

Heterogeneous Graph Networks and Behavioural Sequence Modelling: AI-Driven Computational Frameworks for Financial Fraud Detection

Dr. Andreas Papadopoulos, Associate Professor of Electrical and Computer Engineering, National Technical University of Athens, Greece

1. Introduction

While the world has seen some significant technological advancements during the last two decades in the field of finance, there has also been a parallel increase in fraudulent activities in financial organizations. The advent of complex financial products and changing financial regulations has only made matters worse. Traditional fraud detection methods in financial services, which were once considered adequate, have been failing for some time now. Fraudsters are using more sophisticated methods to commit fraud, resulting in huge losses to financial institutions. Financial institutions have resorted to integrating artificial intelligence in the development of fraud detection systems to enhance the exploration and exploitation of underlying powerful data representatives and improve fraud detection to a great extent. Herein lies the objective of our research. The capabilities of financial systems have multiple definitions. In the world of open banking, innovation towards the reduction and management of digital financial services hinges on customer data. Those organizations with access to a significant amount of customer transaction data are hard to manage and secure against the growing prevalence of financial crimes such as money laundering, terrorist financing, and, of course, fraud. These innovative approaches are pivotal when considering increasing the amount of money that can be invested in financial services without further risk, as the industry continues to gain more and more money. This is why AI-driven financial crime and fraud reduction must become a central feature of anti-money laundering and, more widely, the counter-terrorist financing operating framework. This review aims to understand how AI is being used to check financial transactions in the big data environment of transactions at risk.

1.1. Background and Significance

The integrity of economies is highly dependent on the flow of resources in the country. As the amount of money involved in financial operations is enormous, the reliance on trust is also high, increasing the risk of fraud. Financial fraud involves both conventional and technological facets. Since the values associated with these crimes range from millions to billions, the ability to dodge legal systems is simpler. The cost of fraud reached significant amounts in recent years, indicating that regular fraud cases had been detected. Numerous financial fraud cases involved substantial sums, while many public companies prosecuted officials in various case filings over the same period. Estimates suggest that companies lose a notable percentage of their annual income to occupational extortion activities.

Since financial institutions have several years of transaction records, they contain examples of all kinds of fraud events that have occurred in the past. Therefore, most fraud investigations include an analysis of historical data to uncover frequent transaction categories of suspected entities and to develop the idea that something unusual and potentially fraudulent has occurred that has triggered an alarm or been handled by a machine learning algorithm. On the basis of historical data and case studies, the identification, detection, and deterrence of fraud through procedures and technology systems are best carried out. Event monitoring and detection, as well as responses, are part of the procedure for dealing with the occurrence of financial crime. Banks face a number of difficult challenges when it comes to achieving company-wide compliance with financial crimes-related rules and examination requirements and expectations. Regulators' expectations have grown in recent years with regard to consistency, precision, errors, and objectives for firms' compliance strategies. For the financial services industry and law enforcement, the rise of the web and mobile technologies and the emergence of new payment mechanisms have presented new monitoring and reaction difficulties. Technological advances, algorithms, and fraud When organizations create and put technology solutions in place, they range from basic to complex in nature. There is a big difference between fraud detection in terms of expenditures and techniques when considering custom-made technologies and solutions. In recent years, advances in machine learning, neural networks, anomaly detection, and AI have been used by banking institutions.

1.2. Research Objectives

The research objectives of this thesis are specified below:

- To investigate how AI technologies may contribute to enhancing the method of financial fraud detection for financial institutions. Specifically, this thesis aims to evaluate the discernibility of fraudulent activities by different AI models compared to existing financial fraud detection methods.
- To obtain an understanding of users' experiences and practical challenges for financial institutions in adopting AI-driven financial fraud detection by utilizing a deep learning algorithm. This exploratory study can make contributions to both practical industries and the academic research community.

The increasingly prevalent and intricate nature of financial fraud causes profound social harm, placing various industries and the broader society at risk. Consequently, effective fraud detection mechanisms within financial institutions are widely considered a critical issue in addressing these concerns. To this end, recent studies have begun to examine the application and effectiveness of various AI models in the field of financial fraud detection in order to improve existing detection methods. However, the practical implications in the operations of financial institutions have only been marginally studied. Given this knowledge gap, the central goal of this thesis is to adopt a rigorous analysis of the operation for financial fraud detection driven by AI to contribute to this knowledge.

2. Understanding Financial Fraud

A fraudulent practice by individuals as a means of concealing the real situation of something in order to gain financially is usually called financial fraud. Misleading information or a distortion usually accompanies the practice, and finally, there is harm to the economy. Fraudulent practices not only occur in the financial services industry but also in various other industries such as health insurance, banking, and mail. The adverse impacts of financial fraud and the diversity of existing forms require the development of fraud detection techniques that are more sophisticated, and AI with various algorithms is considered capable of fulfilling these requirements.

Fraud in the financial sector can take various forms, such as phishing. Phishing is a fraudulent act to steal personal information and is usually done through emails, fake

websites, or chat rooms. Identity theft is a form of fraud that involves impersonating others in order to use their identities without their knowledge. A common example is credit card fraud. In this type of fraud, someone uses another person's credit card data illegally. Defaulting refers to borrowers who are unable or unwilling to pay their loans, usually as a result of financial difficulties. This is also called moral hazard. Insider trading is trading in shares by people who are likely to have insider information about the company. Social engineering is a type of fraud that works in an indirect manner with the main objective of interacting directly with personnel in a particular company, starting from the janitor to the highest-ranking officials, for the purpose of stealing sensitive information. It is presented as a teaser to deceive company employees into thinking it is harmless, supporting attackers in obtaining credit card information and permitting malware installation.

Financial institutions, such as banking companies and investment firms, are the most likely to protect themselves through several transaction mechanisms that allow fraud to occur, so fraud detection must be carried out continuously. Technological development also helps fraud patterns to evolve. Therefore, more complex data patterns are also needed to catch these evolving fraud components. The current process of detecting fraud is in a phase with techniques that use artificial intelligence.

2.1. Types of Financial Fraud

Prevalent worldwide, financial fraud comes in many forms. The goal of all is to deceive individuals, organizations, and governments. In the broadest sense, fraud can be personal, corporate, or public. Financial fraud can be categorized in many ways. Many classify fraud into major types of fraud and then enumerate a few categories under each. A popular way to characterize financial fraud is accounting fraud, securities fraud, bank fraud, money laundering, and market manipulation. Recent examples of Ponzi schemes include those of individuals who perpetrated significant frauds.

Ponzi schemes are becoming even more prevalent as the world embraces different types of investments, such as real estate. Real estate investments reached an all-time high, almost double the amount invested four years prior. Investment-related Ponzi schemes generally focus on a particular type of investment, but not always. It is estimated that Ponzi schemes defraud investors of significant amounts annually. Another type of financial fraud prevalent today is mortgage fraud. Major mortgage fraud lawsuits

unfolded over a period of the prior five years. Mortgage fraud activity increased significantly from one year to the next. There was a notable increase in Suspicious Activity Reports in a particular year. The rising value of homes has become a primary driver in the increase in mortgage fraud cases. Additionally, professionals in the real estate and mortgage industries have made lucrative profits by unethically manipulating these two industries. Mortgage fraud harms those targeted by the fraudulent loan and anyone affected by devalued property and the ripple effect fraud has on our economy as a whole. If this trend continues, it threatens the financial stability of our country. In sum, there are various types of financial fraud evaluated in this study that can benefit from the use of AI applications such as expert systems.

2.2. Common Techniques Used in Financial Fraud

Common techniques used in financial fraud have involved deceptive practices that exploit the cognitive limitations of individuals. Phishing attacks rely on psychological manipulation to prompt users to reveal their credentials to attackers, while fraudulent telemarketing manipulates potential victims into disclosing sensitive personal information. Techniques that rely on fabricated information try to exploit the technological limitations of access to existing data. An example would be forging a personal identification document, such as a utility bill, to defeat new customer registration barriers. Finally, in some cases, individuals have been able to use a formal legal right during the settlement of legal disputes to gain control of assets illicitly.

Such fraudulent activities, coupled with the increasing difficulties in the detection of fraud cases, motivated extensive interdisciplinary discussion regarding financial fraud and technical ways to expose it, including the creation of tools capable of early fraud detection. This is due to the fact that the speed of technology development has significantly empowered fraudsters with the capability to create increasingly sophisticated fraudulent strategies in order to bypass rule-based security measures. For example, at present, the fraud technique of 'carding' falls into such sophisticated fraud categories, where criminals purchase products online using stolen credit card information from specific customers who belong to the same geographical location.

3. AI and Deep Learning in Financial Fraud Detection

AI offers a range of technologies capable of analyzing large amounts of data, from artificial neural networks to symbolic reasoning. This basic learning ability must be able

to be turned into action in ways responsive to the real-world constraints and opportunities that arise. AI, combined with more recent advances in deep learning, where AI solutions learn from large, dynamic, and unstructured datasets so that increasingly complex tasks can be solved, is driving a new phase in pattern recognition by enabling algorithms to learn from the raw data, a process of feature learning in an unsupervised manner. As a result, the use of this AI model is increasingly employed in a variety of business sectors such as banking, telecommunications, and logistics, and is inevitably replacing various rule-based approaches that were applied earlier.

A number of financial institutions are placing greater reliance on a new generation of software products that use AI and machine learning methods as a means of dynamically identifying suspect movements and patterns that may indicate fraudulent activities. Although the time-critical nature of fraud prevention activities often focuses on what works now, the real value in using AI as an approach to fraud lies in its ability to look for patterns in the data that are indicative of fraudulent behaviors and therefore supplement alternative fraud prevention tools and investigator skills as part of an integrated approach. Machine learning can rapidly identify these different patterns at a scale well in excess of that of a human operator and, crucially, learn from its classification errors. It is important to note that the AI algorithms which underpin machine learning and unstructured data analysis can introduce their own sets of biases and consequential issues that should be better understood and managed. It is for this reason that the ethical use of AI is an important factor in good governance and business practice.

3.1. Overview of AI and Deep Learning

Artificial Intelligence (AI) is the new electricity that is revolutionizing the world. The plethora of AI technologies has emerged and evolved in the last decade. As a subset of AI, Machine Learning (ML) has attracted much research and innovation. Deep Learning (DL) is a type of ML algorithm that is comprised of artificial neural networks and harnesses technologies such as recurrent neural networks, convolutional neural networks, or restricted Boltzmann machines to process complex datasets, thereby providing the ability to discover data representations from large untapped banks of data. Advances in DL have led to the widespread use of these networks for solving tasks such as visual data analysis, speech recognition, and natural language processing. This,

together with a continuous increase in computational power and greater availability of data, has driven interest in DL to unprecedented levels, making the field of AI grow spectacularly, resulting in a breakthrough in products, services, and applications.

DL is a powerful technology that is currently making a significant impact on different areas, including finance. The deep neural network is composed of a series of interconnected processing layers, each of which is built by many interconnected processing nodes, or 'neurons', analogous to the biological brain structure, that are functionally organized in the form of different 'input', 'hidden', and 'output' layers. Neural networks are capable of handling feature engineering so they can encapsulate a large number of operations using linear and non-linear functions, being useful for input-output mapping or when no expert knowledge is available. Deep learning has, therefore, potential for financial time series prediction, providing automatic learning of the features from the data that may improve financial decisions by uncovering hidden patterns in big data analytics.

Almost all the results of the existing studies indicate that deep learning models are superior to traditional models for time series predictions in the finance sector, as they can successfully capture complex patterns from higher disparity among the financial signal data, even small noise in futures data as well. Mainly, the combined model outperforms pure deep learning models in terms of futures prediction. One of the most important advantages of deep learning is to prevent overfitting and achieve better generalization. However, deep learning techniques have some challenges. Some researchers believe that returns index data with deep learning have some perception in raw data and relations; therefore, researchers can ignore direct predictions regarding signal noise, as they can directly predict financial forecasting, omitting data that require completeness and clarity for the future, as the presented work is still reaping the deep learning's such limitation advantages.

Given the number of networks to layers and the number of neurons to layers, the increase in the complexity of the model can lead to overfitting through over-optimization by training. The size of the overfitting will decrease gradually by incorporating a regularization term into the model. Another direction is employing optimization algorithms in model training. This parameter gets updated in iterations and comes into play under computation. This is beneficial to the model and time-

consuming for training. In addition, potential contributions could be a detailed survey of financial forecasting using deep learning, especially for the non-spective gaps and future perspectives based on datasets. However, this area has not been proposed, and researchers who have completed work reviewing the advancements of forecasting in the financial sector with various machine learning techniques have no deep learning implementations available to choose a suitable environment.

3.2. Applications in Financial Services

Alongside all these challenges, AI and specifically deep learning have significantly emerged as a major leading-edge technology to help financial services organizations provide new, innovative, and secure services to their customers. AI and deep learning have many applications in financial services, including fraud detection and prevention, trading and portfolio management, and risk assessment. They can also be applied in customer service by enhancing customers' experience and personalizing financial recommendations. Focusing on fraud detection and prevention, AI is already being used for various applications: from transaction monitoring tools that are able to analyze customers' financial activities in real-time and automatically alert respective financial institutions for suspicious transactions, to AI-enhanced chatbots answering queries and claims from bank customers. A wide range of deep learning and AI-driven solutions are creating direct financial returns as well as facilitating cost reductions. Transaction monitoring systems are a tremendous example of the successful application of deep learning and AI in financial services. Traditionally, bank staff have reviewed transactions over a certain amount by customer, making notes regarding the purpose of transactions and building profiles of their customers' transactions based on these notes absorbed over time. This process is slow, labor-intensive, and subject to human error. AI has been developed and implemented to learn from the profiles built up and the notes taken by bank staff, applying these to customer behavior in real-time, allowing the discovery of deviations from what is typical, thus reducing the time it takes to review transactions and, in turn, detecting attempted fraud more quickly. In response to the many technological advancements in fraud detection and anti-money laundering, there is now a drive to provide tech-based detection solutions to smaller financial services establishments to fit their size and resources. As regulation is a sizable issue for financial services AI developers, a consultative approach needs to be taken to develop AI technologies that share AI models and data across the industry in an ethical manner.

4. Data Sources and Preprocessing

One of the primary building blocks for effective AI in fraud detection pertains to the data sources from which models can be trained and receive predictions. Numerous studies have focused on proposing or developing models from specific, usually structured data pieces, such as transaction histories, credit card features, or supply chain data. Empirically studying a wide variety of sophisticated and integrated models requires the assembly of a comprehensive set of diverse datasets. Thus, financial transaction fraud has the advantage of already offering some very diverse inputs. First, there exists both structured data (such as relevant time, location, or counterparty information about a financial transaction) as well as transaction-related unstructured data (such as descriptions). Smartphones also provide unstructured technical data (such as IP addresses, metadata, and more). While credit card terminals come with geolocation, an IP lookup can also help to enrich IT-related data. Additionally, and of course of key importance, accounting logs might be available for further metadata on accounting fraud.

Preprocessing is a crucial stage of preparing the data for model building, model testing, or prediction, and is often used synonymously with data preparation. Data cleaning, preprocessing, and transformation steps before mining are viewed separately from the mining process. The overall approach of cleansing, normalizing, formatting, encoding, and integrating data is the same across papers, regardless of the contribution the paper presents. However, the complexity of this step can greatly vary. Normalization focuses on bringing the data to a common scale, which is particularly important for clustering or classification-based techniques when using distance-based algorithms. Given widely varying measures available in many fraud datasets, such as different value ranges, currencies, or time zones, normalization facilitates building a model that is not biased by simply a different value representation. While normalization is particularly important in financial data when training a model, it is important to note that in practice it might not always be feasible to normalize certain fields, typically only one (or a few) per dataset where one of the factors does not reside in the default observer country of a customer, nor in a local currency from which the largest portion of the same regional customer base stems. Also, normalization should not be used if it changes or removes critical information.

4.1. Types of Data Used in Financial Fraud Detection

Diverse types of data are used to detect financial fraud. Structured data mainly include transaction records, financial activities, and customer profile data; unstructured data cover emails, social media activity, etc. Electronic transaction auditing and analysis used to be the dominant method of financial fraud detection, processing structured event data. However, these existing methods have only prevailed in detecting financial fraud. Emerging data analysis techniques such as text mining could also analyze emails and other textual data to detect fraudulent behavior and are increasingly seen as significant in enhancing the detection of financial fraud. Moreover, due to the heterogeneous development of detection, acquiring different sources of event and profile data is currently considered to be of growing importance.

From transaction records of personal and corporate customers to data from the monitoring of social network activities, digital trail pieces are collected in the task of financial fraud detection from various event auditor sources. Moreover, other auditors are directly enacted via applications and services. Typically, these activities of a customer are processed and finally prepared into databases of diverse data known as structured and unstructured data. Automated processing of structured data such as customer profiles and historical transaction records can generate the complex event epilogue for the detection of "who is committing fraud?" As a result, real-time analysis of data such as transaction records of customers and online agents elicits patterns and provides for recognizing event sequences in the task of "what is fraudulent behavior?" Therefore, in a more connected world of interoperating yet independently developed information and communication systems, a pressing need exists for real-time heterogeneous analysis and integration of both structured and unstructured data.

4.2. Preprocessing Techniques

Financial fraud detection analysis deals with a variety of data fields, making it necessary to perform various preprocessing techniques before creating a model for classification. Data cleaning, the first step in data preprocessing, results in the handling of missing values, outliers, and noise. Another data preprocessing method that is usually implemented to optimize the model's performance is normalization, which performs feature scaling. This step is necessary to allow financial fraud data to be converted from a raw data format for AI models' purposes. Dimensionality reduction is a preprocessing

technique that is implemented in data preprocessing to increase the algorithm's performance. This specific preprocessing technique reduces the required data dimension while maintaining the crucial features that influence class labels. Preprocessing affects the performance of an anomaly detection algorithm. The effectiveness of anomaly detection algorithms affects the model's ability to identify and separate anomalous data from normal data. The aforementioned preprocessing techniques have been verified to enhance the performance of financial fraud detection. Although useful, preprocessing cannot escape some difficulties, such as the computational cost calculated from the available resources for a large amount of financial data complexity. Besides, financial data is more complex than other data, and it is indeed time-intensive to gain insights using knowledge from financial fraud data. In a real environment, the adaptive preprocessing technique seems to be quite reasonable. These benefits prove that correct preprocessing techniques can improve the performance of the fraud detection system and demonstrate that it is feasible to use anomaly detection algorithms for financial fraud detection.

5. Deep Learning Models for Anomaly Detection

Autoencoders are feedforward neural networks designed to perform a variety of unsupervised learning tasks. The models are trained to obtain a latent representation of the input that can reconstruct the original one. Anomalies are characterized by a deviation from the expected behavior in the reconstructed input. Autoencoders are also able to capture nonlinear relationships and structures present in data. Since the molecular level structure of fraud is often concealed, autoencoders remain a predictive model to detect fraud. The reconstruction loss is taken as the signal to detect the anomalies in the data. The smaller the reconstruction error, the lower the likelihood that the input is an anomaly. One class of autoencoders that have achieved higher performance in outlier detection is Outlier Exposure. The study demonstrated the benefits of knowing the training data's noise distribution for improved performance in outlier detection. In the evaluation, three different autoencoders are assessed, including undercomplete, denoising, and variational. An example of a two-layer autoencoder is defined as follows: where x_k is the k th input data instance, z is the latent variable, ω_{enc} is the parameter of the encoder, $f_{\psi_{enc}}$ is the function for encoding input x , h is the size of the hidden layer, w_n and b_n are weights and bias terms. The output y is the reconstruction of the input x . Here, two different reconstruction losses are evaluated:

evaluating logistic sigmoid and margin ranking losses with early random and fine-tuning of the presentation. To compare the autoencoders, five deep networks with different architectures and hidden sizes are chosen to evaluate the performances. The preliminary results showed that DAE and VAE based neural networks performed better than one hidden layer based on the architecture design. Moreover, the DAE outperformed the VAE for a one-layer network.

Semi-supervised learning is also possible using autoencoders as long as paired data is used. This methodology proves to be a powerful generative model and has applications such as image generation and reconstruction. A more recent application is on anomaly detection using perturbations and generation of unimodal, multimodal, or nonlinear Gaussian process based distributions using the latent data. Variational autoencoder is another typical unsupervised neural network for generative probabilistic modeling. The learning algorithm requires the maximum likelihood estimation, where a split variable is introduced: $y \sim p\theta(y|x)$ for the Bernoulli distributed output, and $p(x|z)$ and $p(x)$ defined as probabilistic and prior distributions, where z is drawn from $p\theta(\epsilon|x)$ defined as $\epsilon \sim N(0,1)$. The scale is increased by sigmoid $f'(\epsilon) = 1 / (1 + \exp(-\epsilon T))$. Gaussian stochastic losses for computing probability from the probability density function are determined by both the reconstruction term and divergence term, leading to KL divergence from the prior distribution parameter z , $z \sim \prod \epsilon_i \phi$ to the posterior distribution $q\phi(z|x)$.

5.1. Autoencoders

Autoencoders typically have an encoder-decoder architecture that consists of the following two components: 1) an encoder that maps the input data into a presumably low-dimensional space, creating a latent representation; and 2) a decoder that reconstructs the input data from the latent representation built by the encoder. To detect anomalies in financial datasets, autoencoders can build latent representations of the input data. As the autoencoder is trained, it learns the distribution of the normal input data. Thus, if a copy of the data has more errors than usual during the reconstruction process, it will be recognized as an anomaly. The attention mechanism may also be utilized to highlight only the relevant parts of the input data for the reconstruction process. One of the major advantages of the autoencoder method is that it can be trained for a specific problem and can fit the distribution of the input data. It can also be

improved in many ways, such as through the utilization of noise contrastive estimation. Compared to traditional methods, the performance of the autoencoder is much better, yielding superior precision. Challenges might arise while adopting autoencoders, such as the requirement of a large dataset for training, which may be computationally intensive. Real-world applications of autoencoders in a financial environment could also be mentioned. For example, they can be put to use in investment security to detect anomalies in data such as invoices, payments, and so forth by analyzing the data patterns or trends.

5.2. Variational Autoencoders

Variational autoencoders (VAEs) are a more advanced variant of traditional autoencoders. They are based on an underlying probabilistic framework and are particularly known for their generative capabilities. In the financial context, VAEs are primarily used for unsupervised anomaly detection. The probabilistic framework combined with the autoencoding approach to unsupervised learning makes VAEs particularly effective at capturing the distribution of the input data in the latent space, making them suitable for the detection of financial service misuse. In the latent space, variational autoencoders incorporate an encoder and a decoder network, similar to standard autoencoders. However, variational autoencoders differ from their standard counterparts due to their treatment of the underlying latent variables, which are modeled as the mean and variance of a multivariate Gaussian distribution. Through encoding, variational autoencoders approximate the posterior distribution of the latent variables in the Gaussian pattern. This allows variational autoencoders to better model complex data distributions by exploiting the generative capabilities of deep learning and the probabilistic nature of the underlying data as captured by the latent variables. Consequently, variational autoencoders can model uncertainty and offer a wide range of generalization capabilities that go beyond the standard autoencoding approach. The generative component of VAEs can be utilized to generate synthetic data points, which—when interpreted intuitively—will lie in the same manifold as the underlying genuine input space. As the latent space of VAEs follows a multivariate Gaussian distribution, several abnormal data points might lie parallel to the underlying manifold and will be able to capture more potential financial crimes. The probabilistic approach to detection enables VAEs to better model the uncertainty present in the data.

Applications of VAEs can be found in the detection of outliers within the insurance sector, where they are being used to predict possible fraud scenarios. This advanced variant of autoencoders can also be used within the finance sector. VAEs are especially prevalent in the detection of rogue traders, as insider trading can at times be considered a service misuse by financial institutions. VAEs also have use in the anomaly detection of credit/debit card transaction systems, where they can be used in the verification of transactions. Additionally, they are used in public sector surveillance activities detecting fraudulent activities. Given their probabilistic nature, VAEs can help financial institutions make well-informed decisions on behalf of recovering fraudulent cases. These decisions are based on a more sophisticated statistical model. However, it is very computationally expensive to train a VAE and to integrate it into the institution's existing systems.

6. Evaluating Model Performance

Early detection of fraudulent transactions not only decreases the impact on a victim but may also capture fraudsters. The three major evaluation metrics for measuring the predictive model are accuracy, precision, and recall. Accuracy is the percentage of correctly classified instances out of the total number of instances. Precision is the percentage of correctly classified actual fraud instances out of the total predicted fraud instances for the model. Recall is the percentage of correctly classified actual fraud instances out of the actual fraud instances in the dataset.

In a real-world scenario, there is a trade-off between a false positive and a false negative. A false negative will ignore a fraudulent transaction as not fraudulent. False positives occur when the model classifies a transaction as fraudulent whereas it is not. An F1 score is calculated using precision and recall to balance false negatives with false positives. In our case regarding fraud detection, we are interested in fraud instance validation and not class one accuracy. In some cases, it is acceptable to avoid a true fraud prediction. It is acceptable to allow some genuine transactions to be incorrectly classified as fraudulent.

The more penalized class (Class 1 - fraud accounting for 0.1% of the total number of cases) should ideally be the class with the highest precision or a higher percentage. As the models in use can only ever give a possibility of fraud, no model would receive 100% recall or precision. There is always a trade-off between presenting the correct fraud

instances and avoiding showing non-fraudulent a fraudulent status. The Random Forest and our Comb_MLP models at 0.5 probability cut-off presented the highest F1 scores. It is noteworthy that the Comb_MLP model was the highest in precision. This model is the most useful as it presents a good avoidance of false positives, which is necessary in a fraud detection scenario, so as not to suspect genuine activity. All models could still be improved upon by further parameter tuning and also the use of higher data sampling. Removing different clusters of data may help to improve the score. All the models were tested on the same pre-processed data where outliers were removed and scaling was carried out.

The model hyper-parameter tuning may still be performed to further evaluate each model's performance in practice. When new data is used for further reviews of the models in deployment, this may update the model. As the standard has shifted recently with the 3D secure validation, which puts an extra layer, the model using further data from review with 3D secure is a possibility for review. Models in place should also be tracked to improve the data monitoring over time. This will help for future model evaluations. The need to have robust detection of fraud is of concern as the volume of payments isolated from the web has significantly increased recently. It is timely to see the benefits and caveats of different models in isolation and some in combination to predict new fraud instances.

6.1. Key Metrics for Model Evaluation

The key metrics of AI model evaluation in order to estimate the quality of constructed models are presented. It is not enough to have high accuracy as a model benchmark. Moreover, it is essential to minimize false positives. This straightforward value reflects the correct rate of detected fraud cases. The cornerstone of the progress has been the capacity of machine learning techniques, especially deep learning, to process, organize, and gain knowledge from large collected data. Besides this, these AI-based solutions have presented extraordinary outcomes in several aspects.

The confusion matrix allows us to see how many particular cases are falsely or correctly classified. This may produce possible headlines that we try to reach with regard to detection techniques. Several metrics can be used to determine the best model for the desired detection in a fraud detection scenario. Their meaning in the fraud context is briefly described hereafter.

- TN (True Negative): Normal instances that have been correctly classified.
- FN (False Negative): Fraud instances that have been misclassified as being a normal class. Several scenarios might be interesting in terms of decision rules.

- The receiver operating characteristic (ROC) curve is created by plotting the true positive rate (TPR) on the y-axis against the false positive rate (FPR) on the x-axis at various threshold settings. The area under the curve (AUC) indicates the capacity of classifiers to use each threshold. To sum up, in most data problems including fraud, the interplay between precision and recall means that our figures are not meaningful in isolation and trade-offs are to be investigated.

6.2. Comparative Analysis

In this section, a comparative analysis is provided to show the performance results of different AI models applied in the financial fraud detection field. The developed ANN, KNN, DT, RF, XGBM, and LGBM were thoroughly evaluated and validated. It is believed that conducting a cross-model analysis of the most popular AI methods may be important for decision-makers, as it might provide insights into the type of fraud-specific AI model based on the suitable performance metrics such as precision or recall that should be used in a given fraud detection system. Different fraud datasets are utilized by researchers to demonstrate the performance of their proposed fraud detection systems. Some argue that classical methods in credit card fraud detection are still widely used today, but in the face of big data and the deep learning revolution, they are being replaced by new deep learning methods for better results. However, the choice of the developed model must be carefully thought out as per its computational cost, especially in real-time fraud detection scenarios. This is because using a many-layer network in a neural network model is associated with higher computational complexity and massive data preprocessing activities.

Consequently, processors might be unable to meet the required computational speed. Most decision-makers in the field may want to use a shallow network in the neural network models. The study outcomes demonstrated that the RF model showed the best performance in terms of precision, recall, and F1 score. It is supposed to therefore be the “go-to” AI model to be used when the management intends to catch as many fraudulent transactions as possible. Such a finding aligns with several examples demonstrating that the RF tree, in particular among the mentioned AI classifiers, outperforms other models

in fraud detection. Furthermore, the study intended to emphasize that there were additional factors that influenced the predictive performance of the developed AI model. The possible causes differed between the sampled true fraud and the sampled non-fraud transactions, although they were from the same parent population. In this context, the sampled non-fraud transactions may have been chosen endogenously. It thus partially explains why the analysis was unable to surge the predictive performance measured in f-max. Also, the selection of features and the methodology to handle imbalanced data, among others, appeared to impact the model's explanatory power. This provides additional guidance to decision-makers when deploying fraud detection processes and psychological understanding of AI-driven fraud detection problems.

7. Real-Time Implementation and Challenges

In financial scenarios, the ability to implement anomaly detection in real time is crucial, as an immediate response is necessary in order to minimize financial losses. In addition to the abundance of potential features that can be used to describe the underlying transaction context, there are several complications associated with implementing a real-time anomaly detection system. One of these complications is the low-latency requirement needed to keep up with system traffic that needs to be analyzed. Furthermore, data storage, retention, and destruction policies are very complex for systems that store all transactions needing to be analyzed in real time. Moreover, in order to be integrated into existing monitoring systems, any anomaly detection method needs to support a read-only application programming interface. It should be noted that alerting behavior may also be part of this API.

In particular, the implementation of a machine learning model in banking financial products and systems is very challenging due to the existing range of heterogeneous solutions, including large amounts of data generated from various data sources, the need for a fast time to market, and the complexity of adjusting the system to the model output. In line with these observations, several works and reviews have discussed the practical challenges of integrating AI models, including server architecture, computer clusters, and database management systems in banking infrastructure. AI and machine learning services often rely on data retrieval and processing as well as database services. Such services obtain data via a read-heavy transferring system, which can be collocated with dedicated cloud-based storage and computation servers. The core servers,

meanwhile, are used for hard real-time transactional analytics, which cannot be deployed on cloud services quickly enough to offset the total system performance. Most data used in training financial models is less timely and can be computed offline with in-memory applications proprietary to the on-premise data center. Machine learning models are used to predict future fraudulent activity with epoch counters reflecting the transaction sequence numbers in the datasets. In addition, the systems are equipped with technologies that can leverage colocated hardware to increase I/O operations per second. These I/O operations are necessary for read-heavy applications. Alternatively, small-scale cold data for these read-heavy data stores are returned to slow spinning disks to maintain consistency across servers. Another practical requirement in implementing real-time anomaly detection in a financial setting is the need for a system capable of real-time data acquisition and subsequent analysis. While the system should be adapted to real-time dynamics—the speed at which fraudulent tactics evolve and change the tool in use. In particular, systems should have the capability to support several types of analysis and minimal storage duration. In addition, the storage backend may have a local cache in order to avoid disk I/O contention. A front-end cache is required to listen for new data that the backend cache has not yet been able to process. Finally, edge processing may be used to reduce latency and improve scalability. For instance, the cloud may be used to check for data consistency and quality or for analyzing anomalies without causing a performance reduction. In addition, edge clusters may be used to store query buffers and aggregate data.

7.1. Challenges in Real-Time Implementation

The implementation of an AI-driven fraud detection system in a real-time environment faces several challenges. System latency, namely the time taken to process the incoming transaction events by the detection system, must be minimal. An increasing consumer base leads to higher data throughput. The time allocated for the system to respond to the incoming data is called response time. In a real-time fraud detection scenario, with millions of transactions occurring every second, the detection system should be able to process data and respond within milliseconds.

Another significant challenge in the real-time implementation is to identify the most challenging fraud instances using minimal resources. The vast amount of transactional data increases the complexities of detection. Such capacity is in short supply for

traditional systems. Real-time fraud detection systems are expected to be integrated with existing systems that consist of a complex infrastructure. New technologies must match the application programming interfaces of the already present systems to ensure that they can exchange data freely. Processes of manual updating are available for existing rule-based legacy systems. Real-time fraud detection is anticipated to be a stand-alone system. The backbone infrastructure should be robust enough to warrant continuous and effective operation. Consequently, the system requires monitoring and upgrading in the event of a degradation of performance and changing fraud tactics. Both of these factors require enormous computational effort. A bit of human involvement is suitable for ongoing operations. In sum, the requirements of a scalable solution call for a great deal of computational and human resources given the restrictions mentioned above.

7.2. Solutions and Best Practices

Several solutions are available to mitigate the aforementioned challenges and successfully deliver real-time fraud detection. Cloud technologies are now mainstream and can offer virtually limitless computational power. They are deeply integrated into most processing devices, including mobile phones, and are increasing in popularity to power IoT devices. They can greatly help in implementing a scalable fraud detection service that can match the size of its purse. When designing such systems, it makes sense to plan for scalability from the beginning and to continuously optimize the amount of processing required to deliver results. A layered approach to fraud detection also presents itself as a powerful best practice. A firewall, antivirus, antimalware, and spam filter may all protect and alert on various aspects of an attack; a similar set of defense mechanisms, directly or indirectly, can be used to detect financial crime.

Data management best practices also apply in this space. When designing real-time systems, it is important to have a detailed understanding of the data that needs to be processed, how often it is produced, and extensive design considerations to ensure data integrity and that processing does not become a data flow bottleneck. It is important to create 24/7 systems that can be easily maintained and to make sure that the technology stack required for deployment is within the grasp of the organization. Finally, systems in this domain need to be continuously trained. In most AI-driven constructs, the models degrade over time and need to be updated frequently. Integration with model and feature update pipelines and technologies, where the overall system can be updated

without shutting it down, are some of the best practices that can ease the pain typically encountered by organizations to make this happen.

Lastly, it is important to document successful case studies in this space with broad reach. A well-documented case study provides many benefits, including solution validation, real-world application, and much more. Prominent case studies in this space include those related to point-of-sale and automated teller machine fraud detection.

8. Conclusion and Future Directions

This essay has reviewed the current state of the use of AI in financial fraud detection. Overall, the advancement in AI tools has helped significantly in the development of more robust financial fraud management systems. The deep learning technique, with its better performance, has gained popularity among other AI tools. Ongoing improvements in AI technology and adoption within the financial sector have yielded significant results in detecting fraud. Many false positives, those that are legitimate but were classified as fraud, can be removed. Often, false negatives, fraudulent activities incorrectly classified as legitimate, can also be reduced to an acceptable level. Despite these findings, a number of limitations and ethical challenges need to be mentioned. In the future, researchers and practitioners need to ensure that AI technology remains a key focus, with innovative solutions and improvements for AI and deep learning measures. Other areas of application of AI and big data, such as text mining and news sentiment analysis, also have interesting potential. There is potential to create more proactive and resilient fraud detection by ensuring that AI is used to counteract financial crimes on a broader scale that operates within a financial system. However, while making progress, those areas need to be further explored. One promising approach could include the joint work of practitioners and scholars to provide access to data and expertise. In addition, compliance with ethical principles and codes would ensure the proper use of sensitive financial data.

8.1. Summary of Findings

The application of AI in general, and deep learning in particular, to the domain of financial fraud has elevated state-of-the-art performance and detection capabilities over the past several years. In particular, the study of autoencoders and variational autoencoders, or otherwise unsupervised models, is due to their capability of modeling, in some high-mass dense zone, the distributed data. Reasonable AUC scores of 98.222%

for machine learning and 93.870% for deep learning exhibit the high performance that AI models can achieve when optimally trained on highly pre-processed datasets of features that are relevant to the organizations and to the benefits of fraud detection upon financial conduct of any kind. With the ever-increasing number of more advanced state-of-the-art models that are reaching up to 99-100% in terms of precision, the demand for ongoing research in the field is emphasized. Deep learning technologies have increased the levels of prediction accuracy with respect to fraud detection. Achieving a maximum AUC score of 98.222% for machine learning and 93.870% for deep learning, the main findings include demonstrating better performance than traditional and latest state-of-the-art methods, as well as variation adjustments, utilizing highly preprocessed customer transactional banking data and over-sampling with informative synthetic sampling to detect internal fraud occurring within the financial sector and across its many areas.

8.2. Future Research Directions

Despite the research achievements and practical applications of AI in financial fraud detection, many challenges raised by the emergence of new technologies and methodologies remain unaddressed. The interpretations of black-box AI models employed in financial fraud detection systems are crucial for stakeholders to comprehend the AI-based anti-fraud results, make proper decisions based on them, and build trust in such systems. Explainable AI is a newly emerging research topic to interpret and visualize complex, high-dimensional, and non-linear models into understandable and transparent representations, promulgating interpretations and explanations. Integrating financial AI fraud detection with other AI technologies is an interesting future research direction to prevent financial fraud, such as controlling database manipulation and solving principal-agent problems in the financial industry. Future work can consider conducting interdisciplinary collaborations with ethicists and legal professionals to mitigate issues surrounding model interpretability.

Since the business environment and fraudulent practices are continually in a state of evolutionary dynamics, researchers can also explore the real-world practices of integrating AI in financial fraud in emerging economic entities. It is necessary to examine how well fraud models that deploy AI techniques perform within a five or ten-year horizon. The efficiency of these AI-based financial fraud models may be

continuously eroded as criminals and fraudsters continuously alter their behavior and strategies to avoid the increasing fraud prevention tools developed in consort with cutting-edge AI techniques. Moreover, recent results suggested dissatisfaction with the traditional research approach to fraud. To revolutionize the research approach to fraud is to start developing contemporary models and frameworks that may offer fresh insights into the processes underlying this form of betrayal and deception.