AUTOMATED FRAUD DETECTION IN NAMIBIA'S PUBLIC INSTITUTIONS' FINANCIAL TRANSACTIONS USING MACHINE LEARNING: A DEEP LEARNING APPROACH

A thesis submitted in partial fulfilment of the requirements for the degree of Master of Data Science

In the

Faculty of Computing and Informatics, Department of Informatics

At the

Namibia University of Science and Technology

By

Johannes Pandeni Paavo

223128481

Date of Submission:

December 2024

Supervisor: Dr Richard Maliwatu

Co-Supervisor: Prof. Rafael Rodríguez-Puentes

METADATA

TITLE:	Mr
STUDENT NAME:	Johannes Pandeni Paavo
SUPERVISOR:	Dr Richard Maliwatu
CO-SUPERVISOR:	Prof. Rafael Rodríguez-Puentes
DEPARTMENT:	Informatics, Journalism and Media Studies
QUALIFICATION:	Master of Data Science
SPECIALISATION:	Data Science
STUDY TITLE:	Automated Fraud Detection in Namibia's Public Institutions Financial Transactions Using Machine Learning: A Deep Learning Approach
KNOWLEDGE AREA:	Data Science
KEYWORDS:	deep learning, financial fraud, financial transactions, fraud detection, machine learning
TYPE OF RESEARCH:	Applied Research
METHODOLOGY:	Mixed Methods Research
STATUS:	Final Thesis
SITE:	Windhoek Campus
DOCUMENT DATE:	16 April 2025
SPONSOR:	NONE

ABSTRACT

Financial fraud continues to be a significant concern in public-sector financial operations, undermining the credibility of financial statements and eroding public trust. Traditional methods used by financial experts, such as auditing, are frequently ineffective in addressing the growing complexity of fraudulent activities and effectively mitigating associated risks. This study aimed to tackle this issue by creating an automated fraud detection system based on deep learning designed for Namibia's public sector financial transactions. The Ministry of Finance provided the primary data for the study through the Office of the Auditor-General, which included accounts payable records from public entities with large transaction volumes for the fiscal years 2021/2022 and 2022/2023. The task of fraud detection is framed as a classification problem. The study explored three common deep learning models: Autoencoders, Generative Adversarial Networks (GAN) and Convolutional Neural Networks (CNN). These models' performance was evaluated using historical and simulated financial data, focusing on accuracy, inference time, and resource utilisation. A comparative analysis revealed that the CNN model performed exceptionally well, with the highest accuracy (0.95), F1-score (0.98), and lowest false positive rate (0.038). In contrast, the GAN model excelled in inference time (7.17 ms per transaction) and precision (0.99). This study proposes a scalable, data-driven approach to improving fraud detection in large public-sector financial datasets, thereby increasing accountability in Namibia's public financial systems.

Keywords: deep learning, financial fraud, financial transactions, fraud detection, machine learning

DECLARATION

I, Johannes Pandeni Paavo, hereby declare that the work contained in the mini-thesis, entitled Automated Fraud Detection in Namibia's Public Institutions Financial Transactions Using Machine Learning: A Deep Learning Approach, is my original work and that I have not previously submitted it in its entirety or in part to any university or other higher education institution for the award of a degree.

Signature:  ..... Date: 16 April 2025

RETENTION AND USE OF THESIS

I, Johannes Pandeni Paavo, a candidate for the degree of Master of Data Science, accept the Namibia University of Science and Technology's requirements regarding the retention and use of theses/mini-theses deposited in the Library and Information Services.

In terms of these conditions, I agree that the original of my thesis/mini-thesis deposited in the Library and Information Services will be accessible for purposes of study and research, in accordance with the normal conditions established by the Librarian for the care, loan or reproduction of theses/mini-theses.

Signature:  Date: 16 April 2025

DEDICATION

This thesis is dedicated to the young boy I once was, the boy whose burning curiosity still inspires me today. Your aspirations for greatness fuelled this journey, and despite the many obstacles, I WILL NOT FAIL YOU.

ACKNOWLEDGEMENTS

I want to thank God for the guidance and strength that have carried me through this academic journey.

My sincere appreciation goes to my dedicated supervisor, Dr Richard Maliwatu, for his unwavering support, constructive feedback, and invaluable insights. I am also deeply indebted to my co-supervisor, Prof. Rafael Rodríguez-Puentes, whose continued mentorship throughout this research and previous projects has been invaluable. His consistent support and encouragement have been instrumental to my academic development.

I would also like to express my heartfelt gratitude to the financial sector experts from the Office of the Auditor-General Namibia for their invaluable insights and participation in this research. Your contributions helped illuminate the subject.

A special thank you to my boy Jacob Tangeni Paavo for being there, even if only in spirit. Your company, even while doing nothing, meant more than you know!

Looking back on this experience, I am humbled and grateful for the opportunity to grow.

PUBLICATIONS

Paavo, J. P., Rodríguez-Puentes, R., & Maliwatu, R. (2024). Exploring machine learning fraud detection solutions for financial transactions. International Engineering Conference on Sustainable Emerging Innovations and Technological Advancements.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABBREVIATION/ACRONYMS

ACFE	Association of Certified Fraud Examiners
ANN	Artificial Neural Networks
AP	Accounts Payable
CNN	Convolutional Neural Network
GAN	Generative Adversarial Networks
GL	General Ledger
GPU	Graphics Processing Unit
IFMS	Integrated Financial Management System
ML	Machine Learning
NDP 5	Fifth National Development Plan
RAM	Random-access Memory
ReLU	Rectified Linear Unit
ROC-AUC	Receiver Operating Characteristics – Area Under Curve
SDG	Sustainable Development Goals
SVM	Support Vector Machines
XGBoost	Extreme Gradient Boosting

CHAPTER 1 : INTRODUCTION

This chapter provides a brief background and introduction to the research topic, outlining the problem statement, research objectives, aim, and research questions. It also discusses the study's significance, limitations, and delimitations.

1.1 Research Background

Financial fraud occurs when fraudsters use deceitful and unlawful techniques to obtain financial benefits, manifesting across diverse finance domains such as banking, insurance, taxation, government, corporations, and others (Ashtiani & Raahemi, 2021). The Association of Certified Fraud Examiners (ACFE) categorises occupational fraud into three primary categories: corruption, financial statement fraud and asset misappropriation. According to the ACFE, around 10% of instances of white-collar crime entail the fabrication of financial statements. Although financial statement fraud is relatively rare in the corporate sector, it is the most financially detrimental, leading to a median loss of US$954,000 per case (Bloomenthal, 2021). Based on global data from the ACFE (2020), this figure highlights significant financial risks, though the impact on Namibian public institutions may vary depending on local financial reporting and auditing practices.

Ashtiani and Raahemi (2021) reported a significant increase in fraud cases recently, specifically in public-sector financial operations. This highlights the growing importance of fraud detection, given the significant threat to the stability and transparency of such operations. Byrd and Guimbert (2009) emphasise that the inappropriate allocation of public resources, unorthodox transactions, and deceptive conduct can have pervasive consequences, eroding the public's confidence and undermining the credibility of government institutions.

The effective management of public funds is essential for advancing the socioeconomic development and welfare of the country’s citizens (Noja et al., 2019). However, the complexity and sheer volume of financial transactions within institutions often pose challenges for traditional methods such as manual detection. Nonetheless, artificial intelligence, particularly machine learning (ML) technologies, has

proven highly successful in fraud detection (Sheshasayee & Thomas, 2017). Furthermore, as the digital era has progressed, more complex fraudulent schemes have emerged, necessitating advanced detection and prevention tools, as observed by Sheshasayee and Thomas (2017). In this research, it is critical to note that actual fraud and errors are categorised together as anomalies that could signal irregular or questionable transactions. An erroneous transaction may arise from system glitches, procedural failures, or data entry mistakes rather than deliberate deception. Further verification is typically required to determine whether a detected anomaly is a fraud or an error (Rivas, 2023).

1.2 Problem statement

Financial fraud remains an ongoing and persistent challenge for public institutions in Namibia. It presents significant risks to the integrity and transparency of financial statements and a loss of public trust. Conventional approaches employed by financial experts, such as auditors, often prove inadequate to address the escalating complexity of fraudulent activities and mitigate these risks.

1.3 Research objectives, aims and questions

To effectively address the research problem, this thesis developed a set of research objectives, aims and questions that served as the guiding framework. They established specific goals and purposes that guided the research process and facilitated the definition of what the study aimed to accomplish or investigate.

1.3.1 Research objectives

Main objective:

- Develop a deep learning-based automated system to detect fraud in Namibia's public sector financial transactions.

Sub-objectives:

I. Identify relevant financial features and indicators for fraud detection.

II. Investigate resource utilisation challenges and opportunities for the system handling a large volume of public institutions' financial transactions.

III. Evaluate the developed system's fraud detection performance using historical and simulated data.

1.3.2 Research aim

This research aimed to automate fraud detection in financial transactions using advanced deep-learning techniques tailored explicitly for Namibia's public sector financial transactions. This system will help to detect fraudulent activities in public institutions' financial records, thereby enhancing accuracy, speed, and real-time fraud prevention.

1.3.3 Research questions

The following questions were designed to help guide the study.

Main question:

- What deep learning architectures and algorithms are best suited for building an automated fraud detection system for Namibia's public sector financial transactions?

Sub-questions:

I. Are there unique financial indicators specific to Namibia's public sector that should be considered when detecting fraud?

II. What are the resource utilisation challenges and opportunities for a system that processes a high volume of financial transactions in public institutions?

III. How can the system's fraud detection performance be effectively assessed using historical and simulated data to ensure reliability?

1.4 Significance of the study

The significance of this study stems from its approach to addressing Namibia's pressing issue of financial fraud in the public sector. The study developed an automated financial fraud detection system that enhances efficiency in processing large-scale financial transactions, improving financial accountability

and integrity in the government's monetary operations. This system can help auditors and stakeholders identify anomalies faster and more accurately, ultimately restoring public trust in government institutions, aligning with the NDP 5 pillar of good governance and the SDG of promoting peace, justice, and strong institutions.

1.5 Limitations / Delimitations

Future research has the potential to address two significant limitations of this study. Firstly, financial data from public institutions is subject to inherent limitations related to occurrence, accuracy, and completeness, which can impact the robustness and reliability of any model built upon such historical data. Furthermore, simulated fraudulent transactions might not accurately represent the intricacies and variety observed in practical financial transactions and fraud patterns.

On the delimitations, the study exclusively utilised publicly available information and worked with anonymised historical data to address privacy and ethical concerns. While stakeholders were engaged, not all may have been represented. Furthermore, the financial reporting standards and regulatory frameworks were built on existing structures, which may not be able to accommodate future changes that affect the relevance and effectiveness of the developed model.

1.6 Summary of research findings

The CNN outperformed all other models, with the uppermost accuracy (0.95) and F1-score (0.98), indicating that it is the best model. While GAN provided higher precision (0.99) and faster inference times (7.17 ms/transaction), it required more system memory (7.2 GB) and had a higher false negative rate (0.673). Although the Autoencoder is useful for unsupervised tasks, it achieved lower accuracy (0.72) and faced challenges in handling false negatives and false positives, thus making it ineffective for this system.

The study used a Random Forest model to identify key financial indicators unique to Namibia's public sector. The most important features were discrepancies in General Ledger (GL) codes for the same

vendor, duplicate invoices, and high amounts. Unusual purchase patterns, inconsistencies in bank accounts, and large transaction amounts were among the less notable features.

The models' fraud detection performance was assessed using historical financial data from Namibia's Ministry of Finance and simulated data to ensure robustness. When evaluated on historical data, CNN outperformed other models regarding accuracy, F1 score, recall and precision. Simulated fraud patterns confirmed the system's reliability in various scenarios, though further model refinement may be required to optimise the performance of Autoencoder and GAN models.

Several methods were used to manage the large volume of transactions to address computational efficiency challenges. CNN demonstrated superior memory efficiency (1.94 GB RAM) when processing large datasets (430,738 rows in 2.16 seconds), which makes it ideal for real-time fraud detection systems. While faster in inference, GAN required significantly more system memory (7.2 GB), whereas the Autoencoder model was moderate in terms of memory usage and inference time.

1.7 Organisation of the dissertation

The next five chapters of this thesis are structured as follows:

Chapter Two: Literature Review explores various fraud detection techniques in the finance sector, examining the limitations of traditional approaches. It also investigates machine learning solutions for financial fraud detection, such as deep learning techniques. Furthermore, the chapter offers perspectives on future research directions and practical applications.

Chapter Three: Research Methodology provides an overview of the research procedures and methodologies used to achieve the study's objectives.

Chapter Four: Design and Development examines the methodologies used in the design and development of the automated fraud detection system, which was the primary focus of the research project.

Chapter Five: Results and Discussion presents the findings and discussions on the automated fraud detection system deep learning architectures and algorithms best suited for Namibia's public institutions' financial transactions. Specific financial indicators considered, resource utilisation challenges and opportunities, metrics used to assess accuracy, false positive rates, and system performance were discussed.

Chapter 6: Conclusions and Future Work summarises vital findings, contributions to the field, limitations, and future research directions before making conclusive assertions.

CHAPTER 2 : LITERATURE REVIEW

This section provides a systematic literature review, summarising existing knowledge and research on the subject. It also discusses various ML algorithms for fraud detection solutions used to improve the effectiveness of financial statement auditing, such as Deep Learning, a subset within the field of ML. Lastly, the section looks at the implications of ML and provides insights for future research in this domain.

2.1 Introduction

Following the Enron scandal, the largest accounting fraud in history, there has been a greater emphasis on preparing and presenting corporate financial statements containing fraudulent activities, and raising concerns about the reliability of financial reports (Stalebrink & Sacco, 2007). The American Institute of Certified Public Accountants (AICPA) explains fraud in SAS No. 99 as a deliberate action that leads to a significant misrepresentation of financial statements. The AICPA's definition distinguishes between two types of misstatements (AICPA, 2007). Misstatements are caused by deceptive financial reporting, such as accounting record manipulation, and errors related to asset misappropriation, including asset theft. Although manual auditing of financial statements has traditionally been the primary method for detecting fraud, its effectiveness is inherently limited, particularly considering the growing complexity and volume of financial data. As a result, intelligent fraud detection systems are constantly being developed to detect possible indicators of fraudulent financial transactions, thereby providing valuable insights to help stakeholders make informed and timely decisions (Sanad & Al-Sartawi, 2021). These intelligent systems utilise machine learning, which, as defined by Faraji (2022) and Bhavitha et al. (2017), refers to analytical methods capable of detecting patterns in data autonomously without requiring expert manual instruction.

2.2 Financial fraud detection techniques

Identifying occurrences of financial fraud presents a significant and challenging task, as acknowledged by Yao et al. (2018). This can be attributed to using refined analytical tools and techniques required to

detect evolving and complex financial fraudulent patterns, often missed by the traditional methods commonly used by experts like auditors, as argued by Hamal and Senvar (2021). These conventional methods to detect fraud in financial transactions typically include cluster analysis, regression analysis, factor analysis, and discriminant analysis. It is worth noting that, among the numerous challenges, traditional statistical methods are limited by the need to adhere to specific assumptions in the data, such as normality, linearity and independence of variables (Chen et al., 2014). It is also important to mention that there is no universally applicable method for identifying financial transaction fraud, as stated in the literature. However, a meta-analysis of research articles suggests that models driven by machine learning are more effective than traditional methods in financial fraud detection (Gupta & Mehta, 2021). Craja et al. (2020) emphasise that these models can be trained on unseen data, thus enabling them to adapt and detect financial fraud as patterns evolve. Additionally, Gupta and Mehta (2021) present empirical data demonstrating the enhanced precision of machine learning-based detection models compared to conventional methodologies. In contrast, advanced methodologies such as deep learning, a specialised subgroup of machine learning, are employed by researchers to identify instances of fraudulent financial transactions in the era of artificial intelligence (Craja et al., 2020). Craja et al. (2020) assert that deep learning, although built upon neural networks, possesses more layers and is adept at capturing features and addressing complex problems. Moreover, deep learning outperforms traditional machine learning algorithms in efficiency and predictability, especially when handling large volumes of data (LeCun et al., 2015).

2.3 Conventional/Traditional Fraud Detection Methods Limitations

According to West and Bhattacharya (2016), manual fraud detection methods are inefficient, costly, unreliable, and unsuitable in the era of big data. As a result, financial institutions have turned to automated systems that leverage computational and statistical techniques. Besides individual acts of fraud, the other challenge these institutions face is the presence of organised crime rings, where groups collaborate to commit financial fraud (Sudjian et al., 2010). According to Sudjian et al. (2010), individual transactions can easily evade traditional detection systems because they appear legal or involve small

amounts of money. Nonetheless, when viewed as part of a larger pattern of activity, which often involves multiple individuals, the criminal nature becomes more apparent. Thiprungsri and Vasarhelyi (2011) share these sentiments and claim that traditional statistical methods, such as cluster analysis, struggle to detect these intricate fraud patterns, which frequently involve sophisticated techniques. These methods typically rely on superficial linear relationships or normal distribution assumptions, which fail to capture the nuanced behaviours associated with fraudulent activities. Thiprungsri and Vasarhelyi advise auditors to consider new and innovative audit approaches to address these challenges effectively.

Sudjian et al. (2010) also pointed out that the complexity of financial data presents a significant challenge in detecting financial fraud. Financial transaction records include multiple variables such as transaction amount, timing, and user activities such as transaction initiations and approvals, resulting in high-dimensional datasets. They explain that traditional methods, such as factor analysis, may have difficulties effectively managing such complex data, thus leading to lower accuracy and interpretability. According to Makki et al. (2019), traditional statistical methods used in fraud detection, such as logistic regression and linear analysis, excel at classifying transactions as fraudulent or legitimate based on predefined labels. However, unsupervised approaches such as cluster analysis have an advantage in real-world applications because they identify unusual or anomalous patterns in data, making them effective in detecting previously unknown fraudulent activities. This strength is beneficial as fraudsters constantly refine their strategies to exploit detection measures, thus necessitating continuous model retraining and adaptation (Hilal et al., 2022).

2.4 Machine learning financial fraud detection solutions

2.4.1 K-means

Recent studies, including that of Huang et al. (2024), have focused on machine learning methods for overcoming the limitations of traditional fraud detection techniques. K-means clustering has emerged as a notable technique for improving financial fraud detection. This algorithm groups transaction data based on shared characteristics such as amount, frequency, and location, thereby assisting in detecting

suspicious activities (Huang et al., 2024). Variants such as weighted and fuzzy K-means have also been investigated to address issues associated with imbalanced datasets and complex fraud patterns. Wang et al. (2018) supported this approach with a simulation experiment that used Hidden Markov Models (HMM) and K-means methods to detect bank fraud. Huang et al. (2024) also assert that K-means outperforms traditional rule-based techniques by adapting to evolving fraud tactics, thereby increasing detection accuracy and flexibility. Deng and Mei (2009) present an additional K-means solution for fraud detection that combines text dimensionality reduction with document clustering. Their dual Growing Hierarchical Self-Organising Map (GHSOM) technique accurately identifies non-fraud-central spatial patterns, revealing topological structures of fraudulent transactions. Deng and Mei (2009) address the uncertainty of node clustering borders by combining K-means clustering with SOM, which strengthens the robustness of their clustering-based fraud detection method. Zeng et al.'s (2024) experimental results confirm the promise of K-means clustering in financial fraud detection, particularly in high-risk areas, by identifying feature differences between clusters and detecting potential fraud cases. Despite the need for careful initialisation and optimal cluster determination, K-means is still helpful because of its simplicity, scalability, and efficiency. Its effectiveness in real-time applications makes it ideal for managing large datasets and assisting financial institutions with targeted fraud prevention (Huang et al., 2024).

2.4.2 Support Vector Machines

Support Vector Machines (SVM) is another popular ML algorithm for financial fraud detection. They are a supervised technique that seeks to identify the hyperplane that offers the largest margin to divide input training data into two distinct categories (MDPI Books, 2023), thereby effectively distinguishing fraudulent transactions from legitimate ones in a high-dimensional space. This ability to detect patterns in large and complex datasets demonstrates SVM's effectiveness in fraud discovery. Sulaiman et al. (2022) in their credit card fraud detection analysis found that while SVM is effective and accurate when dealing with a smaller subset of features, it becomes difficult to handle larger datasets with more than 100,000 entries. In response, they proposed a hybrid solution that combines Artificial Neural Networks

(ANN) with a joint learning framework. SVM's ability to manage non-linear relationships inherent in financial transactions, aided by kernel functions, improves its accuracy in detecting fraudulent activities (Singh & Jain, 2020). Similarly, Rajak and Mathai (2015) developed a hybrid smart detection system for detecting fraud in credit card transaction processing that combines SVM and Fusion Danger theory to achieve better F-measures and time complexity performance. Health insurance has also not been immune to fraud, which involves deliberate deception or misrepresentation to obtain benefits illegally, usually disguised as healthcare expenses. Kirlidog and Asuk (2012) used SVM to develop an automated medical bill architecture, demonstrating improved real-time medical fraud detection. Equally, Wang and Xu (2018) demonstrated that optimised SVM outperformed alternative models in spotting fraudulent events in online credit card transactions using datasets from commercial banks. Mareeswari and Gunasekaran (2016) used SVM and an existing spike detection algorithm to detect fraudulent credit card patterns in a banking system, overcoming inherent methodological limitations such as scalability issues, extremely imbalanced classes, and time constraints. In a similar case, a method combining One-Class Support Vector Machine (OSVM) and deep learning showed promise in detecting credit card fraud (Jeragh & AlSulaimi, 2018). Sundarkumar et al. (2015) have also improved fraud claims detection in the insurance industry using an OSVM-based undersampling technique. The SVM's resistance to overfitting, especially in high-dimensional financial data, ensures accurate and reliable fraud detection.

2.4.3 Decision trees

Another fraud detection ML algorithm is decision trees, which, according to HaratiNik et al. (2012), are an effective machine learning tool for developing decision support systems. These trees are made up of internal nodes representing binary choices based on features, taking advantage of their intuitive structure and ability to handle complex datasets. Decision tree-based methods have been utilised to detect financial fraud for many years. In 2017, Devi and Kavitha introduced a decision tree approach to categorise credit card transactions as either usual or doubtful. In this case, a decision tree divides transaction data recursively into branches based on feature values like transaction amount, frequency,

and user behaviour. Each split is intended to maximise the separation of different classes, namely fraudulent and non-fraudulent transactions. The method was evaluated using various correctness metrics, and the results demonstrated that the decision tree method surpassed current methods, thereby achieving a notably high level of accuracy.

Decision tree algorithms are effective because they provide human-readable classification rules, allowing analysts to easily understand the decision-making process and identify which factors are most important for fraud detection (Ali et al., 2012). Capable of processing both numerical and categorical data, decision trees are suited to analysing a wide range of financial dataset types. However, Ali et al. (2012) noted that building decision trees can be computationally costly and time-consuming, particularly for large datasets. They are also susceptible to overfitting, where the model captures noise in the training data, resulting in suboptimal decision-making. To address these issues, Breiman (2001) proposed Random Forests, an ensemble technique that integrates several models to create a powerful predictive system. Breiman (2001) concluded that random forests reduce overfitting by constructing multiple individual decision trees from randomly sampled data with replacement, increasing accuracy and efficiency compared to other classifiers such as Discriminant Analysis and Support Vector Machines.

2.4.4 Naive Bayes

In addition to ensemble methods such as Random Forests, Naive Bayes provides another option for detecting financial fraud. Based on Bayes' theorem, Naive Bayes is a probabilistic classifier that predicts membership probabilities for each class in a dataset (Deng, 2010). This model calculates the probability of a data point being assigned to a particular category. The Naive Bayes (NB) model is widely used in financial fraud detection due to its ease of use and efficiency in dealing with large datasets. For instance, a fraud detection system for financial transactions was developed by Deng (2010) using the NB model, and it was tested on a dataset comprising both usual and irregular financial statements. The findings showed that the model could quickly categorise transactions as fraudulent or non-fraudulent by analysing feature probabilities. Similarly, Peng and You (2016) used the Naïve Bayes algorithm to detect

fraudulent transactions in the healthcare industry by analysing medical procedure records and classifying supplier behaviour as normal or abnormal. Naive Bayes' ability to easily update with new data ensures that the model remains relevant and effective as fraud patterns change. Despite its often unrealistic assumption of feature independence in real-world scenarios, research by Gupta et al. (2021), Deng (2020) and Al-Hashedi and Magalingam (2021) has shown that Naive Bayes performs well in practice. Its speed and accuracy make it an important tool for real-time fraud detection, allowing financial institutions to identify and address potentially fraudulent activities quickly.

2.4.5 Deep Learning

Despite the benefits of traditional machine learning algorithms, many researchers are turning to deep learning methods to detect financial fraud. This shift is motivated by deep learning's ability to handle larger datasets efficiently, identify previously unknown patterns, and provide superior computational speed. In their study, Alghofaili et al. (2020) suggested a deep learning approach to fraud detection using Long Short-Term Memory (LSTM) techniques. Their objective was to enhance existing detection techniques and optimise precision, particularly in big data. The researchers performed experiments and compared their results to a previously developed deep learning model such as the Auto-encoder model, and several other machine learning methods. The experimental outcomes demonstrated LSTM's remarkable performance, with an impressive 99.95% accuracy in less than a minute.

Another notable solution in the context of financial reports for goods exports in Brazil involved using an unsupervised deep learning model. This model effectively identified irregularities in at least twenty exporters' financial records (Paula et al., 2016). In a separate study, Bakumenko and Elragal (2022) analysed seven supervised machine learning methods, including deep learning and two unsupervised methodologies. According to their findings, top-performing models in both the supervised and unsupervised categories showed significant potential in detecting specified anomaly categories and effectively selecting data to detect high-risk journal entries. Despite the advantages of deep learning, researchers underscored the substantial data prerequisites linked to this method, which demand a substantial quantity of labelled data. It is crucial to recognise that the responsiveness of these models

to data quality can amplify biases, resulting in predictions that are difficult to interpret due to their "black box" nature.

2.5 Overview of ML - Deep Learning

Machine Learning has made significant advancements in data processing and classification across multiple domains in the past decade, enabling the development of intelligent systems (Youness et al., 2018). These systems' effectiveness depends on the logical and sequential integrity of the data and the timeliness of the feedback they generate. Researchers have proposed numerous approaches and algorithms to address this challenge, particularly those based on deep learning (Bini, 2018).

According to Yoshua et al. (2017), deep learning is a subgroup of machine learning that analyses and models complex data patterns using multi-layered neural networks. Deep learning, which uses data and experience, allows systems to improve their performance over time, making it ideal for anomaly detection, natural language processing, and image recognition applications. Moreover, deep learning stands out for its extraordinary ability to build a dynamic hierarchy of concepts and representations. Each concept within this hierarchy is defined as more fundamental ideas, with increasingly abstract representations deriving from less abstract ones. In essence, as concisely summarised by LeCun et al. (2015), deep learning enables computational models to acquire data representations at different levels of abstraction by utilising multiple processing layers. As a result, it enables computers to address perceptual challenges by using deep artificial neural networks, which utilize several processing layers to recognise complex patterns and structures within large datasets (Rusk, 2016).

2.6 Research gaps

Despite significant progress in financial fraud detection, several deficiencies persist, thereby limiting the effectiveness of current methodologies. A significant challenge is managing imbalances where fraudulent transactions are far less common than genuine ones. Despite the introduction of techniques such as fuzzy K-means, there is a need for more resilient models capable of detecting complex and evolving fraud patterns. This study addresses this gap using deep learning models, specifically

unsupervised methods like Autoencoders, which show improved adaptability in detecting infrequent fraudulent occurrences. Scalability and efficiency are additional constraints for traditional machine learning algorithms such as SVMs and Decision Trees, which frequently struggle with large datasets, resulting in increased computational demands. The present study aimed to enhance fraud detection systems by leveraging the scalability of deep learning architectures, such as Autoencoders and CNNs while optimizing inference times and reducing memory consumption.

Deep learning models, often referred to as "black box" systems, lack interpretability. This lack of transparency poses challenges in finance, where it is essential to explain model decisions in order to build trust. The present study employed interpretability techniques to shed light on the decision making processes of these models, thereby increasing transparency and their applicability in real-world scenarios. Furthermore, while useful for real-time fraud detection, traditional techniques such as Naive Bayes frequently rely on implausible assumptions such as feature independence in complex financial transactions. This study addresses these limitations by utilising deep learning models capable of detecting complex interdependencies between features and adapting to new fraud tactics. By filling existing gaps, the study improves the scalability, interpretability, and efficacy of automated fraud detection solutions, particularly in Namibia's public sector.

2.7 Chapter summary

The financial statement fraud detection landscape highlights a critical need for effective detection methods after high-profile scandals such as the Enron case. Because financial data is becoming more complex, the shift from traditional auditing methods to intelligent systems emphasises the importance of machine learning, particularly deep learning, which was used in the study to achieve its primary goal of developing a deep learning-based automated fraud detection system for Namibia's public sector financial. Deep learning's success in perceptual tasks can be mainly attributed to the automation of the feature engineering process, as well as its supervised models, such as LSTM techniques and unsupervised models, such as Autoencoders, which have demonstrated exceptional accuracy and efficiency in detecting financial fraud, despite data prerequisites and interpretability concerns.

Additionally, this study aimed to address critical research gaps, such as the handling of imbalanced datasets, the scalability of machine learning methods for large datasets, and the interpretability challenges of deep learning models, all critical for real-time fraud detection and fostering trust among financial stakeholders. These advancements underscore the evolving nature of financial accountability and the essential role of technology in ensuring the integrity of financial markets and enhancing trust among stakeholders.

CHAPTER 3 : RESEARCH METHODOLOGY

This chapter describes the research procedures and methodologies utilised to conduct the research, guided by the reviewed literature. It provides a clear framework for the research design, data collection, analysis, and interpretation to address the research objectives effectively. The methodology is based on Saunders et al.'s (2023) research onion, which outlines a comprehensive, staged approach to developing robust research strategies. The critical phases of this framework, demonstrated in Figure 3.1, guided the development of the methodological approach used in this study.

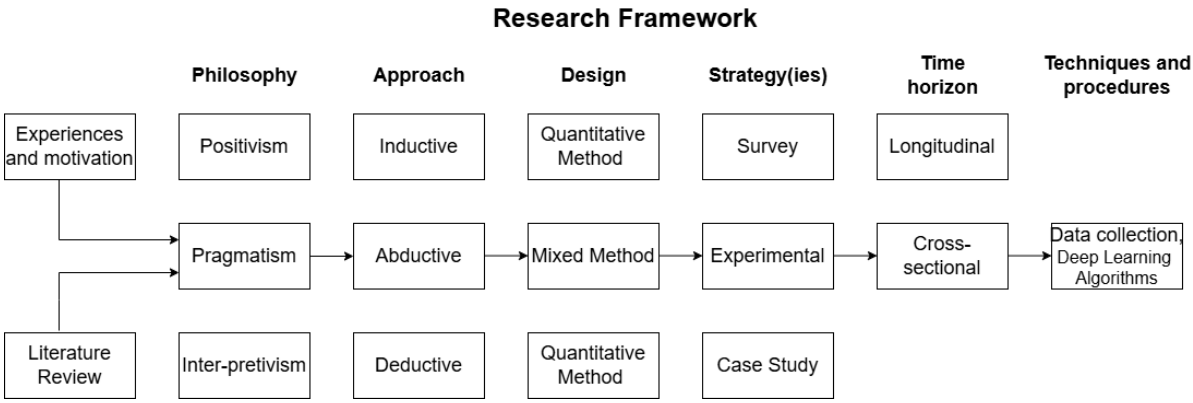

Figure 3.1 The Research Framework (Adapted from Saunders et al., 2016).

3.1 Research philosophy

The pragmatism philosophy, which emphasises practical solutions to real-world problems, was the foundation for this research (Creswell, 2013). In line with this, qualitative and quantitative approaches were used to develop and evaluate CNN, GAN, and Autoencoder models for fraud detection in Namibia's public sector. The models were evaluated using performance metrics such as inference time, memory usage, and detection accuracy to select the best model for the automated fraud detection system. Creswell (2013) emphasises the importance of using multiple data collection methods, combining qualitative and quantitative sources, and focusing on the research's practical implications. This study used expert feedback from domain experts to ensure that the model selection addressed the specific fraud challenges confronting Namibia's public sector. This approach ensured that the final

system was effective and in line with real-world conditions, thereby fulfilling pragmatism's core principles of prioritising practical outcomes.

3.2 Research design

The concept of research design can be broadly defined as a strategic blueprint that addresses research questions while aligning with the study's overall objectives (Sawsan & Jaradat, 2018). This study employed a mixed-methods approach, incorporating an abductive reasoning framework aligned with the pragmatism philosophy to achieve its objectives effectively. The quantitative component involved developing a deep learning-based fraud detection system using the Agile Machine Learning methodology. The system was iteratively refined, and its performance was evaluated through recall, accuracy, F1 score, and precision. Additionally, the quantitative analysis investigated the challenges and opportunities associated with scaling the fraud detection system to handle large volumes of financial data from public entities, thereby addressing the research question of scalability. The qualitative component, on the other hand, involved identifying vital financial indicators and features relevant to fraud detection. This aspect of the research used domain expert insights and literature analysis to guide the system's development. By combining qualitative and quantitative strategies, the research design ensured a comprehensive approach to answering the study's questions about system performance, resource utilisation, and identifying key fraud indicators.

3.3 Research instruments

According to Birmingham and Wilkinson (2003), in their renowned guide for researchers focused on research instruments, effective research instruments should facilitate the achievement of research objectives. Caseware IDEA software was used to sample and extract financial transactions from the government GL, and the fraud detection system was built around the best-performing deep learning model, which was chosen after a thorough evaluation of various algorithms. Custom Python scripts were used to evaluate the system's processing speed, memory usage, and scalability by measuring inference times, RAM usage, and GPU memory utilisation. Furthermore, correlation analysis and

random forest feature importance were used to identify critical financial features, with insights gathered from auditor interviews providing additional context for the fraud detection system. The performance of the fraud detection system was assessed by utilising various evaluation metrics on financial data that included simulated fraudulent transactions.

3.4 Data collection methods

Data collection is the process of gathering information from various relevant sources in order to address the research problem and evaluate the results (Dudovskiy, 2022). The primary data for this study were obtained from the Ministry of Finance and Public Enterprises through the Office of the Auditor-General, by the study's objectives. The obtained data consists of accounts payable transactions from Namibian public institutions' financial general ledger records for the fiscal years 2021/2022 and 2022/2023, and current supplier standing. The public institutions include governmental offices, ministries, agencies, and any other public entity that generates financial statements using the government's Integrated Financial Management System (IFMS). Additionally, this cross-sectional research design included a total of seven interviews with government financial and information systems auditors responsible for auditing these institutions' financial statements. A literature review was also conducted to identify any relevant findings that could contribute to this study.

3.5 Assumptions

The study focused on developing an automated financial fraud detection system using deep learning. The underlying presumptions included:

- The financial data obtained was accurate, complete, and free from significant errors.
- Auditors openly shared their experiences and challenges in detecting financial statement fraud, and the qualitative data collected from them was precise and impartial.

3.6 Data analysis

The first data analysis stage was pre-processing, which included feature engineering and data transformation, followed by data cleansing to correct errors. This process addressed the inherent

limitations associated with the occurrence, accuracy, and completeness of financial historical data while ensuring its quality and dependability. Then, statistical techniques such as correlation analysis were used to identify links between features and instances of fraud, which were graphically represented using charts and other visualisation tools. Furthermore, a comprehensive performance evaluation of the newly developed system was conducted.

3.7 Sampling procedures

A systematic sampling approach was used to ensure that a representative sample of Namibian public entities met the study's objectives. Alvi (2016) defines sampling as selecting a subset from a target population to make statistical conclusions and estimate population characteristics. The study's target population included 40 Namibian public entities, also known as "votes," which represented a variety of public institutions such as ministries and offices. A sample of 15 entities was drawn from this population based on their volume of financial transactions. The study aimed to analyse financial data in areas where fraud detection would have the most significant impact by focusing on entities with the highest transaction volumes. Based on transaction volume, this purposive sampling method ensured that the selected entities provided a robust dataset, thereby increasing the study's ability to identify relevant fraud patterns in Namibia's public sector financial transactions.

3.8 Ethical considerations

Before beginning data collection, the Namibia University of Science and Technology (NUST), Faculty of Computing and Informatics Research Ethics Committee provided ethical clearance, and the Deputy Auditor General of the Office of the Auditor-General granted permission to use government financial data. In addition, during the research period, the following ethical guiding principles were implemented:

- Unless otherwise specified, the researcher maintained the anonymity of interviewed participants throughout the study to ensure that their identities were kept confidential.

- Every research participant provided explicit consent, and no responses were modified in any manner.
- Recognition was given in full for implementing comparable systems approaches; all published data, methods, results, and other relevant resources are openly referenced.
- The study adhered to all applicable public sector research laws and regulations, such as data privacy and ethical research practices.

3.9 Chapter summary

This chapter details the research procedures and methodologies used to achieve the study's objective of creating a deep learning-based automated fraud detection system for Namibian public sector financial transactions. A mixed-methods approach, guided by pragmatism, was used to identify critical financial features and indicators relevant to fraud detection. This analysis was based on insights from industry experts and quantitative data provided by the Ministry of Finance and Public Enterprises for the fiscal years 2021/2022 and 2022/2023, as well as supplier standing data. Various sampling methods, including purposive sampling to target entities with the most transactions, ensured that representative data were collected. Data extraction was performed using Caseware IDEA, and Python scripts were used to assess the system's resource utilisation in handling large volumes of transactions. Statistical techniques, such as correlation analysis, were used to identify patterns between financial features and instances of fraud. The models' performances were evaluated using historical and simulated data, thereby ensuring that the developed fraud detection system is robust and effective. Data pre-processing and cleansing, combined with assumptions based on the literature, ensured that the analysis was accurate and reliable.

CHAPTER 4 : DESIGN AND DEVELOPMENT

This chapter provides an overview of the techniques and procedures used to develop a fraud detection system adapted for Namibia's public institutions, which is the primary goal of this research project. Figure 4.1 shows a high-level overview of the system's model-building cycle.

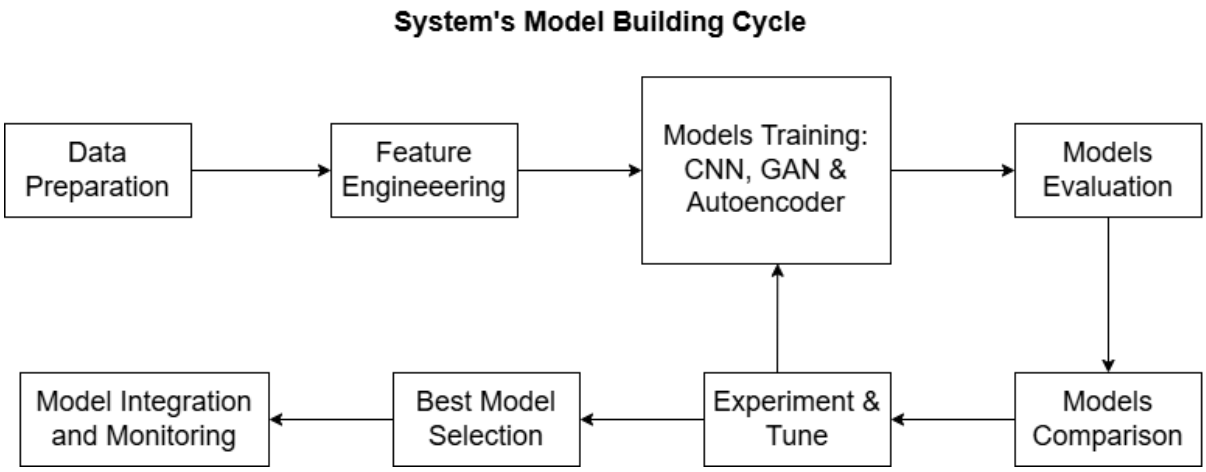

Figure 4.1 System's Model Building Cycle

4.1 System high-level architecture

The fraud detection system uses a modular and scalable architecture to process, analyse, and detect fraudulent patterns in financial transactions. The system combines deep learning techniques and data processing workflows, all based on fraud indicators identified through expert interviews. Each system component serves a specific purpose, ensuring that the solution is adaptable and capable of handling the complexities of large-scale financial data.

4.1.1 Data management

The system begins with the ingestion, cleaning, and pre-processing of AP transaction data and supplier standing data. These datasets are combined using relevant identifiers such as SUPPLIER_NUMBER and Vendor Number, allowing the system to cross-reference transaction details with supplier data. Pre-processing steps included dealing with missing values, standardising numerical and categorical features, and ensuring data integrity. Because of the sensitive nature of the financial data, strict data security

measures were implemented throughout the process to ensure confidentiality and prevent unauthorised access.

4.1.2 Feature engineering module

The system's ability to transform raw data into meaningful features for deep learning models is critical. The fraud detection key indicators shown and explained in Table 1 were developed based on auditor domain knowledge and data patterns. These features were intended to detect potentially fraudulent activities and were critical to the performance of the deep learning models.

Table 4.1 Overview of Feature Engineering in Model Development

Feature Name	Description	Inclusion Reason
Same_Initiated_Approved	Indicates if the transaction initiator and approver are the same	Detect fraud or conflicts of interest in approvals
Bank_Account_Consistency	Checks if the transaction and supplier bank account numbers match	Detect bank account fraud and payment discrepancies
Unusual_Purchase_Pattern	Flag transactions above a certain percentile	Detect unusually large transactions
Diff_GL_Code_Same_Vendor	Indicates whether a vendor has multiple GL codes	Detect abnormal accounting behaviour or misuse of funds
Amount_Discrepancies	Flag transactions that exceed vendors' average amounts	Detect overbilling
Duplicate_Invoices	Flags duplicate invoices (same vendor, invoice number, and amount)	Detect duplicate payments or invoice fraud

High_Amounts	Flag transactions above a certain percentile of all transactions	Detect high-risk transactions

4.1.3 Deep learning models

Three common deep learning models for fraud detection were developed and tested for the system: Convolutional Neural Network (CNN), Autoencoder, and Generative Adversarial Network (GAN). The CNN was designed to detect spatial and temporal patterns in transactions by using key fraud indicators as input features. Its ability to recognise complex relationships in data led to its selection as the most effective deployment model. The Autoencoder was trained to detect anomalies by reconstructing normal transaction patterns, while the GAN was used to generate synthetic fraudulent transactions, thereby increasing the robustness of fraud detection via adversarial learning. The evaluation of these models, detailed in the following chapter, revealed that CNN performed the best across multiple metrics.

4.1.4 Decision support system

The best-performing deep learning model developed in this study is integrated into a decision support system to improve auditor efficiency by flagging transactions based on predefined fraud indicators. The system flags transactions that exhibit patterns associated with fraudulent behaviour, such as unusual purchase patterns, significant discrepancies in amounts, or inconsistencies in account information, for further investigation. However, it is vital to note that the system flags all transactions meeting these criteria without distinguishing between fraudulent and erroneous transactions. Erroneous transactions, which can occur due to clerical errors, data entry errors, or misclassifications, are treated like potentially fraudulent transactions. The system does not evaluate or assign risk levels (high, medium, or low) to flagged transactions, nor does it consider materiality or potential financial impact. Instead, it flags any transaction that meets the fraud indicator threshold, leaving it up to experts such as auditors to investigate and determine whether the flagged transaction is fraudulent or simply an operational error.

By providing immediate predictions as well as the specific indicators that triggered the flag, the decision support system enables auditors to focus their attention on transactions that require further investigation. Although this targeted approach improves the overall auditing process, human expertise is still required to distinguish between fraudulent and erroneous transactions because the system does not make final determinations. Figure 4.2 shows the high-level architecture for the developed automated fraud detection system.

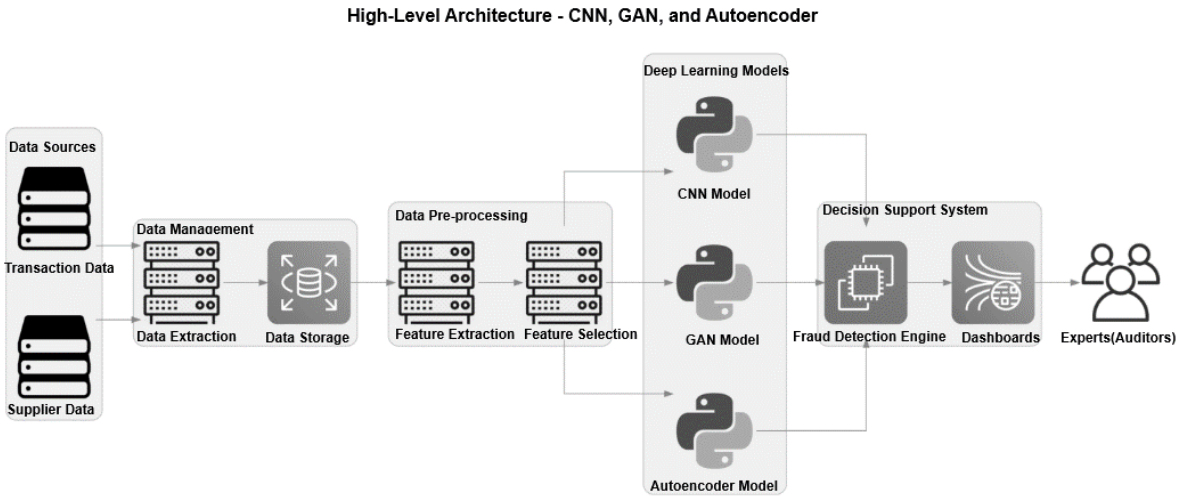

Figure 4.2 High-Level System Architecture for Fraud Detection Using Deep Learning

4.2 Dataset description

4.2.1 Transactions dataset overview

The AP transactions dataset contains 201,850 records for the fiscal years 2021 and 2022. This dataset has 29 features, including critical attributes like VENDOR_ID, AMOUNT, BANK_ACCOUNT_NUM, INITIATED_BY, APPROVED_BY, GL_CODE, and INVOICE_REFNO. These variables are critical for detecting patterns of possible fraud in public transactions. As presented in Figure 4.3, the dataset contains 691 duplicate rows. The data takes up approximately 278 MB of memory, with an average record size of 1.4 KB. The variable types are numeric, categorical, text, and DateTime, and they are essential for modelling the fraud detection system.

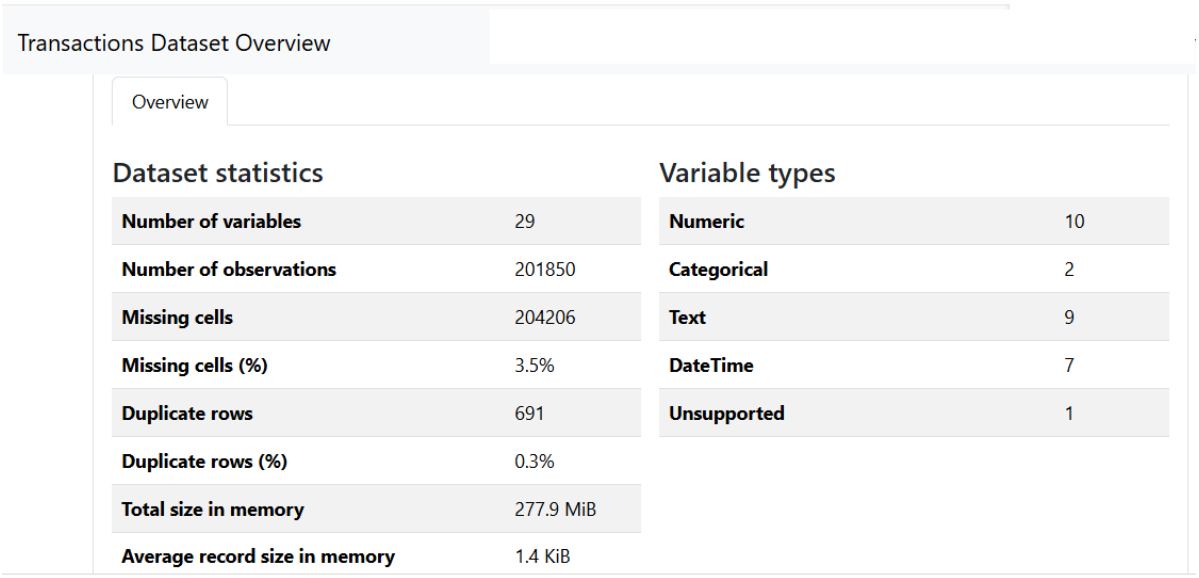

Figure 4.3 Transactions Dataset Overview

The dataset underwent pre-processing, which included removing duplicate rows and converting categorical and numerical features to standard formats. Fraud detection indicators derived from auditor interviews were used to improve the system's ability to detect suspicious transaction patterns. These indicators contributed domain knowledge unique to Namibian public institutions, enriching the dataset with real-world fraud risk considerations. Furthermore, the dataset was heavily synthesised using domain knowledge to address the imbalance, as it was mostly made up of non-fraudulent transactions, which was an important adjustment to ensure a more balanced and accurate detection system.

4.2.2 Suppliers dataset overview

The supplier standing data includes 174,940 records with detailed information about public institution suppliers. This dataset contains 28 features, such as VENDOR_ID and BANK ACCOUNT NUMBER, which were useful for cross-referencing and validating AP transactions. As shown in Figure 4.4, the dataset has 77 duplicate rows. The dataset takes up approximately 228.8 MB of memory, with an average record size of 1.3 KB.

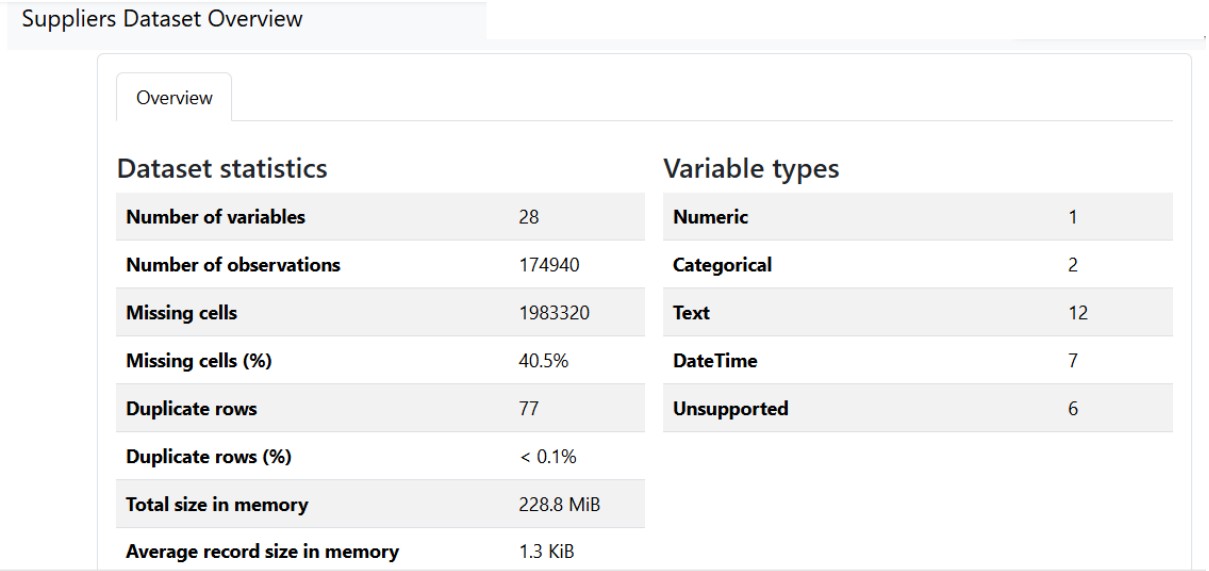

Figure 4.4 Suppliers Dataset Overview

Similar to the AP transactions dataset, the suppliers' dataset was pre-processed, including the removal of duplicate records. These steps were necessary to ensure that the data was appropriate for cross-validation. The system cross-referenced key attributes like VENDOR_ID and BANK ACCOUNT NUMBER to look for inconsistencies like mismatched bank account information or duplicate supplier entries, both of which are common indicators of fraudulent activity.

To summarise, the two datasets were pre-processed and integrated into the fraud detection system, yielding a comprehensive view of financial transactions and supplier standing data. This integration enabled the detection system to effectively identify discrepancies and potential fraud by combining transactional data and supplier metadata. By combining these datasets with domain-specific fraud indicators, the system accurately reflected the complexities of financial fraud in Namibia's public sector.

4.3 Development approach

The deep learning-based fraud detection system was developed and implemented to ensure robustness, efficiency, and flexibility. The system architecture ensured that every component, from data ingestion to immediate fraud detection, was developed independently, allowing for easy maintenance and updates.

4.3.1 Development practices

Several best practices were implemented to ensure a streamlined and efficient development process:

- Version Control: Google Colab was used for development. It provides automatic versioning and the ability to restore previous notebook versions, thereby improving the overall workflow.
- Modular Programming: The system was designed with a modular architecture, breaking down data pre-processing, feature engineering, and deep learning model training into independent components. This method enabled various components' simultaneous development and testing, thereby simplifying debugging and updates.
- Data Pre-processing and Model Training: This included feature standardisation with 'StandardScaler' and merging the transactions and suppliers' datasets, which was designed to handle missing values and inconsistencies. Key fraud indicators were designed to detect suspicious transactions, and the deep learning models were trained with TensorFlow.

4.3.2 Data pre-processing

To prepare the dataset for training, it was split into two parts: 80% for training and 20% for testing. Each feature was standardised with StandardScaler to ensure that it was scaled to a consistent range. This standardisation step was critical for optimising model performance because it accelerated convergence and increased accuracy for the CNN model, stabilised the training process for the GAN architecture, and reduced reconstruction error for the Autoencoder model. For the CNN, the dataset was reshaped into the required format (samples, time steps, and features), with each transaction represented as a sequence of multiple features. Additionally, input values were scaled from -1 to 1, which was critical for stabilising the GAN's training process. The Autoencoder used standardised features as inputs and outputs, enabling the model to learn and reconstruct the original data effectively.

The key fraud detection features used in this study are Same_Initiated_Approved, Bank_Account_Consistency, Unusual_Purchase_Pattern, Diff_GL_Code_Same_Vendor, Amount_Discrepancies, Duplicate_Invoices, and High_Amounts. These features were engineered by

combining transactional and supplier data using SUPPLIER_NUMBER and VENDOR_NUMBER as common identifiers. Each feature was designed to detect distinct fraud patterns, such as inconsistencies in bank account information or unusually large transactions. The fraud labels (is_fraud) were created based on the identified fraud indicators.

4.3.3 Challenges and solutions

Several technical challenges, as summarised in Table 4.2, arose during development. Each challenge was addressed with a specific solution to maintain the system's performance and efficiency:

- Handling Large Datasets: The system had to handle a large volume of data (201,850 transactions and 174,940 supplier records) while remaining efficient. Memory optimisation techniques were used during data ingestion, and missing data points, such as incomplete bank account information, were flagged. The features were standardised with 'StandardScaler', which improved the performance of the deep learning models.
- Immediate Processing: The CNN model was designed to process transactions instantly, thereby providing fast inference and allowing auditors to receive immediate alerts for potentially fraudulent transactions. Thanks to a reshaped input structure (samples, time steps, and features), the CNN processed financial data streams efficiently and without significant delays.
- Data Security and Privacy: While specific encryption protocols were not included in the development environment, using Google Drive for data storage ensured secure and easy access to sensitive financial data. Access to the Google Colab environment was restricted to prevent unauthorised access.

Table 4.2 Challenges and Solutions

CHALLENGE	SOLUTION IMPLEMENTED
HANDLING LARGE DATASETS	Optimised pre-processing using StandardScaler, flagged missing data and applied memory-efficient techniques
IMMEDIATE PROCESSING	Implemented CNN with reshaped input to process transactions instantly
DATA SECURITY AND PRIVACY	Data is stored securely in Google Drive, with restricted access in the Google Colab environment

4.4 Model development

The research developed and compared three deep learning models: CNN, Autoencoder, and GAN. The models were evaluated for their efficacy in detecting fraudulent transactions.

4.4.1 CNN model

The CNN was designed to detect patterns in financial transaction data by leveraging its ability to capture spatial and temporal dependencies. The CNN model was ideal for this task because of its ability to detect complex patterns across multiple features and time-series relationships in financial transactions.

CNN Architecture

The CNN model was built with a 1D Convolutional Neural Network architecture, which included:

- Input Layer: A Conv1D layer with 64 filters, a kernel size of 1, and ReLU activation that extracted feature patterns from the input data.
- MaxPooling Layer: A pooling layer (MaxPooling1D) to downsample the input, lowering computational costs and preventing overfitting.
- Dropout Layer: To reduce overfitting, a Dropout layer with a rate of 0.5 was introduced, which disables neurons at random during training.

- Flatten Layer: The convolution and pooling layer outputs were flattened before being fed into the fully connected layers.
- Fully Connected Layers: To capture complex nonlinear relationships in the data, a dense layer of 100 units and ReLU activation was used. This was followed by an output layer that used a sigmoid activation function to determine whether a transaction was fraudulent or not.

Training and evaluation

The CNN was trained over 20 epochs with a batch size of 32, using the binary cross-entropy loss function to assess model performance and the Adam optimiser for weight updates. Throughout the training process, 20% of the dataset was reserved for validation to track model performance. The final model demonstrated a high ability to detect fraudulent transactions by effectively capturing patterns in financial data. When evaluated on the test set, the model had a loss of 0.80 and an accuracy of 95.4%, indicating strong performance. Figure 4.5 shows the high-level architecture for the CNN model.

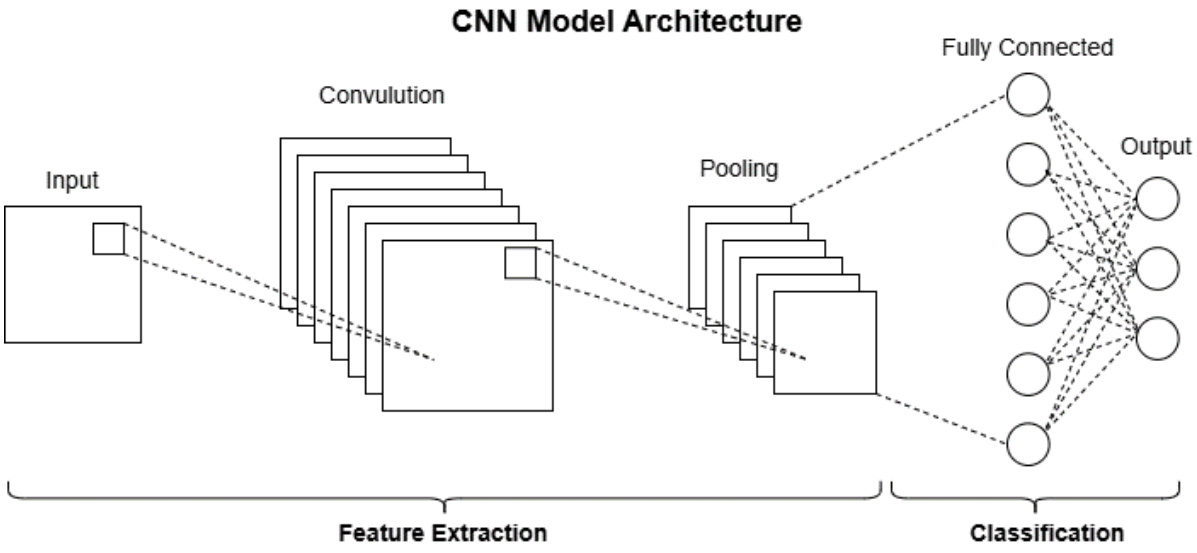

Figure 4.5 High-level Architecture for the CNN Model

4.4.2 GAN model

Through adversarial learning, the GAN generated synthetic fraudulent transactions to improve the model's fraud detection capabilities. The GAN framework is made up of two adversarial neural networks: a generator that creates synthetic transaction data and a discriminator that distinguishes

between real and synthetic data. This adversarial process allows both networks to iteratively improve, producing more realistic synthetic data and increasing the model's ability to detect fraudulent patterns.

GAN architecture

The GAN architecture has two main components:

- Generator: The generator network accepts a random noise vector (latent dim = 100) as input and converts it into a synthetic transaction identical to the original data. The network comprises several fully connected layers with LeakyReLU activations and batch normalisation to help stabilise training and ensure that the generated data closely resembles the actual transaction data.
- Discriminator: The discriminator is a binary classification network that accepts a transaction as input (real or synthetic) and returns a probability indicating whether the transaction is real or generated. It comprises fully connected layers with LeakyReLU activations and dropout layers to avoid overfitting. The discriminator was trained with binary cross-entropy loss to differentiate between real and generated transactions.

Training and evaluation

The GAN model was trained for 500 epochs with a batch size of 64, with the generator producing synthetic fraudulent transactions and the discriminator learning to differentiate between real and synthetic data. A balanced training strategy was used, with the discriminator updated more frequently to ensure effective adversarial learning. Both networks were optimised with the RMSprop optimiser and a learning rate of 0.00005. Following training, the generator generated synthetic fraudulent transactions and real data to address the class imbalance in the fraud detection task. This augmented dataset improved the training data by providing a more balanced representation of fraud cases.

For evaluation, an XGBoost classifier was trained on a dataset that included real and synthetic transactions generated by the GAN. This enabled a performance comparison with traditional supervised learning methods and helped assess the effectiveness of synthetic data in improving model

performance. Synthetic GAN data was intended to improve the classifier's ability to detect fraud by providing a more comprehensive range of fraud patterns, thereby increasing overall fraud detection effectiveness. Figure 4.6 shows the high-level architecture for the GAN model.

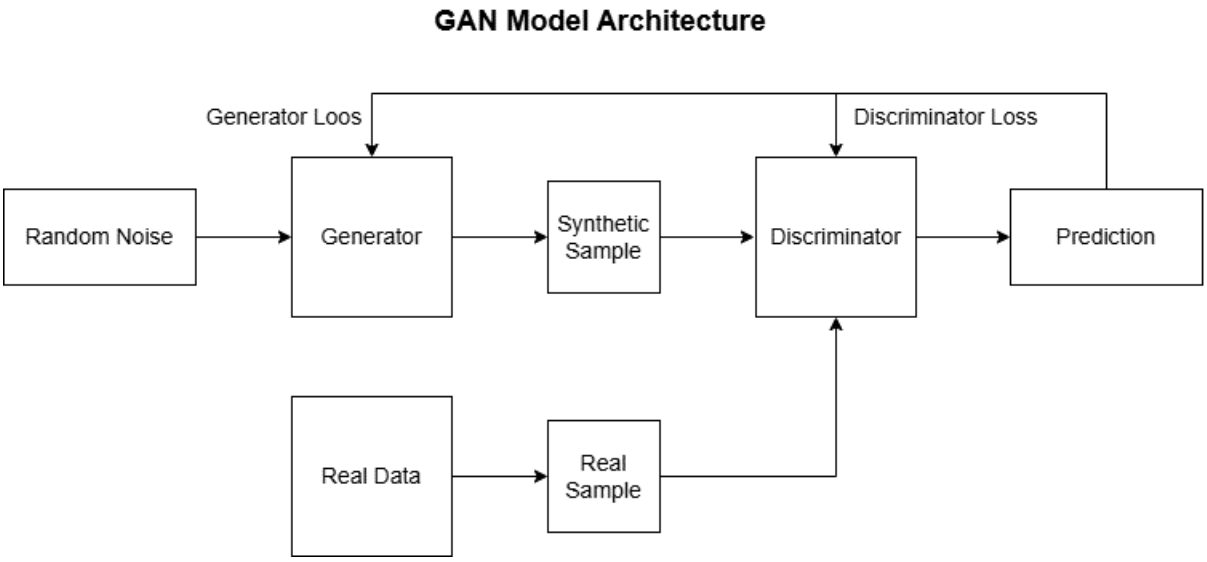

Figure 4.6 High-level Architecture for the GAN Model

4.4.3 Autoencoder model

The Autoencoder was developed as an unsupervised learning model to detect anomalies in financial transaction data by reconstructing normal patterns and flagging deviations as potential fraud. This approach enabled the model to learn the structure of non-fraudulent transactions and detect fraud using reconstruction errors.

Autoencoder architecture

The Autoencoder architecture consisted of two major components: an encoder and a decoder.

- Encoder: The encoder compressed input data into a lower-dimensional space using a dense layer with four neurons and ReLU activation. This encoded representation captured the critical aspects of the transactions.
- Decoder: The decoder reconstructed the input data from the compressed representation by employing a dense layer with the same number of neurons as the input dimensions. The decoder aimed to replicate the input as closely as possible.

Training and evaluation

The Autoencoder model was trained using the mean squared error (MSE) loss function, which computes the reconstruction error between the original input and the reconstructed output. The model was trained for 30 epochs with a batch size of 32. During training, 20% of the dataset was validated to monitor the model's performance. After training, the reconstruction error was calculated for both the training and testing sets, indicating the difference between the original input and the model's output. Transactions were classified using a threshold based on the 95th percentile of reconstruction errors from the training set. Transactions containing reconstruction errors greater than this threshold were flagged as potentially fraudulent. The Autoencoder learnt to reconstruct non-fraudulent transactions, which enabled it to detect anomalies with higher reconstruction errors. When evaluated on the test set, the model successfully distinguished between legitimate and potentially fraudulent transactions. Figure 4.7 shows the high-level architecture for the Autoencoder model.

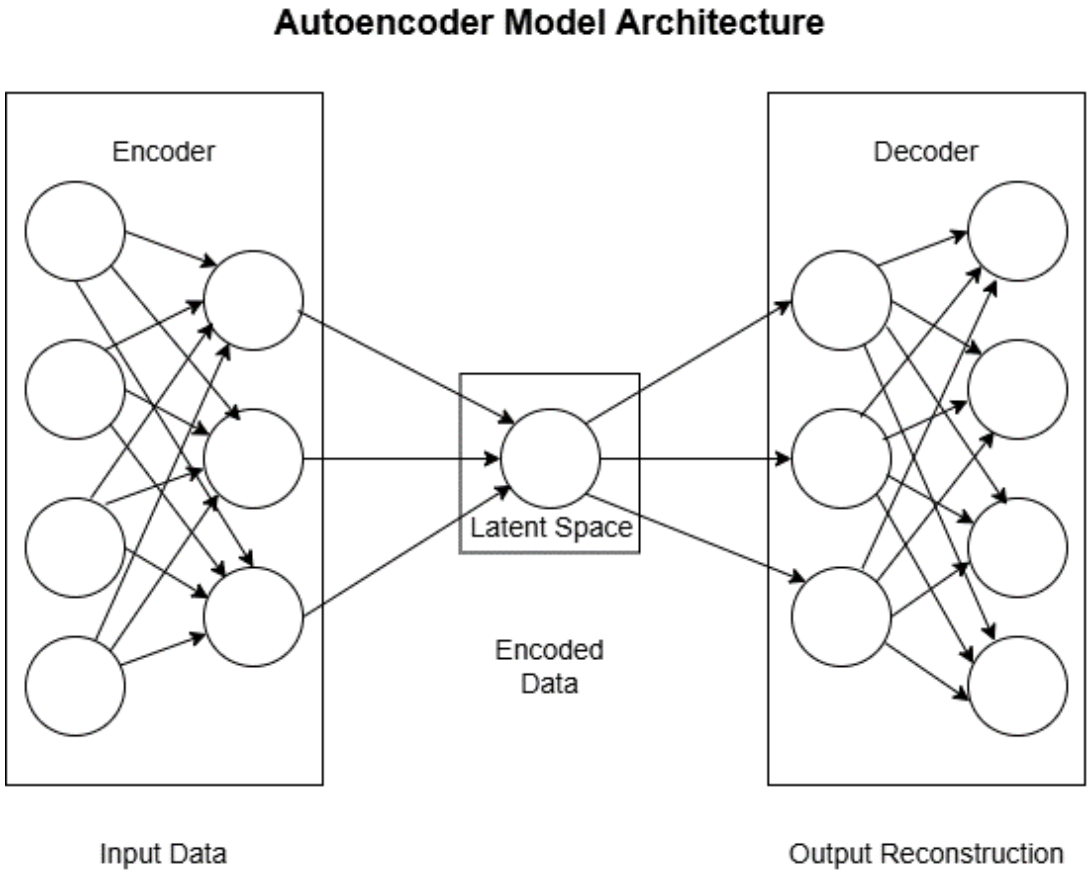

Figure 4.7 High-level Architecture for the Autoencoder Model

4.5 Model deployment

The best model, CNN, was implemented within the fraud detection system to facilitate the immediate classification of transactions as fraudulent or non-fraudulent. The deployment strategy prioritised scalability, instantaneous processing, and system monitoring for continuous updates.

4.5.1 Model serving

The CNN model was deployed using TensorFlow's model-serving framework, allowing instantaneous transaction processing and classification. The system was designed to recalculate key fraud indicators for each incoming transaction, ensuring access to the most recent classifications.

An intuitive user interface was developed for demonstration purposes, allowing users to manually enter transactions and control the analysis flow. This dashboard shows each transaction with all relevant information, including any triggered fraud indicators. Figure 4.8 shows the interface with a colour-coded table, with fraudulent transactions highlighted in red and non-fraudulent ones highlighted in green. This visual distinction helps identify and prioritise high-risk transactions.

A dynamic bar graph also shows the overall distribution of fraud versus non-fraud transactions, making it easier to track trends and focus on areas of concern. This comprehensive approach ensures that users can effectively monitor transaction activity and base their decisions on the system's analysis.

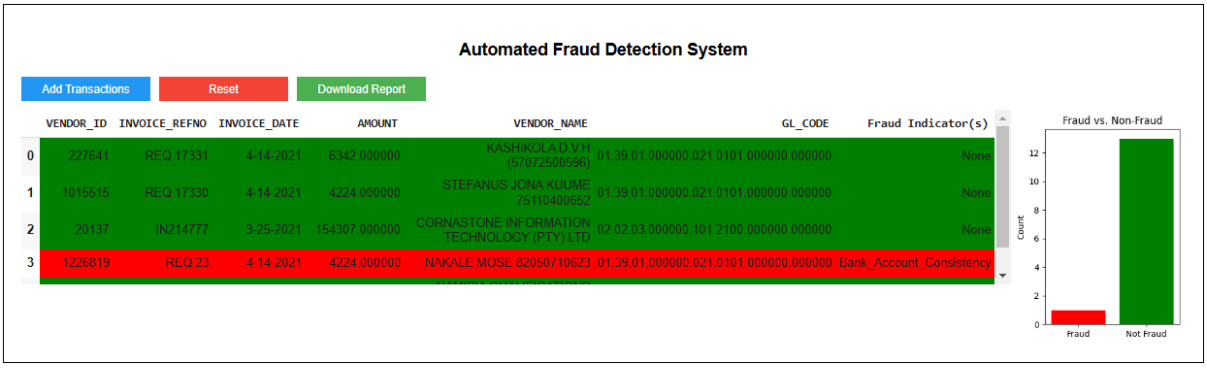

Figure 4.8 User Interface of the Fraud Detection System

4.5.2 Resource utilisation

The system was designed to emphasise efficient resource utilisation to manage the high volume of transactions processed by various public institutions. The backend was optimised to maximise computational efficiency, allowing the system to handle rapid transaction feeds while minimising resource overhead, even during peak transaction times. Memory-efficient techniques, such as standardised features, contributed to the system's ability to handle large datasets without experiencing significant slowdowns. Additionally, the system's modular architecture enabled future integration of additional models or data streams, ensuring flexibility while optimising resource use as new fraud patterns and indicators emerged.

4.5.3 Monitoring and updates

For this research phase, all processes, including metric tracking, were carried out manually. The deployment process involved manually monitoring the CNN model's performance in real-time. The intent was to closely monitor performance degradation by manually tracking key metrics such as accuracy, recall, and false positive rates. Furthermore, a mechanism was developed to allow manual retraining of the model with new transaction data as patterns evolved, ensuring that the model remained relevant and accurate. Although the retraining process is not currently automated, updates are scheduled to occur regularly in response to significant changes in transaction patterns. This adaptive approach ensures that the system can respond to new fraud tactics while improving over time.

4.6 Chapter summary

This chapter covered designing, developing, and deploying a deep learning-based fraud detection system tailored to Namibian public sector financial transactions. Three models were developed: CNN, Autoencoder, and GAN. The system used key fraud indicators from interviews with experienced auditors who audited public institutions.

The system's modular and scalable architecture enabled efficient data processing and immediate fraud detection, making it appropriate for public institutions' large volumes of transactions. The deployment

of the CNN model enabled instantaneous monitoring via an intuitive interface, giving auditors timely alerts and insights into potentially fraudulent cases. Continuous monitoring and retraining were implemented to ensure the system could adapt to changing fraud tactics.

The data used for this project includes accounts payable transactions for the fiscal years 2021/2022 and 2022/2023, as well as supplier data for public institutions. Overall, this system represents a significant step in using deep learning techniques to improve financial fraud detection in Namibia's public sector.

CHAPTER 5 : RESULTS AND DISCUSSION

This chapter presents the results of the three developed models: CNN, GAN, and Autoencoder. The models' efficacy in the fraud detection task was assessed using various metrics, including accuracy, precision, recall, F1-score, and ROC-AUC. The comparative analysis seeks to identify the most effective model for detecting fraudulent transactions.

5.1 Evaluation metrics

The following metrics were employed to thoroughly assess the models' performance:

- Accuracy: The proportion of correct predictions out of all predictions made, calculated using the formula $\frac{TP+TN}{TP+TN+FP+FN}$.
- Precision: Determines the proportion of predicted fraudulent transactions that were fraudulent using the formula $\frac{TP+TN}{TP+TN+FP+FN}$.
- Recall (Sensitivity): The model's ability to identify actual fraudulent transactions is calculated using the formula $\frac{TP}{TP+FP}$.
- F1-Score: Achieves a balance between precision and recall, which is especially useful when dealing with class imbalances, calculated using the formula $2 * \frac{Precision*Recall}{Precision+Recall}$.
- ROC-AUC (Receiver Operating Characteristic - Area Under Curve) demonstrates the model's ability to distinguish between classes at various thresholds.
- Confusion Matrix: Display a summary of the prediction outcome, including true positives, false positives, true negatives, and false negatives.

5.2 CNN Results

5.2.1 Training performance

Figure 5.1 shows the CNN model's training performance, including graphs of accuracy and loss over 20 epochs for the training and validation datasets. Initially, both curves rise steadily, indicating that the model effectively learns from the data. Around the tenth epoch, the accuracy stabilises, reaching about

95.3% for training and validation. While there are minor fluctuations in validation accuracy, this is common and reflects the variability within the validation set. Moreover, the strong correlation between training and validation accuracy indicates that the model is generalising well.

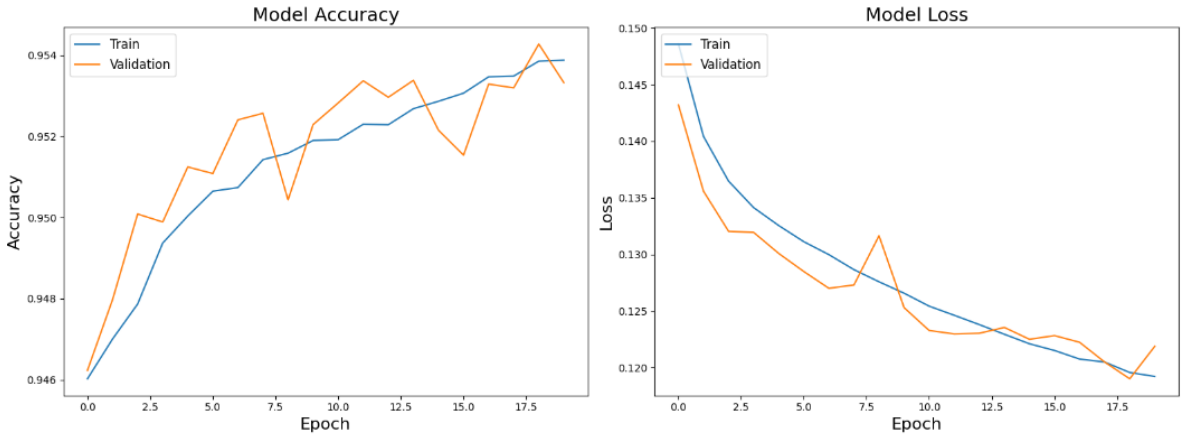

Figure 5.1 Training & Validation Accuracy and Loss of CNN Model

The right graph shows the loss incurred during each epoch for both datasets. The training and validation losses demonstrate a steady decline, signifying that the model has effectively converged. The training loss consistently declines, while the validation loss exhibits a comparable trend, though with slight fluctuations. These variations are anticipated due to the stochastic characteristics of mini-batch gradient descent and do not signify any significant instability.

The persistent reduction in training and validation loss and the stable accuracy curves signify that the CNN model has effectively learnt the data patterns without overfitting. The minor fluctuations in validation metrics indicate that the model may improve with slight hyperparameter adjustments or additional data. Nonetheless, the model's training and validation performance are consistent, suggesting that it is robust and capable of generalising to novel data. This makes the CNN model a reliable approach for detecting fraud in financial transactions, with the potential for further performance improvement.

5.2.2 Testing results

The CNN model's testing phase provided valuable insights into its performance on previously unseen data. The model achieved an overall testing accuracy of 95.4% and a test loss of 0.1217, indicating strong performance after the training phase.

Figure 5.2 shows the confusion matrix, which offers a detailed breakdown of the model's classification performance. The model correctly identified 1,592 fraud instances (true positives) and 80,584 non-fraud cases (true negatives). However, 3,198 non-fraudulent transactions were incorrectly flagged as fraudulent (false positives), while 774 actual fraudulent activities were not detected (false negatives). The relatively low false negative rate indicates that the model can identify most fraud cases, which is critical for fraud detection. However, the higher number of false positives suggests room for improvement in reducing incorrect fraud alarms, which could be addressed by fine-tuning the model or incorporating more distinguishing features.

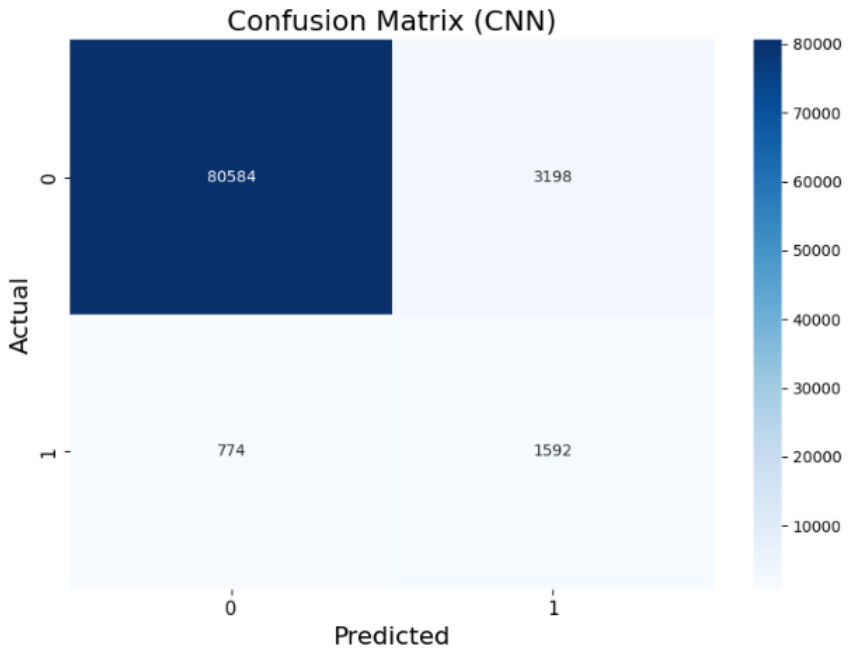

Figure 5.2 Confusion Matrix for CNN Model on Test Data

The classification report confirms the model's effectiveness, with a precision of 0.96 for fraud detection, indicating that the model has a high rate of correctly identifying fraudulent cases. The recall rate of 0.99

suggests that it can detect almost all fraud cases. The F1-score of 0.98 indicates an excellent balance of precision and recall. Despite an overall high accuracy of 95%, the macro average metrics (precision, recall, and F1-score) are significantly lower due to class imbalance, with the dataset containing significantly more non-fraudulent cases than fraudulent ones.

The feature importance graph in Figure 5.3 highlights the key indicators contributing to the model's predictions.

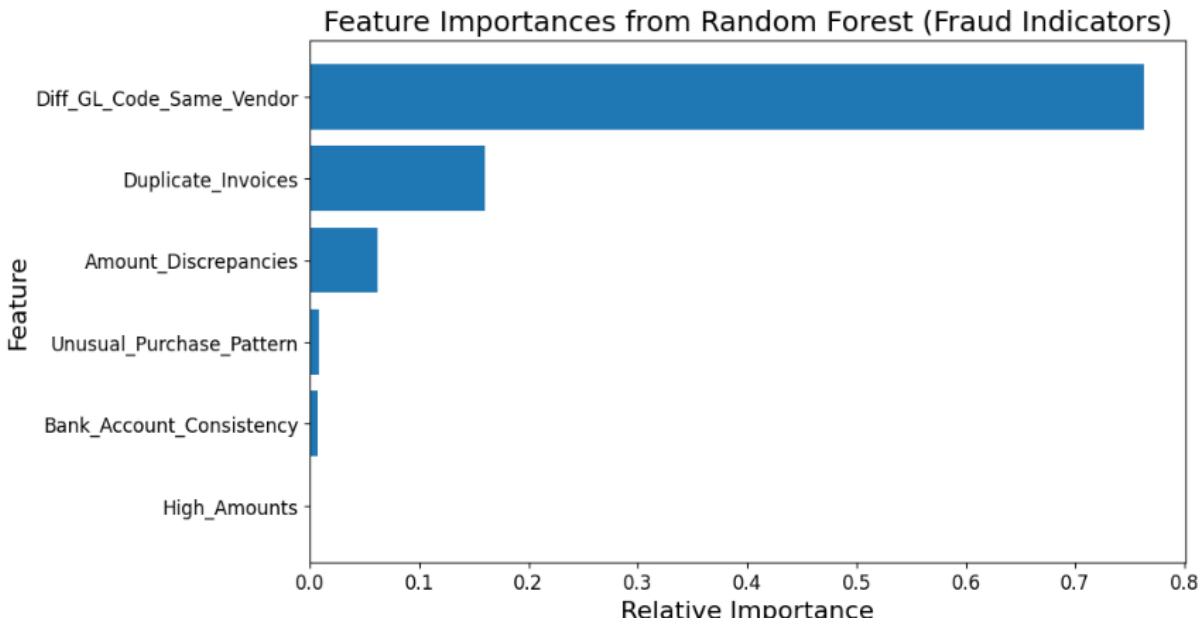

Figure 5.3 Feature Importance from Random Forest for Fraud Detection Indicators

Diff_GL_Code_Same_Vendor was the most critical fraud detection feature, implying that vendors who frequently use different general ledger GL codes may be involved in fraudulent transactions. This feature had the highest relative importance, making it critical for detecting fraud. Duplicate invoices were a notable indicator, followed by amount discrepancies, suggesting that transactions with inconsistent amounts are often linked to fraudulent activity. Notably, features like Unusual_Purchase_Pattern and Bank_Account_Consistency had a lower impact, implying that, while potentially useful, they make only minor contributions to fraud detection. This information can help guide future efforts to fine-tune the feature set, focusing on the most critical factors while improving or re-evaluating the less important ones.

5.3 GAN results

5.3.1 Training performance

Figure 5.4 shows the training performance of a GAN designed to improve fraud detection in financial transactions. The discriminator loss for real and synthetic transactions is initially high but gradually decreases, indicating improved classification performance as training advances. The minimal disparity between real and synthetic losses suggests that the generator effectively produces synthetic fraudulent transactions that increasingly mimic real-world fraud patterns.

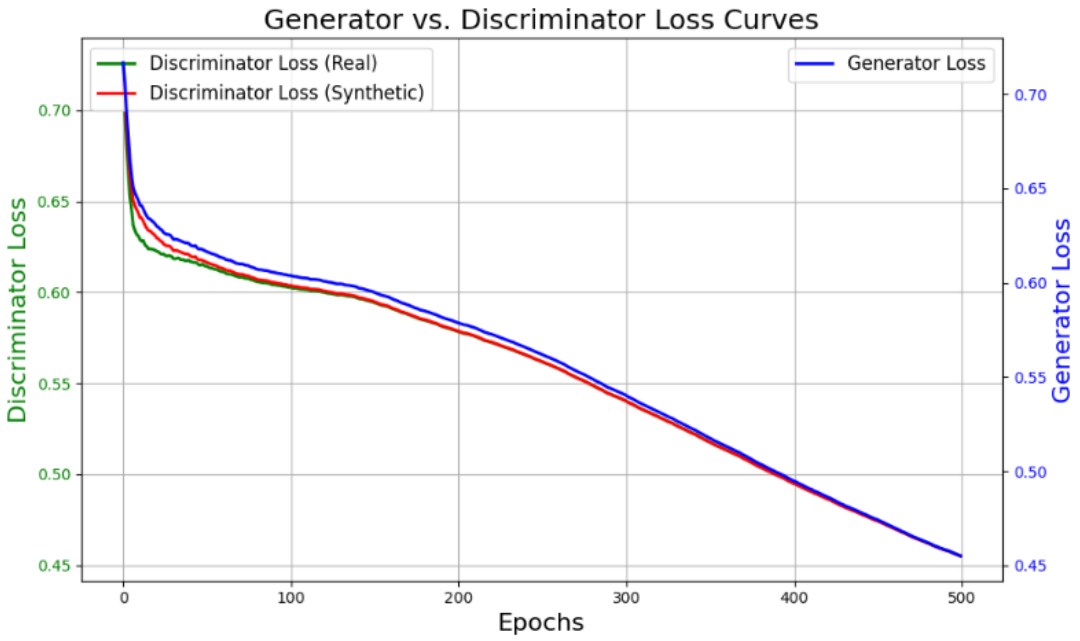

Figure 5.4 Generator vs. Discriminator Loss Curves during GAN Training

The generator loss decreases as the generator's ability to produce realistic fraudulent samples increases. The GAN-generated data is essential to the research objective as it offers a solid foundation for improving fraud detection in the automated system. The convergence of losses after approximately 100 epochs indicates balanced adversarial training, with both networks learning at the same rate. This performance demonstrates that the GAN can produce synthetic fraud data, which may significantly improve the effectiveness of a detection system relative to CNN and Autoencoder models.

5.3.2 Testing results

The confusion matrix in Figure 5.5 indicates that the model accurately identified 3,338 fraudulent transactions and 75,272 non-fraudulent transactions while misclassifying 6,855 fraudulent transactions as non-fraudulent and incorrectly labelling 683 non-fraudulent transactions as fraudulent. The results demonstrate that the model is proficient in identifying non-fraudulent cases; however, further optimisation is necessary to minimise false negatives.

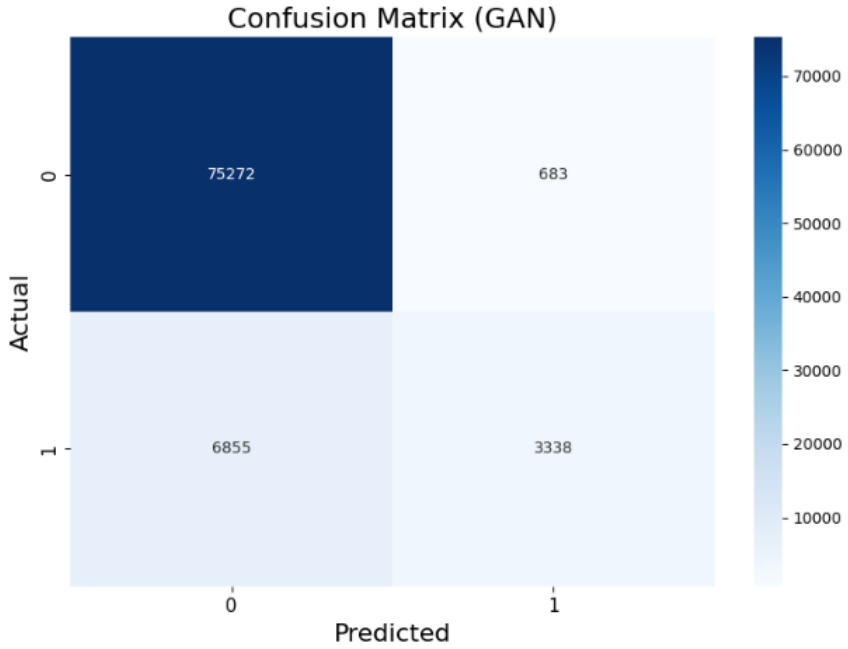

Figure 5.5 Confusion Matrix for GAN Model

The classification report emphasised the model's effectiveness in identifying fraud, achieving a precision of 0.99 and a recall of 0.92 for fraudulent transactions, resulting in a high F1-score of 0.95. This performance indicates that the model is proficient in identifying fraudulent transactions despite missing some actual fraud cases. In non-fraudulent transactions, the precision is 0.33, and the recall is 0.83, indicating that the model has difficulty misclassifying some non-fraudulent transactions as fraud. Despite the dataset's imbalance, the model's overall accuracy of 91% and weighted average F1-score of 0.93 demonstrate its effectiveness.

Cosine similarity and Euclidean distance were employed to compare real and synthetic data produced by the GAN. A cosine similarity score of -0.0057, as shown in Table 5.1, indicates that the direction of

the synthetic data diverges from that of the actual data, underscoring the challenge of replicating authentic fraud patterns. The mean Euclidean distance between real and synthetic data is 3.7323, signifying a substantial disparity in magnitude that may facilitate the identification of abnormal behaviour.

Table 5.1 Cosine Similarity and Euclidean Distance between Real and Synthetic Data

Metric	Value
Cosine similarity	-0.0057
Average Euclidean distance	3.7323

These metrics indicate that although an obvious distinction exists between real and synthetic data, this disparity can be advantageous for detecting potentially fraudulent transactions, thereby improving the model's anomaly detection efficacy.

Figure 5.6 on the next page presents a visual comparison of real and synthetic data for key fraud detection indicators. The Diff_GL_Code_Same_Vendor and Unusual_Purchase_Pattern features show considerable discrepancies, indicating that the GAN had difficulty replicating the intricate patterns present in real data. Conversely, indicators like Duplicate_Invoices and Amount_Discrepancies display minimal variations, signifying that the GAN effectively captured these fraud-related characteristics. Although synthetic data is beneficial, specific critical features necessitate additional refinement to improve the accuracy of fraud detection models trained on GAN-generated data.

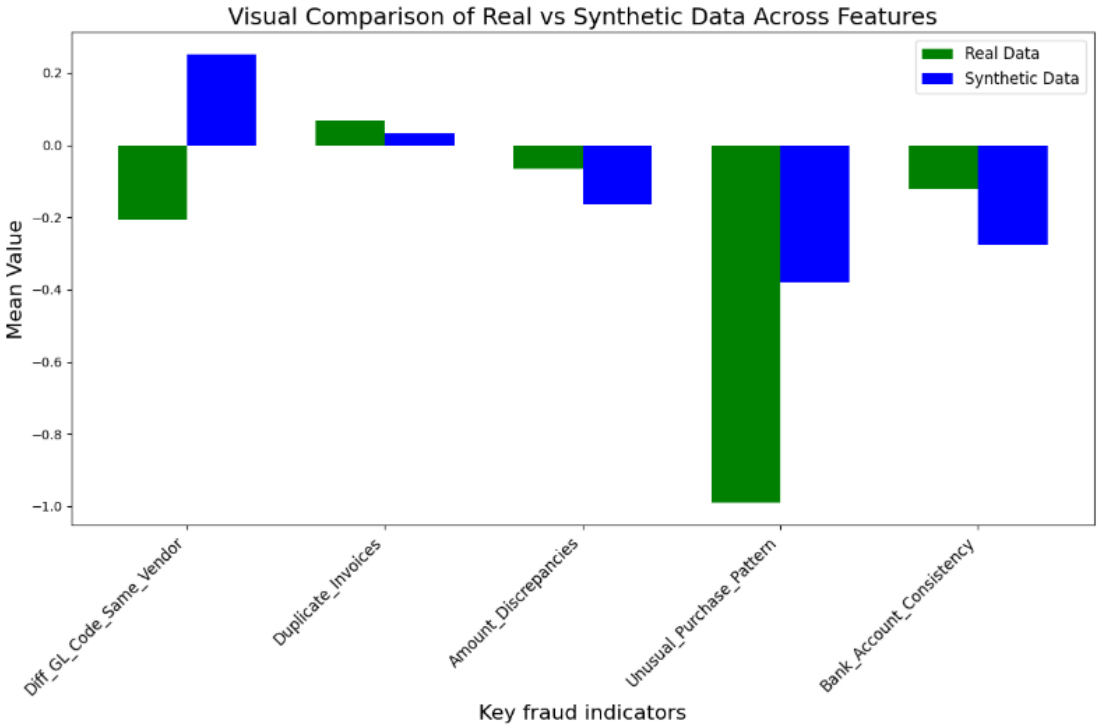

Figure 5.6 Visual Comparison of Real vs. Synthetic Data across Key Fraud Detection Features

5.4 Autoencoder results

5.4.1 Training performance

The Reconstruction Loss Curve in Figure 5.7 shows the Autoencoder model's training performance over 20 epochs. The graph shows the training and validation losses, offering insight into the model's ability to reconstruct the input data.

The training loss starts at around 0.5 and drops sharply over the first 5 epochs, indicating that the model is quickly learning the fundamental structure of the input features. The validation loss follows a similar trend, indicating that the model is effectively generalising to the validation set during the training process.

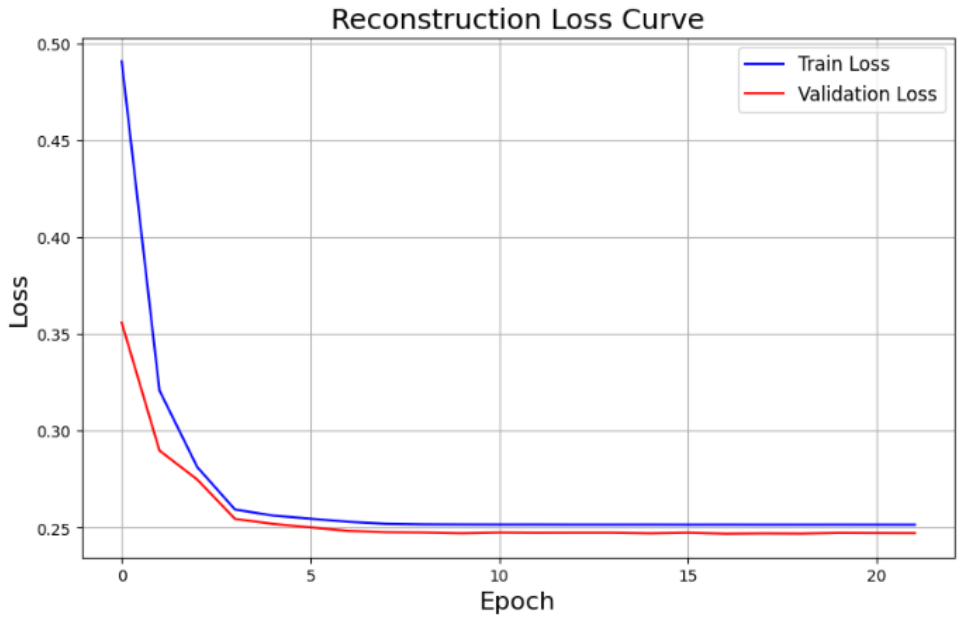

Figure 5.7 Reconstruction Loss Curve of the Autoencoder Model

Following about five epochs, the training and validation losses stabilise at around 0.25. This stabilisation indicates that the model has reached a point of diminishing returns, where additional training produces no significant improvements in reconstruction accuracy. The proximity of training and validation losses suggests minimal overfitting, as the model performs consistently on both datasets.

The smooth convergence of the loss curves indicates that the Autoencoder successfully acquired a compressed representation of the input data with minimal reconstruction error. This is especially important in anomaly detection tasks, such as fraud detection, where the Autoencoder must distinguish anomalies by reconstructing normal patterns and flagging instances with high reconstruction errors.

5.4.2 Testing results

The effectiveness of the Autoencoder model in anomaly detection was evaluated using a reconstruction error threshold of 0.001068, based on the 25th percentile of training reconstruction errors. The confusion matrix in Figure 5.8 shows that the model correctly identified 61,616 non-fraudulent transactions while misclassifying 20,578 non-fraudulent cases as fraudulent (false positives). Furthermore, the model misclassified 3,172 fraudulent transactions as non-fraudulent while correctly identifying only 782 fraudulent transactions. The high rate of false positives indicates that the threshold

used was overly sensitive, causing most transactions to be classified as fraudulent. Although higher percentiles had 100% detection rates, they were impractical for real-world applications. The 25th percentile reached a higher balance, highlighting the inherent challenge of calibrating Autoencoder thresholds to reduce false positives.

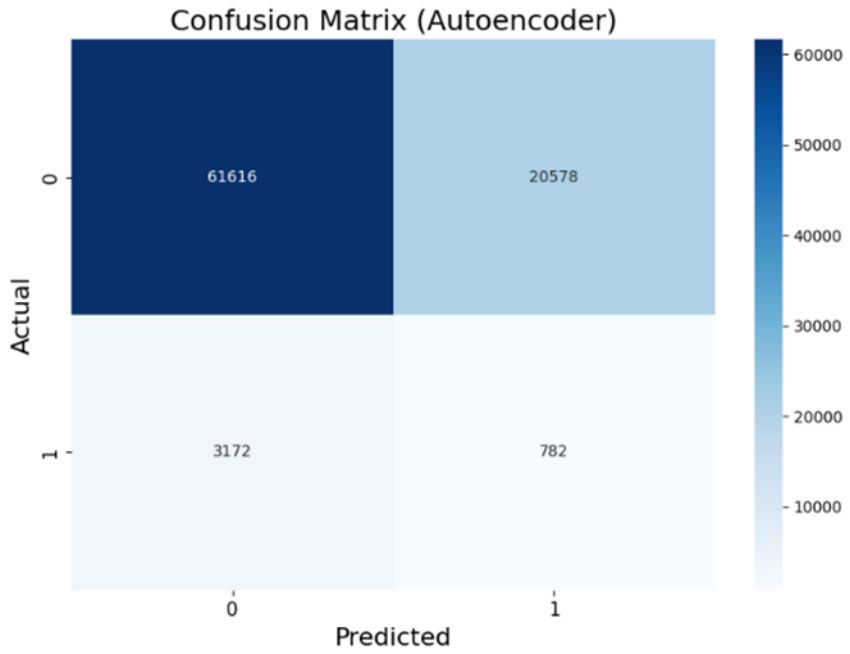

Figure 5.8 Confusion Matrix of Autoencoder Model on Test Data

The classification report emphasises the model's effectiveness, with a precision of 0.95 in detecting fraudulent transactions, indicating that the model correctly predicts fraud 94% of the time. The recall for fraud detection is 0.75, which means that the model only identified 75% of actual fraud cases, leaving out a significant portion. An F1-score of 0.84 indicates a good balance of precision and recall for the fraudulent class. Nonetheless, the model performed poorly in the non-fraudulent class, with a precision and recall of 0.04 and 0.06, respectively, owing to a high false positive rate.

Figure 5.9 shows the Receiver Operating Characteristic (ROC) Curve, highlighting the model's performance. The area under the curve (AUC) is 0.46, significantly lower than the acceptable level, indicating that the Autoencoder has difficulty distinguishing between fraudulent and non-fraudulent transactions. The curve's nearly linear shape emphasises the model's limited ability to differentiate between the two classes, consistent with the high number of false positives recorded.

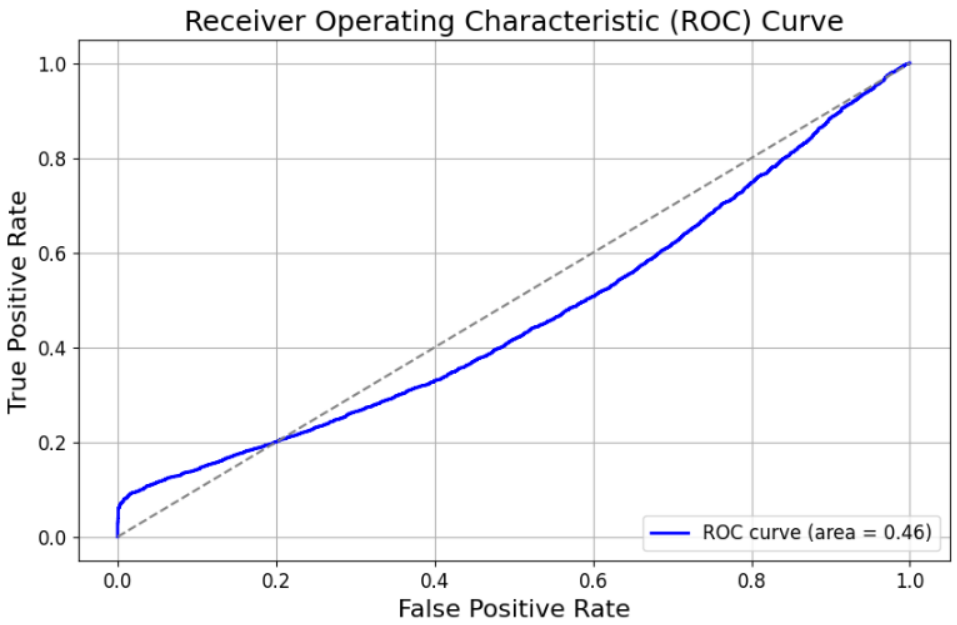

Figure 5.9 Receiver Operating Characteristic (ROC) Curve for Autoencoder Model

In general, the model achieved 72% accuracy, which is moderate given the dataset's significant class imbalance. The macro-average F1-score of 0.45 indicates poor performance in both categories, particularly for fraudulent transactions. The increased number of false positives indicates that the current threshold is too stringent. Modifying this threshold or taking a more sophisticated approach to its determination may improve the model's ability to detect both fraudulent and non-fraudulent transactions. Additionally, further model optimisation or feature engineering may reduce false positives and improve the model's ability to distinguish between the two classes.

5.5 Results overview (CNN, GAN and Autoencoder)

Table 5.2 summarises the performance metrics of the CNN, GAN, and Autoencoder models, including accuracy, precision, recall, F1-score, and ROC-AUC. The CNN model performed the best overall, with an accuracy of 95%, a precision of 0.96, a recall of 0.99, and an F1-score of 0.98. This validates CNN's suitability for immediate fraud detection. The GAN model also performed well, especially in terms of precision (0.99), but its recall was slightly lower at 0.92, resulting in an F1-score of 0.95. Despite its ability to detect anomalies, the Autoencoder model had significantly lower accuracy (72%) and ROC-AUC (0.46), indicating difficulties in distinguishing between fraudulent and non-fraudulent transactions.

Table 5.2 Overview of Model Performance Metrics

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC
CNN	0.95	0.96	0.99	0.98	–
GAN	0.91	0.99	0.92	0.95	–
Autoencoder	0.72	0.95	0.75	0.84	0.46

Figure 5.10 compares the principal metrics of the CNN, GAN, and Autoencoder models. The graph demonstrates that the CNN model consistently surpassed the others in accuracy, recall, and F1-score. Although GAN attained marginally superior precision, it fell short in the recall, affecting the overall balance. Despite its high precision, the Autoencoder model achieved the lowest scores across all other metrics, thus underlining its limitations in this research.

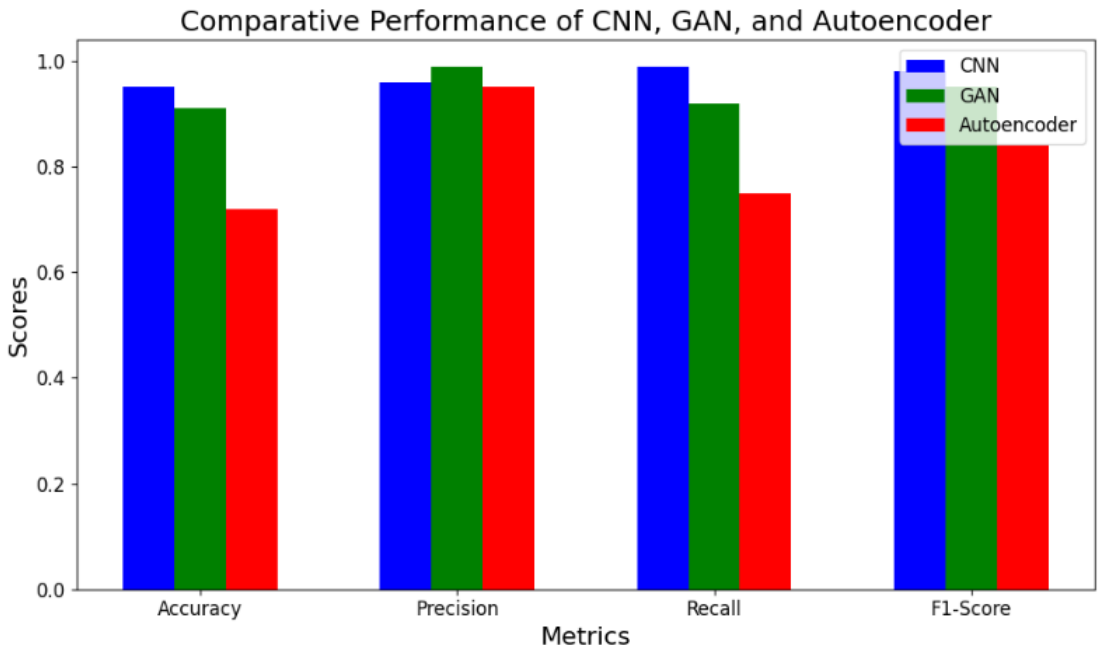

Figure 5.10 Comparative Performance of CNN, GAN, and Autoencoder

5.6 Models' strengths and limitations

The three deep learning models each have unique strengths and limitations. The CNN model demonstrated high classification accuracy, with a false positive rate of 0.038 and a false negative rate of 0.327. Although CNNs are ideal for fraud detection, their tendency to overfit, particularly in imbalanced datasets, necessitates careful tuning. The false negative rate shows that the model missed 32.7% of fraudulent instances, indicating that there is room for improvement in detecting elusive fraud.

GAN was highly effective at producing realistic synthetic data for fraud detection. The GAN model had a false positive rate of 0.009, indicating a high ability to prevent the incorrect classification of legitimate transactions as fraudulent. The false negative rate of 0.673 indicates difficulty in identifying many actual fraud cases. This emphasises the importance of balancing the adversarial dynamics between the generator and the discriminator to improve fraud detection efficacy.

The Autoencoder performed well in anomaly detection, but the model had a false positive rate of 0.250 and a false negative rate of 0.802, indicating that it could not accurately identify any fraudulent transactions. According to the literature review, Autoencoders excel at detecting anomalies, but their high false negative rate suggests a significant limitation in their ability to detect subtle fraud patterns in the datasets used in this research.

Table 5.3 shows each model's false positive and false negative rates, emphasising the trade-offs between their strengths and weaknesses in detecting fraudulent transactions.

Table 5.3 False Positive and False Negative Rates for CNN, GAN, and Autoencoder

Model	False Positive Rate	False Negative Rate
CNN	0.038	0.327
GAN	0.009	0.673
Autoencoder	0.250	0.802

5.7 Processing speed and resource utilisation assessment

5.7.1 Model inference time

Inference time was important in determining the feasibility of real-time fraud detection in Namibia's public sector financial transactions. Table 5.4 shows the average inference duration for each model. The results show that the CNN model, with an average inference time of 60 ms per transaction, is the fastest and best suited for real-time fraud detection in high-throughput environments. The Autoencoder model, with an inference time of 49 ms per transaction (4.22 seconds for 86,148 samples), is even faster, making it an efficient choice for such tasks. While capable of generating synthetic fraud data, the GAN model has an average inference time of 72 ms for every 1000 transactions (0.72 seconds), which translates to approximately 0.72 ms per transaction. Although the GAN operates more slowly than the CNN, it is still suitable for fraud detection in low-throughput scenarios or applications requiring minimal latency.

Table 5.4 Average Inference Time for Models

Model	Average Inference Time (ms/transaction)
CNN	60
GAN	7.17
Autoencoder	49

5.7.2 Resource utilisation analysis

The models' resource utilisation was assessed by evaluating memory consumption and computational efficiency as dataset sizes increased. The combined dataset consisted of 430,738 rows for the CNN, 492,762 for GAN, and 344,590 for the Autoencoder.

The CNN model exhibited exceptional resource efficiency, processing 1,000 rows in 2.16 seconds while consuming moderate memory (1.94 GB system RAM and 387 MB GPU RAM). The Autoencoder demonstrates similar efficiency, processing large datasets in 1.65 seconds per 1,000 rows with

comparable memory usage (2.32 GB system RAM / 385 MB GPU RAM). The GAN model processed data in 2.53 seconds per 1,000 rows but required significant memory during inference, using 7.2 GB of system RAM and 393 MB of GPU RAM, as summarised in Table 5.5.

Table 5.5 Resource Evaluation of Models

Model	Memory Usage (System RAM / GPU RAM)	Dataset Size Processed	Processing Time (per 1000 rows)
CNN	1.94 GB / 387 MB	430,738 rows	2.16 seconds
GAN	7.2 GB / 393 MB	492,762 rows	2.53 seconds
Autoencoder	2.32 GB / 385 MB	344,590 rows	1.65 seconds

These findings highlight the resource efficiency of the CNN and Autoencoder models, which can effectively process large volumes of data while maintaining low memory consumption. While the GAN model delivers competitive processing speeds (2.53 seconds per 1,000 rows), its higher memory demands make it more suitable for real-time fraud detection in environments where memory resources are less constrained.

5.8 Chapter summary

This chapter evaluated the effectiveness of CNN, GAN, and Autoencoder models in detecting fraud in Namibia's public sector financial transactions. The CNN model proved the most effective for real-time fraud detection, with a good balance of precision and recall. It demonstrated strong classification abilities, achieving 95% accuracy while maintaining a low false positive rate of 0.038 and a false negative rate of 0.327, making it ideal for efficiently detecting fraudulent transactions.

The GAN model demonstrated significant potential, particularly in producing realistic synthetic data to improve fraud detection. Despite achieving a high precision of 0.99 in fraud detection, its false negative rate of 0.673 indicated difficulties in identifying some fraudulent instances, implying that further optimisation is required to improve its ability to detect elusive fraud. Furthermore, the cosine similarity

and Euclidean distance metrics between authentic and synthetic data revealed the model's ability to generate valuable anomalies for future fraud detection applications.

The Autoencoder model performed well in anomaly detection; however, it had a false negative rate of 0.802, indicating that it could only detect fraudulent instances. Despite this limitation, the false positive rate of 0.250 indicates that it was reasonably conservative in misclassifying non-fraudulent transactions. This suggests that better tuning and a more impartial training methodology could improve Autoencoder performance in financial fraud detection.

Regarding resource utilisation, both the CNN and Autoencoder models demonstrated exceptional efficiency, effectively managing large datasets with minimal memory consumption and high processing speeds. While more memory-intensive, the GAN model remains a viable option for large-scale applications, particularly with potential optimisations to enhance its efficiency.

In short, the CNN model outperformed the other models in accuracy and efficiency, whereas GAN and Autoencoder showed complementary strengths useful in specific contexts or hybrid models. These findings lay the groundwork for developing a scalable and efficient fraud detection system for Namibia's public sector financial transactions and identify areas for improvement, particularly in terms of reducing false negatives and optimising the balance of precision and recall.

CHAPTER 6 : CONCLUSION AND FUTURE WORK

This final chapter summarises the study's key findings, returning to the research questions to highlight their significance and contributions to the field. It also looks into the study's limitations and makes practical recommendations. Finally, potential future research directions are proposed, laying the groundwork for further investigation and advancement of deep-learning financial fraud detection systems.

6.1 Revisiting the research questions

This thesis presented a deep learning-based automated fraud detection system for Namibia's public sector financial transactions. Three standard models (CNN, GAN, and Autoencoder) were developed and thoroughly evaluated for classification accuracy, inference time, and resource utilisation, with the results indicating each approach's strengths and limitations. The CNN model emerged as the best-performing option and was ultimately selected for use in the fraud detection system, demonstrating its effectiveness in identifying fraudulent transactions.

The feature importance analysis further supported the system's development by identifying critical financial indicators for fraud detection. The presence of multiple general ledger codes for the same vendor emerged as the most significant feature, strongly indicating possible fraud. Repeated invoices and mismatched transaction amounts were also important indicators of fraud. While features such as unusual purchase patterns, inconsistencies in bank account usage, and massive amounts had a lower significance, they remained relevant. These findings emphasise the importance of prioritising high-impact features to improve the fraud detection system’s accuracy and reliability.

The resource utilisation analysis highlighted critical aspects of the system's efficiency in processing large volumes of financial transactions. The CNN model achieved the best balance of computational efficiency and memory usage, quickly processing transactions using minimal system and GPU memory. The Autoencoder model performed similarly, processing data efficiently but with slightly higher

memory consumption, whereas the GAN model required more memory, indicating that it is best suited for fraud detection in low-throughput environments.

In terms of performance, the CNN model outperformed the other models by achieving the highest accuracy (95%) while maintaining precision (0.96) and recall (0.99), making it the best choice for real-time fraud detection. It also demonstrated efficient resource utilisation and fast inference times, making it suitable for use in large-scale financial transaction systems. However, the CNN had a moderate false negative rate (0.327), indicating that some fraudulent transactions were missed and could be avoided with additional tuning. The GAN model similarly demonstrated promise, particularly for generating synthetic data, with a high precision of 0.99. However, its false negative rate of 0.673 indicates that it missed many fraud cases, highlighting the difficulty in fully capturing complex fraud patterns. Despite this, GAN-generated synthetic data can potentially improve fraud detection systems by augmenting datasets and increasing model generalisation. On the other hand, while Autoencoder is effective at detecting anomalies, it has a false negative rate of 0.802, indicating that it can only detect fraudulent transactions while missing many non-fraudulent transactions. Despite this, the false positive rate of 0.250 suggests that it was reasonably cautious in misclassifying non-fraudulent transactions. These limitations demonstrate that Autoencoder requires additional tuning and a more balanced training methodology to perform better in real-world fraud detection scenarios.

In summary, the findings support the notion that a multi-model approach combining the strengths of CNN and GAN could provide a more robust and scalable fraud detection system. In contrast, an Autoencoder may require further refinement to be effective in this setting.

6.2 Contributions to the field

This study contributes to the field by presenting a deep learning framework for automated fraud detection in public sector financial transactions using CNN, GAN, and Autoencoder models. CNN's superior accuracy and resource utilisation efficiency demonstrate its effectiveness in instantaneous fraud detection. At the same time, the use of GAN-generated synthetic data suggests the potential for

improving model training when fraudulent data is scarce. The findings show that a hybrid methodology combining CNN and GAN may improve detection effectiveness, which has important implications for developing more resilient fraud detection systems with the potential of incorporating technologies like blockchain to improve financial systems' security and transparency.

6.3 Limitations

The main limitation of this study was the dataset's imbalance, with a significantly lower number of fraudulent transactions than legitimate ones, which may have impacted the generalisability of the models, particularly the Autoencoder and GAN, resulting in higher false negative rates. Furthermore, while the models were tested for scalability using custom Python scripts, the testing environment may not fully reflect the complexities of real-world deployment with fluctuating transaction volumes, and the GAN model's resource-intensive nature (7.2 GB memory usage) may present challenges for large-scale applications. The use of Python scripts also limited the ability to test the system's full-scale performance, as more advanced tools such as Gatling could provide more detailed insights into scalability and performance. Furthermore, while the GAN's synthetic data was useful, it had difficulties replicating certain critical fraud patterns, potentially reducing its overall effectiveness in detecting fraud. Finally, the evaluation focused on key metrics such as accuracy and precision while ignoring aspects such as interpretability, which are critical in high-risk environments such as public sector financial systems.

6.4 Future work

The findings of this research not only address the immediate objectives but also point to several areas for future research and development, highlighting opportunities to deepen understanding, refine applications, and explore broader implications. These possibilities are detailed below:

- Model Optimisation: To reduce false negative rates in CNN and GAN models, consider optimising hyperparameters and regularisation techniques.

- Improved GAN-Generated Data: While GAN was effective at generating synthetic data, there were discrepancies in replicating critical fraud detection features. Future research could focus on improving the GAN's ability to capture these complex patterns, possibly through advanced GAN architectures or training with more data.
- Integration of Hybrid Models: Combining CNN classification with GAN-generated synthetic data can improve detection rates. A hybrid approach that combines the strengths of both models could result in a system capable of detecting more subtle fraud cases while remaining scalable and efficient.
- Blockchain Integration: Integrating blockchain technology could help ensure data integrity by providing a transparent and immutable audit trail, thereby increasing trust in public financial systems. Future research could look into how blockchain can work with fraud detection models to improve security.
- Further Testing: While the models demonstrated efficient resource utilisation and feasibility in simulated environments, deploying the system in a real-world financial environment will provide valuable insights into its practical application. Further testing with real transaction data, including continuous model retraining, may improve the system's robustness and adaptability to changing fraud patterns.
- Exploring Transfer Learning: Future research could look into transfer learning from other domains with extensive fraud detection research to better understand financial fraud. Pre-trained models can help capture generalised fraud patterns, reducing the need for extensive retraining on sector-specific data.
- Exploration of Unsupervised and Reinforcement Learning: Further research could greatly benefit from a more thorough investigation of unsupervised learning techniques and their potential integration with reinforcement learning. By leveraging these approaches, we could enhance anomaly detection capabilities, allowing models to recognise fraud patterns with minimal labelled data and to adapt dynamically to evolving fraud tactics.

In conclusion, this thesis lays the groundwork for automated fraud detection in Namibia's public sector, with CNN serving as a solid foundation and GAN providing valuable synthetic data generation. Future efforts should concentrate on improving model performance, integrating hybrid approaches, and testing in real-world scenarios to improve the system's ability to detect fraudulent activities.

REFERENCES

AICPA. (2007). *AICPA audit and accounting manual as of July 1, 2007: Nonauthoritative technical practice aid.* In K. R. Biser, C. Cole, K. L. Illuzzi, & L. L. Pombo, Eds.) American Institute of Certified Public Accountants. https://egrove.olemiss.edu/aicpa_guides/967

Alghofaili, Y., Albattah, A., & Rassam, M. A. (2020). A financial fraud detection model based on the LSTM deep learning technique. *Journal of Applied Security Research, 15*(4), 498-516. doi:https://doi.org/10.1080/19361610.2020.1815491

Al-Hashedi, K. G., & Magalingam, P. (2021, May). Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019. *Computer Science Review, 40*, 100402. doi:https://doi.org/10.1016/j.cosrev.2021.100402

Ali, J., Khan, R., Ahmad, N., & Maqsood, I. (2012). Random Forests and decision trees. *International Journal of Computer Science Issues (IJCSI), 9*(5), 272. www.ijcsi.org

Alvi, M. H. (2016). *A manual for selecting sampling techniques in research.* Munich Personal RePEc Archive. https://mpra.ub.uni-muenchen.de/70218/1/MPRA_paper_70218.pdf

Ashtiani, M. N., & Raahemi, B. (2021). Intelligent fraud detection in financial statements using machine learning and data mining: A systematic literature review. *IEEE Access, 10*, 72504-72525. doi:10.1109/ACCESS.2021.3096799

Bakumenko, A., & Elragal, A. (2022). Detecting anomalies in financial data using machine learning algorithms. *Systems, 10*(5), 30. doi:https://doi.org/10.3390/systems10050130

Bhavitha, B. K., Rodrigues, A. P., & Chiplunkar, N. N. (2017). Comparative study of machine learning techniques in sentiment analysis. *In 2017 International conference on inventive communication and computational technologies (ICICCT)* (pp. 216-221). Institute of Electrical and Electronics Engineers (IEEE). doi:https://doi.org/10.1109/ICICCT.2017.7975191

Bini, S. A. (2018). Artificial intelligence, machine learning, deep learning, and cognitive computing: What do these terms mean, and how will they impact health care? *The Journal of Arthroplasty, 38*(8), 2358-2361. doi:https://doi.org/10.1016/j.arth.2018.02.067

Birmingham, P., & Wilkinson, D. (2003). *Using research instruments: A guide for researchers* (1 ed.). Routledge. doi:https://doi.org/10.4324/9780203422991

Bloomenthal, A. (2021). *Detecting financial statement fraud.* Investopedia. https://www.investopedia.com/articles/financial-theory/11/detecting-financial-fraud.asp

Breiman, L. (2001, October). Random forests. *Machine Learning, 45*, 5-32. doi:https://doi.org/10.1023/A:1010933404324

Byrd, W. A., & Guimbert, S. (2009). Public finance, security, and development: A framework and an application to Afghanistan. *World Bank Policy Research Working Paper*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1327275

Chen, S., Goo, J. Y.-J., & Shen, Z.-D. (2014). A hybrid approach of stepwise regression, logistic regression, support vector machine, and decision tree for forecasting fraudulent financial statements. *The Scientific World Journal, 2014*. doi:https://doi.org/10.1155/2014/968712

Craja, P., Kim, A., & Lessmann, S. (2020). Deep learning for detecting financial statement fraud. *Decision Support Systems, 139*, 113421. doi:https://doi.org/10.1016/j.dss.2020.113421

Creswell , J. W. (2013). *Qualitative inquiry & research design: Choosing among five approaches* (3rd ed.). In M. Masson (Ed.). Lincoln: SAGE Publications, Inc.

Deng, Q. (2010). Detection of fraudulent financial statements based on a Naïve Bayes classifier. *2010 5th International Conference on Computer Science & Education* (pp. 1032-1035). Hefei, China: IEEE. doi:10.1109/ICCSE.2010.5593407

Deng, Q., & Mei, G. (2009). Combining self-organising map and K-means clustering for detecting fraudulent financial statements. *2009 IEEE International Conference on Granular Computing* (pp. 126-131). IEEE. doi:10.1109/GRC.2009.5255148

Devi, V. J., & Kavitha, K. S. (2017). Fraud detection in credit card transactions by using classification algorithms. *2017 International Conference on Current Trends in Computer, Electrical, Electronics and Communication (CTCEEC)* (pp. 125-131). Mysore. doi:10.1109/CTCEEC.2017.8455091

DISCOVERPHDS. (2023, 10 6). *Community blog*. https://www.discoverphds.com/blog/research-instrument

Dudovskiy, J. (2022). *The ultimate guide to writing a dissertation in business studies: A step-by-step assistance* (6 ed.). Research-methodology.net.

Faraji, Z. (2022). A review of machine learning applications for credit card fraud detection with a case study. *SEISENSE Journal of Management, 5*(1), 49-59. doi: https://orcid.org/0000-0002-3264-5378

Gatling Corp. (2023). *Documentation*. https://gatling.io/docs/gatling/

Gupta, A., Lohani, M. C., & Manchanda, M. (2021, September 02). Financial fraud detection using naive bayes algorithm in a highly imbalanced dataset. *Journal of Discrete Mathematical Sciences and Cryptography, 24*(5), 1559–1572. doi:https://doi.org/10.1080/09720529.2021.1969733

Gupta, S., & Mehta, S. (2021). Data mining-based financial statement fraud detection: Systematic literature review and meta-analysis to estimate data sample mapping of fraudulent companies against non-fraudulent companies. *Global Business Review*. doi:10.1177/0972150920984857

Hamal, S., & Senvar, O. (2021). Comparing performances and effectiveness of machine learning classifiers in detecting financial accounting fraud for Turkish SMEs. *International Journal of Computational Intelligence Systems, 14*(1), 769-782. https://www.atlantis-press.com/journals/ijcis/issue/498

HaratiNik, M. R., Akrami, M., Khadivi, S., & Shajari, M. (2012). A hybrid model for credit card fraud detection. *6th International Symposium on Telecommunications (IST)* (pp. 1088-1093). Tehran, Iran: IEEE. doi:10.1109/ISTEL.2012.6483148

Hilal, W., Gadsden, A. S., & Yawney, J. (2022, May 01). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications, 193*. doi:https://doi.org/10.1016/j.eswa.2021.116429

Huang, Z., Zheng, H., Li, C., & Che, C. (2024, March). Application of machine learning-based K-Means clustering for financial fraud detection. *Academic Journal of Science and Technology, 10*(1), 33-39. doi:https://doi.org/10.54097/74414c90

Jeragh, M., & AlSulaimi, M. (2018, October). Combining auto encoders and one-class support vector machine for fraudulent credit card transactions detection. *2018 Second World Conference on Smart Trends in Systems, Security and Sustainability (WorldS4)*, 178-184. doi:10.1109/WorldS4.2018.8611624

Kirlidog, M., & Asuk, C. (2012, October 2024). A fraud detection approach with data mining in health insurance. *Procedia - Social and Behavioral Sciences, 62*, 989-994. doi:https://doi.org/10.1016/j.sbspro.2012.09.168

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*(7553), 436–444. doi:https://doi.org/10.1038/nature14539

Makki, S., Assaghir, Z., Taher, Y., Haque, R., Hacid, M.-S., & Zeineddine, H. (2019, July). An experimental study with imbalanced classification approaches for credit card fraud detection. *IEEE Access, 7*, 93010-93022. doi:10.1109/ACCESS.2019.2927266

Mareeswari, V., & Gunasekaran, G. (2016, February). Prevention of credit card fraud detection based on HSVM. *2016 International Conference on Information Communication and Embedded Systems (ICICES)*, 1-4. doi:10.1109/ICICES.2016.7518889

MDPI Books. (2023). *Special issue reprint: Applied machine learning.* Applied Sciences. doi:https://doi.org/10.3390/books978-3-0365-7907-8

Noja, G. G., Cristea, M., Sirghi, N., Hategan, C.-D., & Paolo, D. (2019). Promoting good public governance and environmental support for sustainable economic development. *International Journal of Environmental Research and Public Health, 16*(24), 4940. https://www.mdpi.com/1660-4601/16/24/4940

Paavo, J. P., Rodríguez-Puentes, R., & Maliwatu, R. (2024). Exploring machine learning fraud detection solutions for financial transactions. *International Engineering Conference on Sustainable Emerging Innovations and Technological Advancements.* https://ejournal.bumipublikasinusantara.id/index.php/ajeee/article/view/608/0

Paula, E. L., Ladeira, M., Carvalho, R. N., & Marzagão, T. (2016). Deep learning anomaly detection as support for fraud investigation in Brazilian exports and anti-money laundering. *In 2016, 15th IEEE International Conference on machine learning and Applications (ICMLA)* (pp. 954-960). Institute of Electrical and Electronics Engineers (IEEE). doi:10.1109/ICMLA.2016.0172

Peng, H., & You, M. (2016). The health care fraud detection using the pharmacopoeia spectrum tree and neural network analytic contribution hierarchy process. *2016 IEEE Trustcom/BigDataSE/ISPA* (pp. 2006-2011). Tianjin, China: IEEE. doi:10.1109/TrustCom.2016.0306

Rajak, I., & Mathai, J. K. (2015, September). Intelligent fraudulent detection system based on SVM and optimised by danger theory. *International Conference on Computer, Communication and Control (IC4)*, 1-4. doi:10.1109/IC4.2015.7375705

Rivas, M. (2023, August 30). Deputy director of information systems audits at the Office of the Auditor-General. (J. P. Paavo, Interviewer) Windhoek, Namibia.

Rusk, N. (2016). Deep learning. *Nature Methods, 13*(1), 35-35. doi:https://doi.org/10.1038/nmeth.3707

Sanad, Z., & Al-Sartawi , A. (2021). Financial statements fraud and data mining: A review. *In the European, Asian, Middle Eastern, and North African Conference on Management & Information Systems. 239*, pp. 407-414. Springer International Publishing. https://link.springer.com/chapter/10.1007/978-3-030-77246-8_38

Saunders, M., Lewis, P., & Thornhill, A. (2023). *Research methods for business students* (9 ed.). Pearson.

Sawsan, A., & Jaradat, R. (2018). Clarification of research design, research methods, and research methodology: A guide for public administration researchers and practitioners. *Teaching Public Administration, 36*(3), 237-258. https://journals.sagepub.com/doi/abs/10.1177/0144739418775787

Sheshasayee, A., & Thomas, S. S. (2017, February). Implementation of data mining techniques in upcoding fraud detection in the monetary domains. *2017 international conference on innovative mechanisms for industry applications (ICIMIA*, 730-734. doi:10.1109/ICIMIA.2017.7975561

Singh, A., & Jain, A. (2020, February 02). An empirical study of aml approach for credit card fraud detection—Financial transactions. *International Journal of Computers Communications & Control, 14*(6), 670-690. doi:10.15837/ijccc.2019.6.3498

Stalebrink, O. J., & Sacco, J. F. (2007). Rationalisation of financial statement fraud in government: An Austrian perspective. *Critical Perspectives on Accounting, 18*(4), 489-507. https://www.sciencedirect.com/science/article/abs/pii/S1045235406000190

Sudjian, A., Yuan, M., Kern, D., Nair, S., Zhang, A., & Cela-Díaz, F. (2010, February). Statistical methods for fighting financial crimes. *Technometrics, 52*(1), 5-19. https://www.jstor.org/stable/40586676

Sulaiman, R. B., Schetinin, V., & Sant, P. (2022, May 05). Review of machine learning approach on credit card fraud detection. *Human-Centric Intelligent Systems, 2*(1), 55-68. doi:https://doi.org/10.1007/s44230-022-00004-0

Sundarkumar, G. G., Ravi, V., & Siddeshwar, V. (2015, December). One-class support vector machine-based undersampling: Application to churn prediction and insurance fraud detection. *2015 IEEE International Conference on Computational Intelligence and Computing Research (ICCIC)*, 1-7. doi:10.1109/ICCIC.2015.7435726

Thiprungsri, S., & Vasarhelyi, M. A. (2011). Cluster analysis for anomaly detection in accounting data: An audit approach. *The International Journal of Digital Accounting Research, 11*, 69-84. https://core.ac.uk/download/pdf/60648535.pdf

Wang, X., Wu, H., & Yi, Z. (2018). Research on bank anti-fraud model based on K-Means and Hidden Markov Model. *2018 IEEE 3rd International Conference on Image, Vision and Computing (ICIVC)* (pp. 780-784). Chongqing, China: IEEE. https://ieeexplore.ieee.org/document/8492795

Wang, Y., & Xu, W. (2018, January). Leveraging deep learning with LDA-based text analytics to detect automobile insurance fraud. *Decision Support Systems, 105*, 87-95. doi:https://doi.org/10.1016/j.dss.2017.11.001

West, J., & Bhattacharya, M. (2016, March). Intelligent financial fraud detection: A comprehensive review. *Computers & Security, 57*. doi:https://doi.org/10.1016/j.cose.2015.09.005

Yao, J., Zhang, J., & Wang, L. (2018). A financial statement fraud detection model based on hybrid data mining methods. *In 2018 International Conference on Artificial Intelligence and Big Data (ICAIBD)* (pp. 57-61). Institute of Electrical and Electronics Engineers (IEEE). doi:10.1109/ICAIBD.2018.8396167

Yoshua, B., Goodfellow, I., & Courville, A. (2017). *Deep Learning* (Vol. 1). MIT Press.

Youness, A., Lahby, M., & Attioui, A. (2018). An efficient real-time model for credit card fraud detection based on deep learning. *In Proceedings of the 12th International Conference on Intelligent Systems: Theories and Applications* (pp. 1-7). Association for Computing Machinery. doi:https://doi.org/10.1145/3289402.3289530

APPENDICES

Appendix A: ETHICAL CLEARANCE

NAMIBIA UNIVERSITY OF SCIENCE AND TECHNOLOGY

FACULTY RESEARCH ETHICS COMMITTEE (F-REC)

DECISION/FEEDBACK ON THE RESEARCH PROPOSAL

Dear **Paavo, Johannes Pandeni (223128481)**

RESEARCH TOPIC: AUTOMATED FRAUD DETECTION IN NAMIBIA'S PUBLIC INSTITUTIONS FINANCIAL TRANSACTIONS USING MACHINE LEARNING: A DEEP LEARNING APPROACH

Supervisor (if applicable): Dr R. Maliwatu

Qualification registered for (if applicable): Master of Data Science

(Reference number of applications: **FACULTY RESEARCH ETHICS COMMITTEE REGISTRATION NUMBER: FREC - 14/24)**

Re: Ethical screening application No: **FREC - 14/24**

The Faculty of **Computing and Informatics** Ethics Screening Committee of the Namibia University of Science and Technology reviewed your application for the above-mentioned research. The research as set out in the application has been:

Approved	X

(Indicate with an X, and N/A if not applicable and proceed)

We would like to point out that you, as a researcher, are obliged to maintain the ethical integrity of your research, adhere to the ethical guidelines of NUST, and remain within the scope of your research proposal and supporting evidence as submitted to the F-REC. Should any aspect of your research change from the information as presented to the F-REC, which could affect the possibility of harm to any research subject, you are under the obligation to report it immediately to your supervisor or F-REC as applicable in writing. Should there be any uncertainty in this regard, you must consult with the F-REC.

We wish you success with your research and trust that it will make a positive contribution to the questfor knowledge at NUST.

Any ethical issues that need to be highlighted?	Why are these issues important?	What must/could be done to minimize the ethical risk?
No	N/A	N/A

Recommendation: The application is approved.

Sincerely,

Prof. Suama L Hamunyela
Chairperson: Faculty Ethics Screening
CommitteeTel: +264-61-207-2922
CC: Co-supervisor: None

NAMIBIA UNIVERSITY OF SCIENCE AND TECHNOLOGY
2024 -04- 11
OFFICE OF THE DEAN
FACULTY OF COMPUTING & INFORMATICS

APPENDIX B: OFFICE OF THE AUDITOR-GENERAL CONSENT LETTER

Republic of Namibia

OFFICE OF THE AUDITOR-GENERAL

Tel: (264) (061) 2858000
Fax: (264) (061) 22430
WINDHOEK
www.oag.gov.na

Private Bag 13299

9000

Our Ref: *Your Ref:* *Enquiries:*

12 February 2024

TO WHOM IT MAY CONCERN

CONSENT FOR ACCESS AND USE OF FINANCIAL RECORDS FOR RESEARCH PURPOSES

This letter serves to grant permission to Johannes Paavo, a final-year Master of Data Science student at Namibia University of Science and Technology (NUST), to access and use relevant financial records from our office for his research project titled "Automated Fraud Detection in Namibia's Public Institutions Financial Transactions Using Machine Learning."

The data provided will be used exclusively for academic purposes, with strict anonymization procedures in place to ensure the confidentiality of sensitive information.

Your cooperation in this matter is appreciated.

Yours sincerely,

DEPUTY AUDITOR-GENERAL OFFICE NAMIBIA 0 FEB 2024 Private Bag 13299

12/02/24

Mr. G. MENETTE **Date**

DEPUTY AUDITOR-GENERAL

APPENDIX C: ETHICS INFORMED CONSENT FORM

Informed Consent Form for Auditors of the Office of the Auditor General, Namibia.

This informed consent form is for auditors who are accountable for auditing the financial statements of Namibia's public institutions and who are invited to participate in this academic research titled "Automated Fraud Detection in Namibia's Public Institutions Financial Transactions Using Machine Learning: A Deep Learning Approach".

Name of Principal Investigator: Mr. Johannes Pandeni Paavo
Name of Organisation: Namibia University of Science and Technology
Name of Sponsor: None
Name of Project and Version: Automated Fraud Detection in Namibia's Public Institutions Financial Transactions Using Machine Learning: A Deep Learning Approach.

This Informed Consent Form has two parts:

- **Information Sheet (to share information about the study with you)**
- **Certificate of Consent (for signatures if you choose to participate)**

You will be given a copy of the full Informed Consent Form.

Part I: Information Sheet

Introduction

I am a Master of Data Science student working for the Office of the Auditor-General, and I am inviting you to take part in a research project aimed at developing a robust deep learning-based automated fraud detection system for Namibia's public sector financial transactions. Before deciding whether to participate, you must understand the study's purpose, procedures, risks, benefits, and confidentiality. Please read the following information thoroughly and ask any questions you may have before making your decision. You also do not have to decide whether or not to participate in the research today, and you can discuss it with anyone you feel comfortable with.

Purpose of the research

The primary aim of the research is to automate fraud detection in financial records using advanced deep-learning techniques, specifically tailored for Namibia's public sector financial transactions. This system will be a cutting-edge tool for detecting instances of fraudulent activity in the financial records of these public institutions. We believe that you can either help us identify relevant financial features and indicators for fraud detection or help with the investigation of scaling challenges and opportunities for the system handling a large volume of public institutions' financial records.

Type of Research Intervention

This research project will necessitate your active participation in an in-depth interview session that will last approximately one hour. Your valuable insights and contributions during this interview will be critical in furthering our understanding of the subject and will significantly contribute to the success of this research endeavour.

Participant Selection

We are inviting you to participate in this research because we believe your experience as an Auditor at the Office of the Auditor-General, an office tasked with auditing the financial statements of Namibian public institutions, can significantly contribute to our understanding and knowledge of fraud in public institution financial records.

Voluntary Participation

Participation in this study is completely voluntary and you are free to withdraw at any time without any negative consequences. Your decision regarding participation will also not affect your relationship with the researcher or the institution.

Procedures

If you agree to participate, the researcher will interview you about unique financial indicators specific to Namibia's public sector that should be considered when detecting fraud. We will also explore the scaling challenges and opportunities that the fraud detection system may present when dealing with large volumes of financial statements from public institutions. The interview will take place in any available boardroom at the time and will last approximately one hour. Please be assured that your responses will be kept strictly confidential and used only for the purposes of this research. If you decide not to answer a question during the interview, simply express your preference, and the interviewer will move on to the next question. During the interview, only the interviewer will be present unless you wish to have someone else accompany you. Your information will be treated with the utmost confidentiality, and no one else except the researcher will have access to the information documented during your interview. The entire interview will be tape-recorded for accuracy and research purposes, but no one will be identified by name on the tape to ensure your privacy. The tapes will be securely stored on a password-protected phone to safeguard your information. These recordings will be retained for a period of 20 weeks as necessary for research purposes and then will be permanently deleted.

Duration

The research project is expected to last 8 months in total. Throughout this time period, follow-up interviews will be conducted on an ad hoc basis, based primarily on your availability and subject to your

prior request. Each interview session will last approximately one hour. These interviews will be scheduled at your convenience and in accordance with your advance request.

Risks

There are no known risks associated with participating in this study, and you have complete autonomy to choose not to respond to any question that makes you uncomfortable or choose not to answer.

Benefits

There will be no direct benefit to you, but your participation is likely to assist us in identifying and understanding relevant financial features and indicators for fraud detection, as well as in the investigation of scaling challenges and opportunities for the system handling a large volume of public institutions' financial statements, how they might align with existing literature, and how they might contribute to scientific knowledge about concerns associated with land administration.

Reimbursements

Participants in this study will receive no monetary incentives or compensation for their participation. Nonetheless, as a thank you for your time and participation, we will gladly offer you a bottle of water and some confectionery items during the interview.

Confidentiality

Your participation in the interview process is strictly confidential and anonymous. All data will be securely stored, ensuring that no personal information about the participants is disclosed in any way that could potentially reveal their identities when the findings from the interviews are disseminated.

Sharing the Results

Concerning the dissemination of our findings, we are deeply committed to keeping you informed. Once the research is completed, we intend to provide each participant with a comprehensive summary of the findings. Our commitment to sharing our discoveries with academic and professional circles through publications and conference presentations demonstrates our strong belief in transparency and community engagement. Please know that any information you provide during the research will be kept strictly confidential, and your anonymity will be strictly maintained.

Right to Refuse or Withdraw

This is a reconfirmation that your participation in this study is entirely voluntary, and you are free to withdraw at any time with no negative consequences. You will have the opportunity to review your comments near the end of the interview, and you may request changes or removal of any portions if you disagree with my notes or believe there was a misunderstanding. Please know that your decision to participate or not will have no bearing on your relationship with the researcher or the institution.

Who to Contact

Namibian University of Science and Technology, Faculty of Computing and Informatics Research Ethics Committee responsible for safeguarding the well-being of research participants, has thoroughly reviewed and granted approval for this research study. If you have any questions about this study before or after participation, please feel free to contact the researcher, Johannes Pandeni Paavo, at j3paavo@gmail.com/+264818797347 or the HoD of Informatics at mmaravanyika@nust.na.

This proposal has been reviewed and approved by the Faculty of Computing and Informatics Research Ethics Committee, which is a committee whose task it is to make sure that research participants are protected from harm. If you wish to find out more about the REC, contact Dr Munyaradzi Maravanyika at mmaravanyika@nust.na or at +264 61 207 2263. It has also been reviewed by the Research and Ethic Committee of NUST, which is supporting the study.

Part II: Certificate of Consent

(This section is mandatory)

I have read the foregoing information, or it has been read to me. I have had the opportunity to ask questions about it and any questions I have asked have been answered to my satisfaction. I consent voluntarily to be a participant in this study.

Print Name of Participant_____

Signature of Participant _____

Date _____

Day/month/year

If illiterate [1]

I have witnessed the accurate reading of the consent form to the potential participant, and the individual has had the opportunity to ask questions. I confirm that the individual has given consent freely.

Print name of witness_____ Thumbprint of participant

Signature of witness _____

Date _____

Day/month/year

Statement by the researcher/person taking consent

_____

[1] A literate witness must sign (if possible, this person should be selected by the participant and should have no connection to the research team). Participants who are illiterate should include their thumb print as well.

I have accurately read out the information sheet to the potential participant, and to the best of my ability, made sure that the participant understands that the following will be done:

1. Will be interviewed voluntarily.
2. I will be allowed to make changes or remove any portions of my notes if they disagree or believe there was a misunderstanding.
3. Will be provided with a comprehensive summary of the findings.

I confirm that the participant was allowed to ask questions about the study, and all the questions asked by the participant have been answered correctly and to the best of my ability. I confirm that the individual has not been coerced into giving consent, and the consent has been given freely and voluntarily.

A copy of this ICF has been provided to the participant.

Print Name of Researcher/person taking the consent_____

Signature of Researcher /person taking the consent_____

Date _____

Day/month/year

APPENDIX D: INTERVIEW GUIDE

This guide was utilised to conduct interviews with auditors specialising in government financial and information systems auditing to obtain valuable qualitative insights into the practices and challenges of financial fraud detection in Namibia's public sector. The duration of the in-depth interview is intended to last approximately one hour, although it may differ depending on the respondent's answers and the need for additional questions to obtain further clarification.

Introductory Questions:

1 Could you please tell me about your role and responsibilities as an auditor at the Office of the Auditor-General?
2 How many years have you spent auditing the financial transactions of public institutions?
3 Could you provide some background on the audit process for financial transactions of public institutions in Namibia, including the key objectives and challenges you face in your work?

Sub-Objective 1: Identify relevant financial features and indicators for fraud detection.

1 Can you give an overview of the typical financial transactions examined during the audit process of Namibian public institutions?
2 . In your experience, what are the most common financial indicators or features that auditors look for when determining the presence of fraud in these financial records?
3 Could you please describe any specific patterns or anomalies in financial data that auditors frequently discover to be indicative of fraud in public institutions?
4 . What tools or methodologies do you currently use to detect and investigate financial transaction fraud?
5 Are there any financial features or indicators that are unique to Namibian public institutions that should be considered when detecting fraud?
6 How important do you think data quality and accuracy are in detecting financial fraud, and what challenges do auditors face in this regard?

Sub-Objective 2: Investigate resource utilisation challenges and opportunities for the system handling a large volume of financial transactions from public institutions.

1 . How does your current system handle large volumes of financial transactions, and what challenges do you face with system resources?

2 How do large datasets affect the efficiency of the system you use to manage financial transactions in audits?

3 Do you currently use any strategies or tools to optimise system resources when processing large amounts of financial data?

4 What improvements or changes would you recommend to better manage resource utilisation in your system when dealing with large-scale financial transactions?

Sub-Objective 3: Evaluate the developed system's fraud detection performance using historical and simulated data.

1 . How effective is your current system for detecting fraudulent financial transactions based on historical data?

2 . What challenges do you encounter when using historical data to evaluate the accuracy of fraud detection in financial transactions?

3 How does your system perform when detecting potential fraud with simulated data versus real-world data?

4 . What improvements would you recommend to improve your system's fraud detection performance based on your experience with both historical and simulated data?

Final Question:

1 Is there anything else you'd like to say or any additional information you think we should know about you or anything we did not cover during this interview?

APPENDIX E: AUTHORSHIP CERTIFICATE

UNAM INTERNATIONAL ENGINEERING CONFERENCE 2024

CERTIFICATE OF PAPER AUTHORSHIP

This is to certify that

Johannes P. Paavo, Rafael Rodriguez-Puentes, Richard Maliwatu

Jointly authored the paper titled

Exploring Machine Learning Fraud Detection Solutions for Financial Transactions

Which was presented by
Johannes P. Paavo
at the
University of Namibia International Engineering Conference on Sustainable Emerging Innovations and Technological Advancements (UNAM-IEC24)

Held on 02 – 04 December 2024

We gratefully acknowledge the contribution made by the authors to the success of the conference

Dr. Leokadia N. P. Ndjuluwa
Conference Chair
UNAM School of engineering and the built environment

Prof. Chinwuba Arum
Conference Editor-In-Chief
UNAM School of engineering and the built environment

APPENDIX F: LANGUAGE EDITING CERTIFICATE

ACET Consultancy
Anenyasha Communication, Editing and Training
Box 50453 Bachbrecht, Windhoek, Namibia
Cell: +264814218613
Email: mlambons@yahoo.co.uk

16 April 2025

To Whom It May Concern

LANGUAGE EDITING – JOHANNES PANDENI PAAVO

This letter serves to confirm that a research project titled *AUTOMATED FRAUD DETECTION IN NAMIBIA'S PUBLIC INSTITUTIONS' FINANCIAL TRANSACTIONS USING MACHINE LEARNING: A DEEP LEARNING APPROACH* was submitted to me for language editing.

The research was professionally edited, and track changes and suggestions were made in the document. The research content or the author's intentions were not altered during the editing process, and the author has the authority to accept or reject my suggestions.

Yours faithfully

PROF. (DR) NELSON MLAMBO
PhD in English
M.A. in Intercultural Communication
M.A. in English
B. A. Special Honours in English – First class
B. A. English & Linguistics