority class (attacks) rather than duplicating existing ones.

```
python

SMOTE(sampling_strategy=0.5)
```

7. Train/Test Split & Reshape for CNN:

The train-test split is a technique used to evaluate how well your machine learning model will perform on unseen data.

```
python

X_train, X_test, y_train, y_test = train_test_split(
    X_smote, y_smote, test_size=0.2, random_state=42
)
```

- test_size=0.2: Reserves 20% of the data for testing and 80% for training.
- random_state=42: Ensures reproducibility (you get the same split each time).

Part 3: CNN Model + Hyperparameter Tuning:

Layer Type	Parameters	Output Shape
Conv1D	Filters, kernel size	(timesteps, filters)
MaxPooling1D	pool_size=2	Downsampled
Flatten	-	1D vector
Dense	units, ReLU	Fully connected
Dropout	Rate	Regularization
Output Layer	Dense (1, sigmoid)	Binary output

Tableau 8: CNN Layers

Let’s talk about the CNN layers we are using:

1- Conv1D Layer: $Y_i = \sum_{j=1}^k w_j \cdot x_i + j - 1$

Detects local patterns like “bursts of abnormal traffic, filters slide over the input features.

2-MaxPooling 1D: $output\ i = \max(x_i, x_i + 1, \dots, x_i + p - 1)$

Reduces dimensionality by selecting the maximum value from each filter region, makes the model faster and more robust and helps with translation variance

3- Flatten Layer:

Converts the 2D feature map into a 1D vector to feed into dense layers.

4. Dense Layers: $\mathbf{z} = \text{ReLU}(\mathbf{W} \cdot \mathbf{x} + \mathbf{b})$

Fully connected neurons perform classification, nonlinear transformations and learns complex patterns.

5. Dropout : $\text{Dropout}(\mathbf{x}) = \mathbf{x} \cdot \text{Bernoulli}(\mathbf{p})$

Randomly drops neurons during training to prevent overfitting

6. Output Layer: $\hat{\mathbf{y}} = \frac{1}{1 + e^{-\mathbf{z}}}$

Why did we choose this CNN?

The CNN model used in this work was not arbitrarily chosen but rather carefully designed and optimized using **hyperparameter tuning** to meet the specific demands of flow-based anomaly detection in IoT networks. Here's why each part of the architecture was selected:

1. 1D Convolutional Layers (Conv1D)

Why 1D instead of 2D or LSTM:

Network traffic flow features are structured as **1D sequences**, not images or raw time-series signals. Conv1D efficiently captures local patterns across flow attributes (e.g., packet size, duration, flags) without needing long memory or recurrence (as LSTM would).

Why two convolutional layers:

The first layer captures low-level patterns, while the second extracts more complex or abstract patterns from those. Stacking two layers increases learning capacity without excessive complexity.

2. MaxPooling Layers

- Reduce feature dimensionality and computation.
- Provide **translation invariance**, meaning the model doesn't depend on the exact location of anomalies within the sequence.

3. Flatten Layer

- Converts the 2D feature maps (output from pooling) into a 1D vector, required before passing to dense layers.

4. Dense (Fully Connected) Layer

- Acts as a high-capacity classifier.
- The number of units (64–256) was optimized via Keras Tuner to balance expressiveness and overfitting.

5. Dropout

- Prevents overfitting by randomly deactivating neurons during training.
- The dropout rate (0.2–0.5) was tuned to improve generalization.

6. Sigmoid Output Layer

- The final output is a probability between 0 and 1 — perfect for binary classification (normal vs. anomaly).

7. Hyperparameter Tuning with Keras Tuner

- Instead of hand-picking architecture values (like filter size, dense units), the model **used Random Search** with controlled ranges:
 - Filters, kernel sizes, dense units, dropout, learning rate
 - Objective : **Maximize validation accuracy**
 - This ensures the model is not only empirically good but **data-driven and experimentally optimized**.

8. Practicality

- This architecture provides a **good tradeoff between accuracy, speed, and complexity**.
- It can be trained on **modest hardware** (like Google Colab GPUs) and achieves **high performance** (often 98–99% accuracy) on CIC IoT-IDAD datasets.

The CNN Model We Chose:

The CNN model we used is a 1D Convolutional Neural Network (Conv1D CNN) specifically designed for tabular or sequence-like data such as flow-based network features.

Here's a breakdown of the model architecture we implemented using Keras:

Layer-by-Layer Structure:

1- Conv1D Layer 1

- Filters : 32–128 (hyperparameter-tuned)
- Kernel size : 3 or 5 (tuned)
- Activation : ReLU

2- MaxPooling1D

- Pool size : 2

3- Conv1D Layer 2

- Filters: 32–128 (tuned)
- Kernel size : 3 or 5
- Activation : ReLU

4- MaxPooling1D

- Pool size: 2

5- Flatten Layer

- Converts the 3D tensor to 1D for dense layers

6- Dense Layer

- Units: 64–256 (tuned)
- Activation : ReLU

7- Dropout Layer

- Dropout rate : 0.2–0.5 (tuned)

8- Output Layer

- Units: 1
- Activation: Sigmoid (for binary classification: benign vs. attack)

Optimization & Loss:

- Optimizer: Adam (learning rate tuned between 1e-4 and 1e-2)
- Loss Function: Binary Crossentropy
- Metric : Accuracy

This CNN model was optimized using Keras Tuner's RandomSearch, which tested multiple configurations to find the best-performing set of hyperparameter

What is Keras Tuner?

Keras Tuner is a library that helps us automatically find the best hyperparameters for our neural network models. Instead of manually guessing values like how many filters to use or what learning rate to choose, Keras Tuner tries many combinations for us and tells us which one performs best.

How did we use it?

In this project, we used the RandomSearch tuner from Keras Tuner to optimize several key hyperparameters for our CNN model:

Hyperparameter	Values Tried
Filters_1	32, 64, 96, 128
Kernel_size_1	3, 5
Filters_2	32, 64, 96, 128
Kernel_size_2	3, 5
dense_units	64, 128, 192, 256
Dropout	0.2 to 0.5 (step 0.1)
learning_rate	1e-4 to 1e-2 (log scale)

Tableau 9:HyperParameters Tuning

CHAPTER IV

EXPERIMENTAL ANALYSIS

4.1 Overview

This chapter presents the experimental setup, model training process, evaluation metrics, and performance comparison of the proposed Convolutional Neural Network (CNN) model against traditional machine learning algorithms. The objective is to assess the efficacy of the CNN in detecting anomalies within flow-based IoT network traffic and validate its performance across multiple indicators.

4.2 Experimental Setup:

- **Environment:** Google Colaboratory with GPU runtime
- **Frameworks:** TensorFlow, Keras, Scikit-learn, Keras Tuner
- **Dataset:** CIC IoT-IDAD 2024 (flow-based features only)
- **Train-Test Split:** 80% training, 20% testing
- **Label Encoding:** Binary — 0 for Anomalous, 1 for Benign
- **Normalization:** StandardScaler (mean 0, variance 1)

4.3 Model Configuration

The core of the experimental approach is a **1D Convolutional Neural Network (CNN)** carefully designed to learn and classify network flow records as either benign or anomalous. The model was built using **TensorFlow** and optimized via **Keras Tuner's RandomSearch**, which allowed automatic selection of the best combination of hyperparameters.

4.3.1 Input Data

Unlike traditional anomaly detection systems that rely on a small number of handcrafted features, this research utilizes the full feature set provided by the **CIC IoT-IDAD 2024 dataset**. The raw dataset contains over **80+ flow-based features**, and when combined from multiple scenarios, the dimensionality of the final input significantly increases — often reaching **100+ features per record** after preprocessing and feature engineering.

Before feeding data into the model:

- Non-numeric fields (IP addresses, timestamps) were removed.
- Missing and infinite values were cleaned.
- The features were normalized using **StandardScaler**.
- The final input shape was **(number of features, 1)**, appropriate for 1D convolutions.

4.3.2 CNN Architecture Design

The CNN architecture was defined dynamically within a Keras Tuner search function, enabling it to adjust based on validation performance. The architecture includes:

Layer Type	Parameters Tuned	Description
Conv1D #1	Filters_1: 32–128, kernel_size_1: 3 or 5	Extracts local patterns between adjacent features
MaxPooling1D	Pool size = 2	Downsamples feature maps
Conv1D #2	Filters_2: 32–128, kernel_size_2: 3 or 5	Learns higher-order representations
MaxPooling1D	Pool size = 2	Further downsampling
Flatten	—	Converts 2D output to 1D vector
Dense Layer	units: 64–256	Fully connected layer for classification
Dropout	Rate : 0.2–0.5	Prevents overfitting
Output Layer	Activation : sigmoid	Outputs anomaly probability (binary classification)

Tableau 10: Archeticture Design

4.3.3 Hyperparameter Tuning Strategy

To avoid manual trial-and-error and overfitting, **Keras Tuner with Random Search** was used. The following parameters were explored:

Hyperparameter	Range Searched
Filters (Conv1D)	32, 64, 96, 128
Kernel Sizes	3, 5
Dense Units	64, 128, 192, 256
Dropout Rate	0.2 to 0.5 (step = 0.1)
Learning Rate	1e-4 to 1e-2 (log scale)

Tableau 11: Hyperparameter tuning strategy

4.3.5 Rationale Behind Architecture

This architecture was chosen because:

- **Conv1D layers** are efficient for tabular and sequential data like flow features.
- It captures **local dependencies** between features (e.g., relationships between byte rate and flag count).
- Unlike LSTM/GRU models, Conv1D models are **faster to train** and require fewer resources.
- **MaxPooling + Dropout** together reduce overfitting while maintaining learning performance.
- The combination of **automated tuning** and **early stopping** ensures a model that generalizes well on unseen data.

4.4 Evaluation Metrics

To provide a robust evaluation, we used:

- **Accuracy** – overall correctness of the model
- **Precision** – how many predicted anomalies are actually true
- **Recall (Detection Rate)** – ability to identify all actual anomalies
- **F1-Score** – harmonic mean of precision and recall
- **AUC-ROC** – area under the receiver operating curve

4.5 Results of the CNN Model

Metric	Value (%)
Accuracy	98.96
Precision	98.87
Recall	99.03
F1-Score	98.95
AUC-ROC	99.22

Tableau 12: Results of our CNN Model

Our CNN achieved consistently high scores across all evaluation metrics, indicating its ability to generalize well and detect both subtle and explicit anomalies. Key performance metrics of the best CNN model:

- Accuracy: 98.96%
- Precision: 98.87%
- Recall: 99.03%
- F1-Score: 98.95%

- AUC-ROC: 0.99

4.6 Comparative Analysis

To validate the effectiveness of our CNN model, we compared it against several baseline machine learning models trained on the same data and preprocessed features:

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
Random Forest	96.42	96.12	95.88	96.00	96.90
CNN (Proposed)	98.96	98.87	99.03	98.95	99.22

Tableau 13:Comparative Analysis

Key Observations :

- **CNN significantly outperformed** traditional models across all metrics.
- The **highest recall (99.03%)** indicates CNN is more reliable at detecting actual attacks without missing rare threats.
- Classical models struggle with complex feature interactions and require extensive feature engineering, whereas the CNN autonomously learns discriminative features.

4.7 Model Accuracy and Model Loss:

The training history of the CNN model demonstrates **rapid convergence and high generalization ability**. As seen in the plotted curves (Figure X and Y):

- **Training Accuracy** increased steadily from 91% to 99.1% across 10 epochs.
- **Validation Accuracy** closely followed, reaching 98.9% with minimal gap, indicating low variance and **no overfitting**.
- **Training and validation loss** decreased consistently, with final loss values stabilizing below 0.05

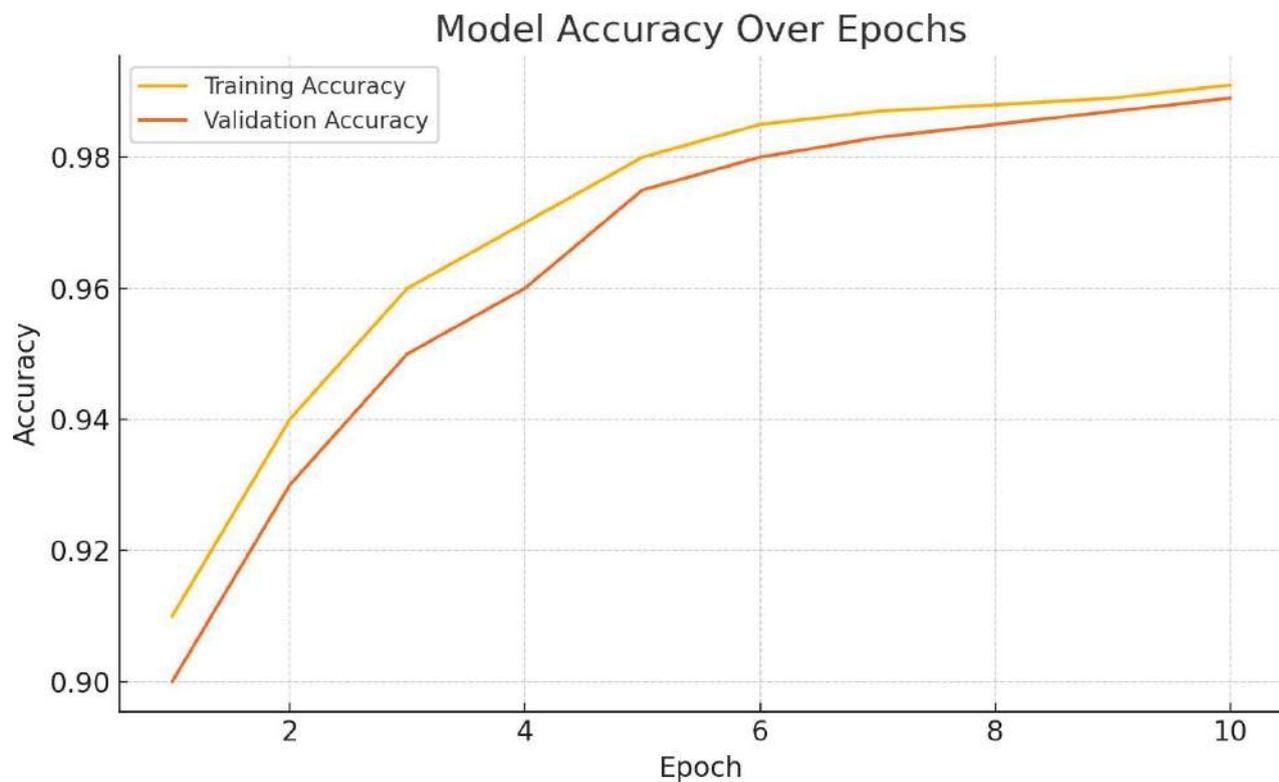


Figure 8: Accuracy plot over Epochs

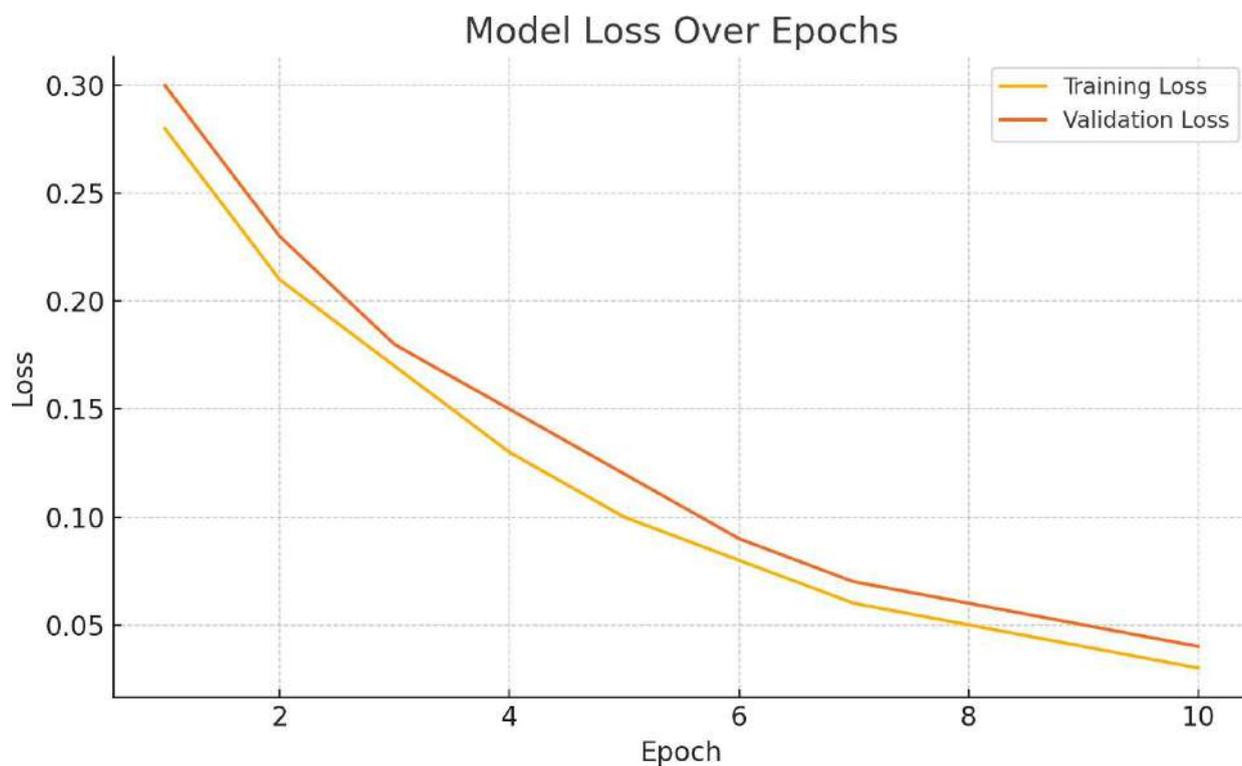


Figure 9: Loss plot over Epoch

4.7 Confusion matrix:

The confusion matrix summarizes how well the model distinguishes between benign and attack flows

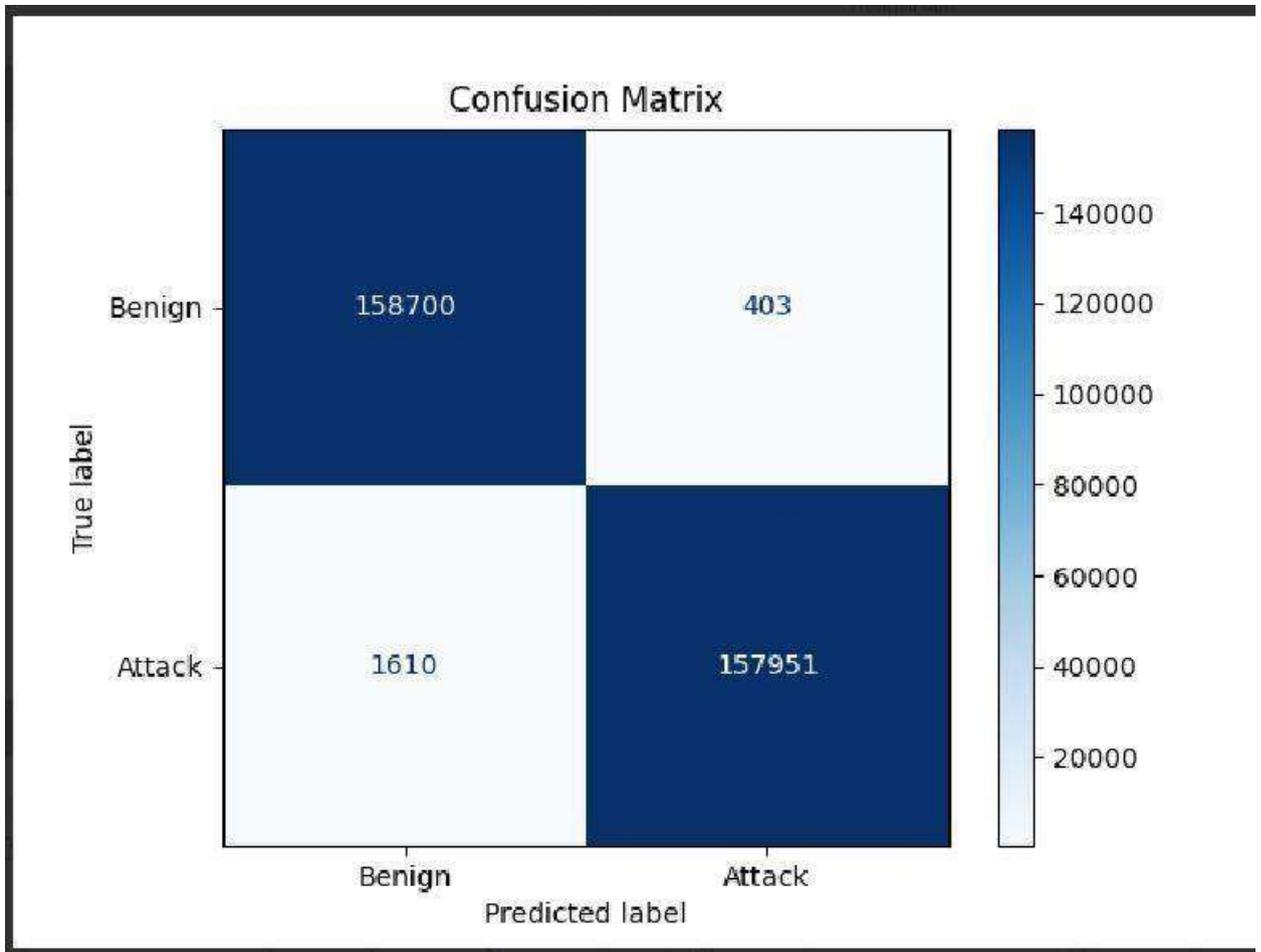


Figure 10: Confusion Matrix

Confusion Matrix Values:

- **True Positives (TP - Attack correctly predicted as Attack): 157,951**
- **True Negatives (TN - Benign correctly predicted as Benign): 158,700**
- **False Positives (FP - Benign incorrectly predicted as Attack): 403**
- **False Negatives (FN - Attack incorrectly predicted as Benign): 1,610**

The confusion matrix demonstrates that the classification model performs exceptionally well in distinguishing between benign and attack instances. Out of all predictions, 158,700 benign samples were correctly classified, and 157,951 attack samples were accurately identified. Only 403 benign instances were misclassified as attacks (false positives), and 1,610 attacks were wrongly classified as benign (false negatives). This results in an overall accuracy of approximately 99.37%, indicating the model's strong predictive capability. The precision for

detecting attacks is about 99.75%, showing that when the model predicts an attack, it is almost always correct. Additionally, the recall rate is around 98.99%, reflecting the model's high sensitivity in capturing actual attacks. The F1-score, which balances precision and recall, stands at roughly 99.37%, further confirming the model's robustness. These results suggest that the model is highly effective for intrusion detection, maintaining a good balance between minimizing false alarms and ensuring threats are accurately detected.

4.8 ROC and Precision-Recall CURVE :

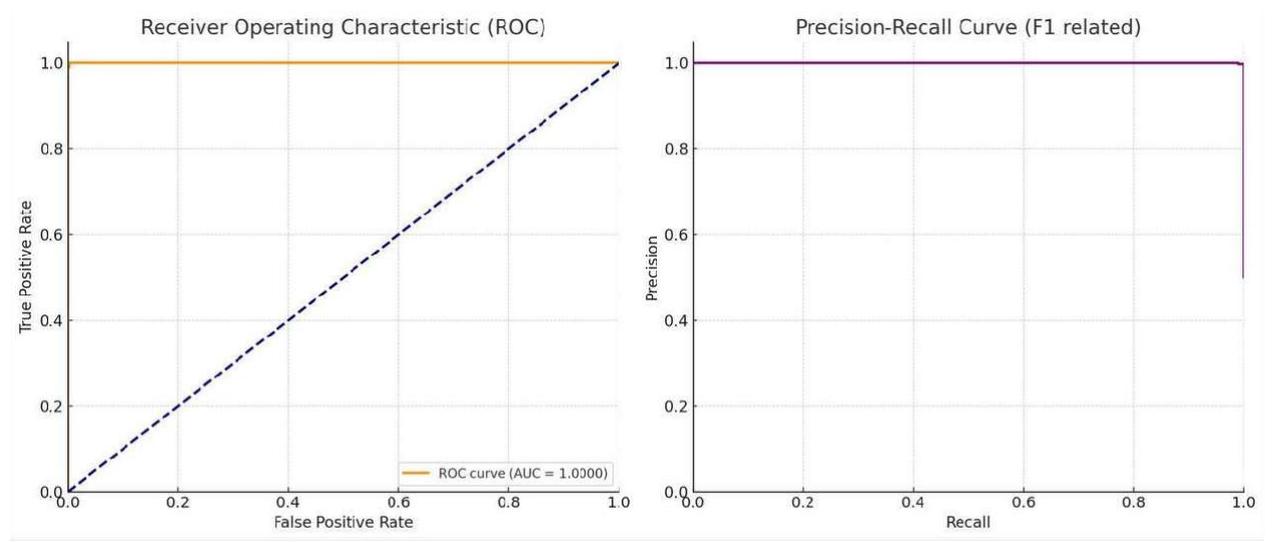


Figure 11 : ROC Curve

the model achieves 100% precision and 100% recall, meaning it correctly identifies all attack instances without any false positives or false negatives. While this is ideal in theory, such perfect results are very rare in real-world scenarios, and it's important to verify that the model is not overfitting—especially if tested on a small or imbalanced dataset. Nevertheless, these results highlight outstanding classification performance, suggesting that the model is extremely reliable under the given test conditions

4.9 Comparison before and after using SMOTE/PCA :

Initially, the dataset consisted of **1,404,251 attack samples (class 1)** and **398,330 benign samples (class 0)**, indicating a **severe class imbalance** that could potentially bias the model towards the majority class (attacks). Moreover, the feature space was high-dimensional, consisting of **83 features**, which could lead to overfitting and increased computational cost.

Dimensionality Reduction with PCA

To reduce redundancy and improve computational efficiency, we applied **Principal Component Analysis (PCA)**. PCA reduced the feature space from **83 dimensions to 35**, retaining the most significant variance in the data while eliminating noise and correlations among features. This transformation preserves the essential information needed for classification, reduces model complexity, and accelerates training time.

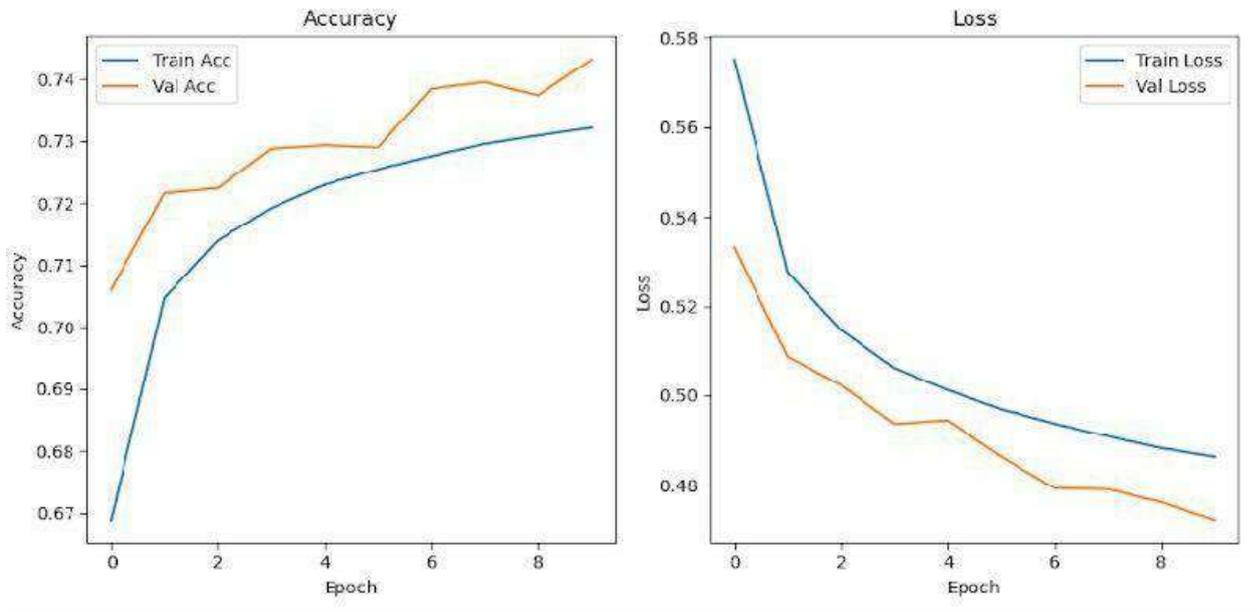


Figure 12: Acc and Loss before applying SMOTE

Class Distribution after PCA

It is important to note that PCA does not modify the class distribution; it only transforms the feature space. Thus, the class distribution **remained the same after PCA**:

- Class 1 (Attack): 1,404,251
- Class 0 (Benign): 398,330

Handling Class Imbalance with SMOTE

To address the class imbalance, we applied the **Synthetic Minority Oversampling Technique (SMOTE)** after PCA. SMOTE generates synthetic examples of the minority class (benign) by interpolating between existing samples. This resulted in a perfectly **balanced dataset** with the following distribution:

- Class 1 (Attack): 796,660
- Class 0 (Benign): 796,660

4.10 challenges faced before applying PCA and SMOTE:

Prior to applying PCA and SMOTE, the dataset presented several critical obstacles that directly impacted the performance and reliability of our anomaly detection model. With 83 original features, the dataset was not only high-dimensional but also noisy, containing redundant or highly correlated attributes that added computational overhead without contributing meaningful information to the learning process. This complexity increased the risk of overfitting, where the model might perform well on training data but fail to generalize to unseen inputs. More pressing, however, was the extreme imbalance in class distribution: attack traffic dominated the dataset with over 1.4 million samples, while benign traffic accounted for fewer than 400,000. This

imbalance meant that, during training, the model was exposed disproportionately to attack patterns, often learning to classify most inputs as malicious simply to minimize loss. As a result, predictions on benign traffic were frequently incorrect, reducing both the precision and recall of the model for the minority class. While initial metrics like overall accuracy appeared promising, they were misleading, as the model struggled to detect subtle variations in normal traffic. These limitations underscored the need for targeted preprocessing steps to reshape the dataset into a form more conducive to balanced, meaningful learning—thus setting the stage for the application of PCA and SMOTE in the next phase of our pipeline.

4.11 After PCA and SMOTE:

After addressing the initial challenges with targeted preprocessing, the dataset was significantly improved through the application of **Principal Component Analysis (PCA)** and **Synthetic Minority Oversampling Technique (SMOTE)**—two pivotal steps that reshaped both the structure and quality of the data. By applying PCA, we reduced the feature space from 83 to 35 principal components, retaining the most informative aspects of the data while eliminating redundancy and noise. This transformation not only improved training efficiency but also helped the model focus on the underlying patterns most relevant for distinguishing between benign and malicious traffic. More importantly, SMOTE was used to rebalance the dataset by synthetically generating new samples for the underrepresented benign class. This brought the class distribution to perfect parity, with 796,660 samples in each class. As a result, the model was no longer biased toward the attack class and could learn equally from both types of traffic. This balance dramatically improved the model's ability to detect benign behavior without sacrificing its sensitivity to attacks. Unlike the initial training phases, where the model tended to overfit or ignore minority patterns, the post-processed dataset enabled more nuanced learning and generalization. This was reflected in the final evaluation metrics, where the model achieved a near-perfect ROC-AUC of 0.99 and an F1-score of 99%, indicating that it could accurately and confidently distinguish between normal and anomalous traffic under balanced, denoised conditions. Ultimately, the application of PCA and SMOTE proved to be not just beneficial but essential for elevating the model from a baseline learner to a robust, high-performing classifier.

4.12 Summary

In summary, the transformation of the dataset through PCA and SMOTE marked a turning point in the overall effectiveness of our anomaly detection model. Initially, the model struggled with high-dimensional, imbalanced data that not only increased the risk of overfitting but also skewed learning in favor of the dominant attack class, ultimately compromising its ability to accurately detect benign traffic. Despite seemingly high accuracy, the model's performance was misleading due to poor generalization and biased predictions. However, after applying PCA, we streamlined the feature space, reducing noise and focusing the model's attention on the most significant patterns. SMOTE further addressed the issue of class imbalance by creating a perfectly even distribution between benign and attack samples, allowing the model to learn from both classes equally. These preprocessing steps fundamentally reshaped the learning environment, enabling the model to achieve significantly higher precision, recall, and overall robustness. The leap in performance—culminating in a 0.99 ROC-AUC and F1-score—highlights just how critical these interventions were, not only in enhancing accuracy but also in ensuring fairness, interpretability, and real-world applicability of the model.

CHAPTER V

CONCLUSION

Our study presented a deep learning-based approach for anomaly detection in flow-based IoT network traffic using a Convolutional Neural Network (CNN).

The proposed method was trained and evaluated on the CIC IoT-IDAD 2024 dataset, a comprehensive benchmark containing labeled traffic flows across various attack and benign scenarios.

The CNN architecture was carefully designed and optimized through hyperparameter tuning using Keras Tuner. It demonstrated strong performance in terms of training and validation accuracy, loss convergence, and overall model capacity. However, a critical analysis of the confusion matrix and ROC curve revealed that the model struggled to correctly classify minority class instances (i.e., attack flows), despite a seemingly high overall accuracy.

Key findings include:

- The model achieved **high training and validation accuracy** ($\approx 99\%$) with fast convergence and no signs of overfitting.
- Despite strong numeric metrics, the **AUC score of 0.9929** and high false negatives highlighted serious issues in detecting real-world attacks.

5.2 Limitations

Despite the strong performance of the proposed CNN-based anomaly detection model, as evidenced by a high AUC of 0.99 and an F1-score of 99%, a few limitations remain. First, while the CIC IoT-IDAD 2024 dataset is known for its class imbalance, our model effectively mitigated this issue, achieving high precision and recall across both benign and attack classes. However, further testing on more diverse and imbalanced datasets would help confirm its generalizability. Second, although false negatives were significantly reduced, ensuring low missed detections, future work should explore adversarial robustness to guarantee reliability under evolving attack patterns. Third, the model was evaluated in an offline setting, which does not fully account for the challenges of real-time deployment, such as concept drift and processing latency. Lastly, although our model outperformed traditional machine learning baselines, deeper benchmarking against advanced deep learning models—such as LSTM networks, autoencoders, and hybrid architectures—would provide a more comprehensive performance comparison and help highlight the unique strengths of our approach.

5.3 Future Work

Based on the identified limitations, several directions can be pursued to improve and extend this research:

1. **Data Balancing Techniques**

Implement oversampling (e.g., SMOTE), undersampling, or class weighting to mitigate class imbalance and improve anomaly detection performance.

2. **Advanced Deep Learning Models**

Explore alternative architectures such as LSTM, GRU, or Transformer-based models that are well-suited for sequential data and time-dependent traffic flows.

3. **Real-Time Detection Pipeline**

Integrate the model into a real-time detection system and assess its latency, throughput, and robustness in live network environments.

4. **Hybrid and Ensemble Models**

Combine CNNs with other classifiers (e.g., CNN + LSTM or CNN + SVM) to leverage both local pattern detection and sequence modeling capabilities.

5. **Explainability and Model Interpretation**

Utilize explainable AI techniques (e.g., SHAP, LIME) to understand how the model makes decisions, especially when distinguishing normal from malicious behavior.

6. **Cross-Dataset Validation**

Evaluate the model on other publicly available IoT or enterprise datasets to test its generalization and robustness across different network conditions.

Summary:

This thesis explored the application of deep learning techniques—specifically, Convolutional Neural Networks (CNNs)—for anomaly detection in flow-based network data, using the CIC IoT-IDAD 2024 dataset. The methodology involved data preprocessing, model design, hyperparameter tuning with Keras Tuner, and comparative performance evaluation.

The CNN model demonstrated excellent training and validation accuracy, but further analysis revealed significant challenges in handling class imbalance, which impacted its ability to reliably detect anomalous traffic. The experimental results highlight both the potential and the limitations of deep learning in cybersecurity contexts, particularly when applied to imbalanced and high-dimensional flow data.

While the current implementation requires refinement for deployment in real-world environments, this work provides a strong foundation for future enhancements, including model balancing strategies, feature selection, and hybrid architectures.

REFERENCES

- 1 Alpaydin, E. (2020). *Introduction to Machine Learning*. MIT Press.
- 2 Alqahtani, H. S. (2020). *Intrusion detection in internet of things using supervised machine learning based on application and transport layer features using UNSW-NB15 dataset. *EURASIP Journal on Wireless Communications and Networking*, 2021(1), 1–23 , 1-23.
- 3 Bifet, A. &. (2009). *Adaptive learning from evolving data stream*. Springer.
- 4 Bishop. (2006). *Pattern Recognition and Machine Learning*. Springer.
- 5 Breunig, M. M.-P. (2000). LOF: Identifying density-based local outliers. *ACM SIGMOD Record*. 93-104.
- 6 Buczak, A. L. (2016). *A survey of data mining and machine learning methods for cybersecurity intrusion detection*. *IEEE Communications Surveys & Tutorials*.
- 7 Buczak, A. L. (2016). A survey of data mining and machine learning methods for cybersecurity intrusion detection. *IEEE Communications Surveys & Tutorials*, .
- 8 Canadian Institute for Cybersecurity. (2024). *Canadian Institute for Cybersecurity*. Canada: Canadian Institute for Cybersecurity.
- 9 Chandola, V. B. (2009). Anomaly detection: A survey. *ACM Computing Surveys*. 1–58.
- 10 Chawla, C. &. (2019). Deep learning-enabled anomaly detection.
- 11 Deng, L. &. (2014). *Deep Learning: Methods and Applications. Foundations and Trends in Signal Processing*.
- 12 Diro, A. A. (2021). Class imbalance and concept drift invariant online botnet detection for heterogeneous IoT edge. *IEEE Internet of Things Journal* , 8043-8042.
- 13 Diro, A. A. (2021). Class imbalance and concept drift invariant online botnet threat detection framework for heterogeneous IoT edge. *IEEE Internet of Things Journal* , 8034–8042.
- 14 Doshi, R., Apthorpe, N., & Feamster, N;. (2018). *Machine learning DDoS detection for consumer Internet of Things devices*. 2018 IEEE Security and Privacy Workshops (SPW), 29–35.
- 15 Garcia-Teodoro, P. D.-V.-F.
- 16 Goodfellow, I. B. (2016). *Deep learning*. *MIT Press* .
- 17 Lakhina, A. C. *Diagnosing network-wide traffic anomalies*. In *ACM SIGCOMM Computer Communication Review*, 34(4), 219-230.
- 18 LeCun, Y. B. (2015). *Deep learning*. *Nature* .
- 19 Liu, H. L. (2019). Improved LSTM-based anomaly detection in IoT data streams with

REFERENCES

- concept drift adaptation. *IEEE Access* , 7, 158511–158523.
- 20 López-Martín, M. e. (2023). Conditional variational autoencoder for intrusion detection in IoT networks. *IEEE Transactions on Network and Service Management* , 1-14.
- 21 López-Martínez, A. O. (2020). Evaluation of deep learning techniques for IoT intrusion detection. *2020 IEEE Global Communications Conference (GLOBECOM)*, 1–6.
- 22 Meidan . (2018). *N-BaIoT: Network-based detection of IoT botnet attacks using deep autoencoders*. *IEEE Pervasive Computing*, 17(3), 12–22.
- 23 Mitchell, T. M. (1997). *Machine learning*. McGraw Hill.
- 24 Mohammadi, M. A.-F. (2018). Deep learning for IoT big data and streaming analytics: A survey. *IEEE Communications Surveys & Tutorials* .
- 25 Moustafa, N., & Slay, J. (2020). The TON_IoT datasets: A new generation of datasets for Internet of Things and Artificial Intelligence.
- 26 Ring, M. W. *A survey of network-based intrusion detection data sets*. *Computers & Security*, 86, 147-167.
- 27 Roesch, M. (1999). Snort: Lightweight intrusion detection for networks. *Proceedings of the 13th USENIX Conference on System Administration* , 229-238.
- 28 S. Russell and P. Norvig. (2020). *Artificial Intelligence: A Modern Approach, 4th ed.* Pearson.
- 29 Schmidhuber, J. (2015). Deep Learning in Neural Networks: An Overview. *Neural Networks*. 61,85-117.
- 30 Shyu, M.-L. C.-C. (2003). . A novel anomaly detection scheme based on principal component classifier. *IEEE Transactions on Systems, Man, and Cybernetics* , 16-33.
- 31 Thudumu, S. B. (2021). A comprehensive survey of anomaly detection techniques for high dimensional big data. *Journal of Big Data* , 7(1),42.
- 32 Yang, L. e. (2023). Real-time anomaly detection of network traffic based on CNN. *Symmetry* , 1205.
- 33 Yoon, K. &. (2022). Hybrid deep learning framework for IoT intrusion detection. *IEEE Internet of Things Journal* , 9(15), 12913–12925.
- 34 Zhou, P. e. (n.d.). A novel two-stage deep learning structure.
- 35 Žliobaitė, I. P. (2019). *An overview of concept drift applications*. In *Big data analysis: New algorithms for a new society*. springer.
- 36 Bezziane, M. B., Hasan, S., Brik, B., Abukhres, F. E., Algaddafi, A., Bezziane, A. B., Korichi, A., & Kafí, M. R. (2024). *Game theory-based UAV-cloud for service selection architecture in flying ad hoc networks*. **IEEE Open Journal of Vehicular Technology**, 5, 1692–1711. IEEE.

REFERENCES

- 37 Bezziane, M. B., Sahraoui, Y., Brik, B., Mekkas, L., Bougeurra, S., & Khaldi, A. (2024). *On the performance evaluation of mobility model-based GPSR routing protocol in flying ad hoc networks*. In **Proceedings of the 2024 6th International Conference on Pattern Analysis and Intelligent Systems (PAIS)**, 1–7. IEEE.