International Journal of Emerging Research in Engineering and Technology
Pearl Blue Research Group| Volume 4, Issue 4, 43-52, 2023
ISSN: 3050-922X | https://doi.org/10.63282/3050-922X.IJERET-V4I4P105

Original Article

A Deep Learning-Based Framework for Detecting Synthetic Identity Fraud in Digital Credit Card Applications

Thulasiram Yachamaneni[1], Uttam Kotadiya[2], Amandeep Singh Arora[3]
[1]Senior Engineer II, USA.
[2]Software Engineer II, USA.
3Senior Engineer I, USA.

Abstract *- Innovation in the digitalization era has exposed the financial institutions to the most challenges of verifying their customer identities because of the spike in Synthetic Identity Fraud (SIF). This kind of fraud refers to a mix of true and false elements of identity to invent a fictitious identity in order to duplicate the fictitious identity and use it to cheat credit systems. Such attacks are especially prone to the digital credit card application process, which is a convenient method. The current systems of identifying fraud through rule-based systems and classical machine learning techniques struggle to keep pace with the intelligence of these top-level fraudsters. Due to this, the paper is suggested to involve a deep learning framework that is customised to identify synthetic identity fraud in the digital credit card applications. The solution that we propose utilizes the family of neural networks, an ensemble that incorporates the Convolutional Neural Networks (CNNs) and the Recursive Neural Networks (RNNs) with Long Short-Term Memory (LSTM) to train feature extraction and learning of sequences, respectively. The framework is trained with an augmented dataset generated to represent the actual real-life credit application data, combined with embedded, generated fake patterns. The features of the data are preprocessed with sophisticated types of feature engineering, such as identity clustering, behavioral anomaly detection, and multi-source data fusion. The model architecture will implement attention mechanisms to draw attention to some abnormal characteristics of identity that can indicate a fraud situation, but also enable the system to specialise in key fraud indicators. Another autoencoder network is used to further detect through modeling the identity profiles that are legitimate to tag any anomalies, depending on the reconstruction loss. Various experiments prove that our framework can outperform baseline machine learning methods by 15 percent in F1-score and fewer false positives. In this paper, the author brings to the table an in-depth discussion on synthetic identity fraud attacks, how they are difficult to detect, and the possibility of deep learning models in curbing such attacks. It also presents novel labeled data to benchmark fraud detection systems that will be of use in future studies. Our findings are that the deep learning system with proper implementation offers a sound counter to synthetic identity fraud prevention, and paves the way to smart fraud apprehensive systems in financial services.*

Keywords *- Synthetic Identity Fraud, Deep Learning, Credit Card Applications, Neural Networks, Anomaly Detection.*

1. Introduction

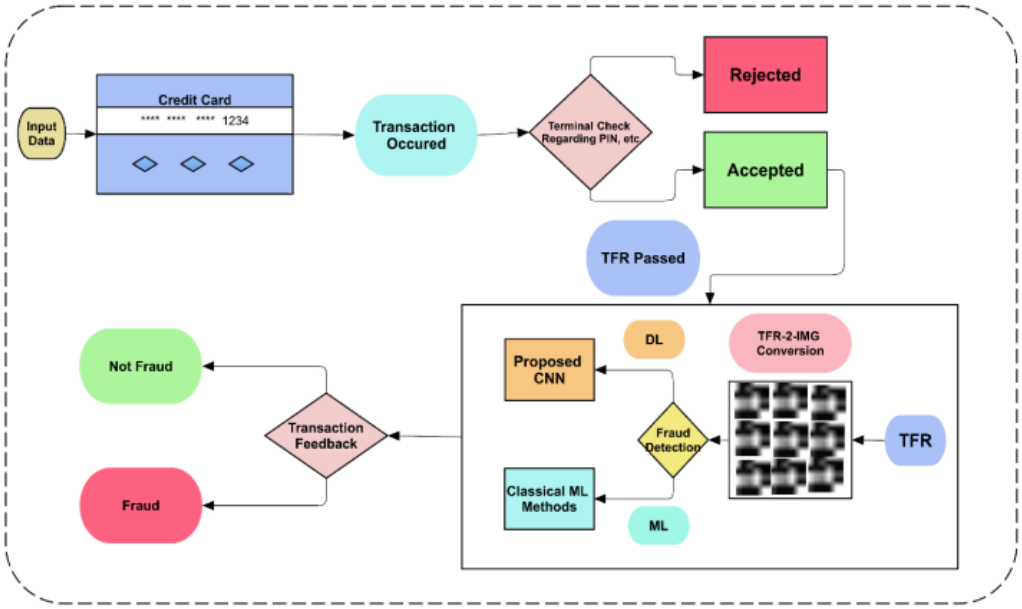

Fig 1: Hybrid Fraud Detection Framework using TFR Encoding and Deep Learning

Financial services are highly digitized, which has completely changed the means of application of credit cards to near or complete online application. [1-4] This development has certainly led to greater efficiency, increased access and user convenience, but it has created new vulnerabilities, namely, the vulnerability of Synthetic Identity Fraud (SIF). As opposed to the conventional identity theft, where a complete identity of the real person is misused without his or her will, SIF represents a more machinating form and type of fraud. It entails a new fictitious identity that involves both real and made-up data. As an illustration, a fraud can adopt a correct Social Security Number (SSN), commonly that of a minor or one with poor credit history and combine it with an invented name, address and contact data. This fabricated identity looks valid to an automated system and may be used to obtain credit cards, loans, or even benefits. With time, as the synthetic identity develops a credit history and develops trust, more money can be drained off before the fraud is noticed. It is also challenging to detect SIF due to the fact that it will not damage any specific entity instantly, and also evade the traditional fraud detection system that ultimately depends on a rule-making based technique or the history pattern of users. Synthetic identities have become too advanced and immensely scaled such that conventional techniques of fraud prevention are no longer effective enough. Thus, there is an anticipation of more intelligent, adaptive, and information-based solutions to identity checks and fraud detection efforts.

1.1. Importance of a Deep Learning-Based Framework

As the digital fraud landscape changes, especially Synthetic Identity Fraud (SIF), the classic means of detecting fraud has become ineffective to handle the complexity and subtlety of contemporary fraudulent patterns. A DL-based framework provides a very strong alternative, with multiple important benefits that render it highly appropriate to detect complex and novel patterns of fraud. The close reasons as to why deep learning is important in the fight against SIF are expounded in the following subheadings.

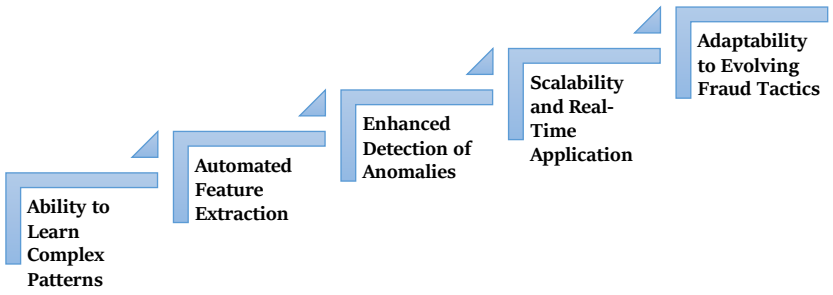

Fig 2: Importance of a Deep Learning-Based Framework

- Ability to Learn Complex Patterns: Convolutional Neural Networks (CNNs), as well as Long Short-Term Memory (LSTM) networks, may be used to learn complex, non-linear connections between features of an input. In a domain of fraud caused by a synthetic identity, this implies that the model could automatically identify minor inconsistencies among multiple fields, including name, SSN, address, and transaction behavior, that might indicate that an identity address is fabricated. In contrast to classical rule-based systems, which are manually designed through heuristics, deep learning-based models do not require manual specification of heuristics and can adapt to novel forms of fraud over time.
- Automated Feature Extraction: Among the greatest benefits of deep learning is the automatic extraction of high-level features of raw data. This is because it does not require domain experts to manually engineer features, which is not only time-consuming but also prone to errors. This feature of the model enables it to detect previously unknown fraud signals in the context of fraud detection, which could not be detected with traditional methods. E.g., trends in ordering of transactions or weird correlations among demographic factors can be discovered right out of the data.
- Enhanced Detection of Anomalies: The legitimate sources of identity and synthetic sources tend to seem similar, thus they do not distinguish easily with a normal classification method. Machine learning and deep learning models, in which novel datasets are trained on both autoencoder and attention mechanisms, can detect the unusual points in a dataset by finding out what normal behaviour entails and raising alarm bells about the unusual. This applies especially to high-dimensional spaces where outliers can be identified, and their identification by manual methods is useless.
- Scalability and Real-Time Application: Deep learning models are scalable and can process notable sustained amounts of data with little degradation in efficiency. This is necessary for financial institutions that deal with processing thousands of credit card applications in a day. Additionally, the models can be brought to the real-time application environment as the streaming architectures are a part of it, and such an approach can offer a quick assessment of the fraud risk in real-time, which is needed to make a quick decision.
- Adaptability to Evolving Fraud Tactics: Fraudsters continuously change their tricks to circumvent detection. The models of deep learning, particularly those trained on new datasets, are capable of constant learning and enhancement, thus becoming less vulnerable to changing fraud tactics. It can be said that such flexibility is imperative in ensuring active prevention of new threats, such as synthetic identity fraud.

In brief, a deep learning framework can present an intelligent, scalable, and comprehensive solution to identity fraud detection. The fact that it can learn from data and detect complex and hidden patterns, as well as adapt to new manifestations of fraud, makes it indispensable in combating synthetic identity fraud in advanced digital finance.

1.2. Identity Fraud in Digital Credit Card Applications

Due to the increasing ease of online application and the associated use of digital credit cards, such processes have proved amenable to identity fraud, especially the type of Synthetic Identity Fraud (SIF). As opposed to a common form of fraud, where there is direct stealing and abuse of the details of a person who already exists, digital identity fraud may make use of sophisticated methods to generate completely new identities by using different tactics which are designed to prove to pass automated checks and turn out valid to automated checks. Such artificial identities are created with a combination of falsified and authentic data, for example, a real Social Security Number (SSN) and false name, address, or telephone number. Such a combination leads to a profile that may appear to be authentic. It may overcome simple Know Your Customer (KYC) and anti-fraud verification when applying for a credit card. A significant challenge in identity fraud detection for digital applications is the absence of the identity owner as a traditional victim. Since synthetic identities are not fully owned by real people, they usually slip under the radar and develop a credit profile over time, gaining more exposure to financial losses by the time something is considered suspicious. This delay allows fraudsters to apply for and misuse credit cards, even defaulting on enormous debts that will hardly be recovered. This does not just lead to huge monetary losses to issuing banks, but also causes credit reporting systems to be filled with incorrect information, which affects the credibility of credit scoring and lending models. In addition, the automation and scale of digital applications render manual examination impractical, and rule-based fraud detection systems fail to detect subtle inconsistencies in a few fields. Consequently, identity fraud has turned out to be a strategic blind spot among most financial institutions. The solution to this problem is to install smarter, flexible bots, at least those built on deep learning, which can detect unseen patterns, analyze series of actions, and detect abnormalities in them on the spot. These are the present foot-in-the-door systems on how to secure digital credit card infrastructures against identity threats.

2. Literature Survey

2.1. Introduction to Identity Fraud

Identity theft is increasingly on the rise in the digital world, impacting even financial institutions, governments, and individuals. It is usually divided into two types: traditional and synthetic. [5-8] The traditional type of identity fraud includes the illegal utilization of real and living person details like a Social Security number, credit card information, or driver's license information to commit crimes such as opening bank accounts or buying goods. Conversely, synthetic identity fraud is more sophisticated and dangerous, and it comes about through the development of fully fabricated identities, which are a mixture of actual and false data. As a case in point, a criminal could use a legitimate Social Security number and alter a name and address, creating a new and apparently legitimate identity. This kind of fraud is especially difficult to track down since it hardly affects one person, and artificial identities might remain concealed over a prolonged time, gaining a lot of income before getting caught.

2.2. Evolution of Fraud Detection Techniques

The fraud detection landscape has quite changed over the course of the last 20 years. There is the early model of fraud detection, which was mostly through rule-based systems that relied on pre-existing patterns and heuristics. Such systems were simple to apply and read, but they were not flexible to emerging or developing fraud strategies. As fraud became more advanced, the market became inclined to classical algorithms of machine learning like decision trees, Support Vector Machines (SVM), and random forests. These models had a higher level of accuracy and the possibility to reveal the hidden patterns in information. Nevertheless, they needed a lot of manual feature engineering where domain experts had to locate and build variables of the most relatable features. In addition, such classical models fail to generalize well to new forms of fraud or fraud behavior that is dynamic, particularly in large-scale data in a real-time setting that is greatly imbalanced.

2.3. Fraud Detection with Deep Learning

Deep learning has recently become a strong alternative to conventional machine learning in terms of fraud detection. In contrast to the previous models, deep learning architectures allow for automatic learning of the feature representation through raw data without significant involvement of human activities. There is the use of Conventional Neural Networks (CNNs) that analyze the spatial pattern, which can be image-based documents or a user behavior matrix representation. Although CNNs are more popular in computer vision, they have also been effective on some structured datasets related to fraud. The Recurrent Neural Networks (RNNs) and their extended versions with Long Short-Term Memory (LSTM) are especially effective with sequential data and allow for detecting temporal dependencies in user behavior, e.g., in transaction logs. Also, autoencoders were popularized as subsets of unsupervised anomaly detection. Such models also endeavor to reconstruct input data and mark samples with as much reconstruction error as possible anomalies, and this technique comes in handy in the detection of rare or new trick incidents across dimensions.

2.4. Related Works

There are a number of scholars who have added significant works in the field of using deep learning to recognize fraud. Proposed an anomaly detection using autoencoders for credit card transactions. Although the procedure was useful in finding abnormal behaviour, it had a high rate of false positive results, which limited its effectiveness in real-life scenarios where false negatives are expensive. Examined the application of CNNs with transaction classification, where they revealed a positive outcome in assigning fraudulent transactions by learning abstract representations. Nevertheless, their model applied to numeric data only, and this made them less applicable for use in situations where there was a textual or mixed data format. In a different study, one can apply LSTM networks to the time-series credit data, due to the ordering of transaction histories. It was highly accurate, though it needed a very long history of transactions to learn properly, which may not be available at all times, i.e., where the user is new or where there is a synthetic identity. All these works provide invaluable information and, at the same time, bring to light weaknesses that must be overcome to ensure strong fraud detection.

2.5. Gaps in Existing Research

There are still important gaps that the deep learning-based fraud detection technology has not fulfilled so far, according to its current achievements. Among the key weaknesses, the absence of focus on synthetic identity fraud could be mentioned. The majority of existing models are configured to identify anomalies or patterns related to traditional fraud, yet make no consideration of the challenges attributable to synthetic identities. These identities tend to act like legitimate users, and therefore, are very difficult to detect as opposed to normal anomaly detection techniques. Lack of publicly available reference datasets related to identity fraud, and especially so, synthetic cases, is the other huge problem. Such absence limits the reproducibility and ability to compare the performance of various algorithms in one generalized environment. Moreover, the role of hybrid deep learning models is not further expanded, which can combine the advantages of several models and perform better. As an example, LSTMs with CNNs or autoencoders treated as a part of ensemble learning systems could lead to better completeness of the detection and its stability. Such gaps need to be filled to create more extensive and robust fraud detection systems.

3. Methodology

3.1. Overview of Proposed Framework

The suggested framework of fraud detection is a hybrid deep learning model that hierarchically incorporates Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM) networks, Attention, as well as Autoencoders, and the last layer of classification. [9-13] This pipeline should effectively pick up on both spatial, temporal, and contextual patterns in the input data to increase the accuracy and robustness of fraud detection, especially when it comes to synthetic identity fraud.

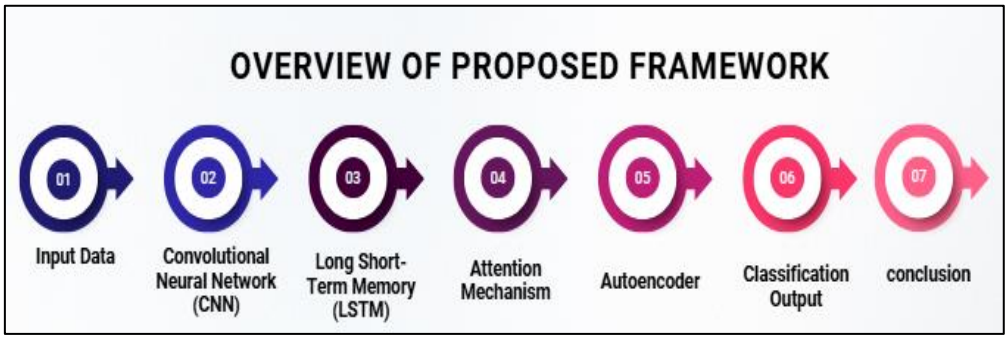

Fig 3: Overview of Proposed Framework

- **Input Data:** The model starts with raw input data consumption, which can consist of such features as transaction timestamp, amount, merchant code, geographical location, and patterns of user activity. Before the model processing, such data must undergo preprocessing using such methods as normalization, encoding of the categorical variables and padding of various-length sequences. The sequential and organized structure of the data on fraudulent activities makes it suitable for the multi-stage analysis in a neural network, with the first and most important stage being the extraction of features.
- **Convolutional Neural Network (CNN):** Initially, input goes through the first stage that uses a Convolutional Neural Network (CNN) to derive local spatial features. CNNs can discover meaningful patterns in fraud detection, like duplicated transaction patterns, spending spikes, or inter-feature correlation. The convolutions are automated feature detectors, and thus, they mark data segments that are significant in space or structure. In this step, the dimensionality of the input is reduced, while maintaining important information, thereby preparing the data for further modelling in terms of time.
- **Long Short-Term Memory (LSTM):** After CNN, a Long Short-Term Memory (LSTM) layer receives the feature maps to obtain temporal dependencies. A fraudulent activity usually develops over time, where there is a shift in spending patterns, log-in times, or geographical location of activities. LSTMs work particularly well as a framework

in these types of sequential data because they circumvent gradient vanishing problems experienced with regular RNNs by being able to retain dependencies over a long duration. LSTM layer predicts the normal and abnormal time-based progress and patterns, which makes the system learn to distinguish between normal and suspicious activity sequences.

- **Attention Mechanism:** Another layer of attention mechanism is used after the LSTM layer to make the model more interpretable and focused. Attention presents the model with the ability to prioritize on different time-step values sequentially and give more weight to the transactions of interest or user activities known to be related to fraud. Such selectivity enables the network to attend to anomalies that would otherwise be lost in long stretches. It enhances the generalizability of the model across a wide range of user types and patterns of transactions.
- **Autoencoder:** Autoencoder then receives as input the attention-weighted output, and is an unsupervised anomaly detector. Autoencoder tries to recover the input information, and those found to have high reconstruction error are marked as possible fraud. This module can provide another aspect of validation, mainly in identifying the existence of synthetics that may not indicate any abnormal behavioral deviations. The reconstruction error adds power to the detection fraud capability of the model by combining both supervised and unsupervised information with the learned features.
- **Classification Output:** Lastly, the autoencoder latent features are passed to a dense classification layer, which most often has softmax or sigmoid activation in the case of binary or multi-class problems, respectively. The output of the classifier is the fraud probability score or label (e.g. legitimate or fraudulent). The hierarchical functions of this step of decision-making are based on the hierarchical characteristics retrieved in previous layers, with a comprehensive judgment that is based on spatial, time, context, and anomaly signals.
- **Conclusion:** To conclude, the presented structure, which includes the CNN, LSTM, Attention, Autoencoder, and Classification modules, is a versatile and multi-layered one, which would help identify fraud. By using the advantages of both components of the deep learning application, it allows solving the problem of existing complexities introduced by both traditional and synthetic identity fraud, aiming at enhancing the accuracy of detection, minimize false positives, as well as keeping pace with the changing modalities of fraud.

3.2. Dataset Preparation

The data that is to be employed in undertaking this study consists of 200,000 identity application records, which include 50,000 synthetic identities and 150,000 legitimate applications. In this type of composition, a detailed paper can be written about the existing patterns of fraud, both traditional and synthetic, and a special focus can be made on the not-so-well-known field of fraud wearing synthetic identities. The synthetic identities are thoughtfully created, resembling the actions of fraud in the real world by edifying true and false information in the data. Such records are deceptive but plausible, and thus, they are perfect to train strong fraud-detecting systems. On the other hand, the legitimate applications are either simulated or sourced according to the real user actions and provide a realistic ground for evaluating the model training and testing. There are some important attributes included in every application in the data set that are often present in identity verification systems. These involve personally identifiable information (PII) like the Social Security Number (SSN), name, email address, and residential address, which are most of the time targeted by fraudsters. Also, employing history and credit score would have a socio-economic context, which would be helpful to differentiate between falsity and authenticity. Such attributes are especially precious towards the study of synthetic fraud because, in this case, sudden shifts or flaws in employment and credit conduct tend to be warning signs. Besides, transaction history is provided in the dataset, and this is vital in the identification of the behavioral patterns over time. Their record on transactions will show buying patterns, when the due date of payments was made, and how frequently the account is used, characteristics that can be proven handy when anomaly detection is being done using sequential models such as the LSTMs. The data is preprocessed before fitting it into the model through missing value imputation, normalization and encoding of categorical data so that there is consistency and scalability. The proposed hybrid deep learning framework can be applied to both training and validating using this dataset as a balanced, feature-rich backbone that would enable it to learn intricate patterns that occur in legitimate and fraudulent identity behaviors.

3.3. Data Preprocessing

Machine learning and deep learning models are very sensitive to data preprocessing, and in a fraud detection process, the data type is different, and the information sets are huge. As such, data preprocessing methods need to be more effective in enhancing the effectiveness of models and performances. The proposed framework includes three crucial preprocessing steps that are conducted to normalize, tokenize, and build a time window. Such measures ensure that the data are ordered, normalised, and have time correlation, which is then delivered into the model.

- **Normalization:** All the numerical values (e.g. credit score, dollar amounts of transactions, and number of transactions) are placed through a normalization procedure in order to provide a uniform range of values throughout the dataset as well. It is a necessary step since deep models of learning are scale-sensitive, features whose range is large have a tendency to overwhelm the learning process and affect convergence adversely. The application of techniques like min-max scaling or z-score normalization, is applied to map the features to a common scale, usually [0, 1] or a central mean with unit standard deviation. This prevents not only higher stability of the models and the speed of training, but also to minimize the impact of outliers.

- **Tokenization:** In case of textual or categorical features, such as names, emails, and names of employers, the tokens are used to convert the raw text into a machine-ready format. This is by decomposing strings into meaningful components or tokens. These tokens are sometimes transformed into numerical vectors using one-hot encoding, label encoding, or embedding layers, in other cases. Tokenization allows the model to identify trends and anomalies in textual data, including strange domain names in the email address or discrepancies in naming, such as can dof cate or mabe, both of which may be signs of synthetic identity or even fraud.

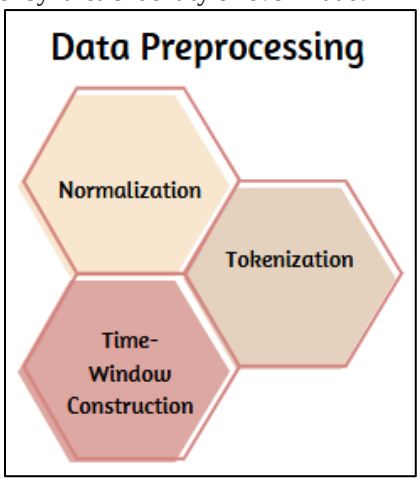

Fig 4: Data Preprocessing

- **Time-Window Construction:** As temporal data, transaction data has time-window construction to order a series of user behavior into time frames of fixed duration. The transaction history of each user is subdivided into chronological windows, i.e., hourly, hourly, etc., in the context of the application. This enables the model to study patterns as time progresses and identify unusual patterns of activity. Time-windowing has a special significance when feeding data to sequential models such as LSTMs, because it maintains the order and timing of events, which tend to be the key indicators of fraud.

3.4. Model Components

Through its hybrid architecture, the proposed model incorporates numerous elements of deep learning that focus on learning various features of fraud patterns, including spatial mismatch, temporal patterns, feature significance, and anomalies [14-18]. This strata-based system makes the model more productive in identifying both traditional and synthetic forms of identity crimes.

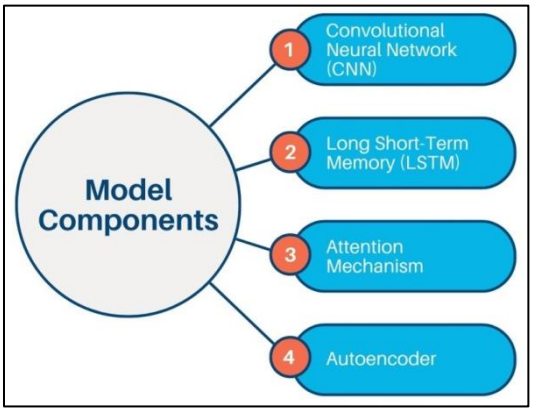

Fig 5: Model Components

- **Convolutional Neural Network (CNN):** The Convolutional Neural Network (CNN) is adopted as the initial component within the architecture to identify the spatial correlation of the input features. The unusual or fluctuating pattern in identity applications, characterised by the use of mismatched names, unrealistic addresses or jobs, or incompatible employment details, indicates numerous cases of fraudulent identities being applied. CNNs can discover such local patterns and co-occurrences through their filtering, which scans across feature maps. Such layers render the input data in the form of a set of high-level spatial features, which happen to be particularly valuable when the input format is structured, such as a table or a grid encoding of identity data.
- **Long Short-Term Memory (LSTM):** After CNN, a temporal sequence in data is modelled by the use of Long Short-Term Memory (LSTM) networks. LSTMs are more suitable to follow long-term dependencies, which is why they can

be used to analyze transaction history, patterns of logins and time series in which identities are created or amended. To illustrate, a drastic alteration in the frequency of transactions or a very sudden change, such as spending, are some of the signs that indicate fraud. These sequences are fed to the LSTM to reveal low-level temporal anomalies that may not become visible on isolated samples.

- **Attention Mechanism:** To further enhance the capacity of the model to stick to the most important information, the Attention Mechanism is implemented on top of the LSTM layer. This is a mechanism that computes the weights of attention, which establish how important a time step or a feature in a sequence is, with respect to each other. The model will be more comprehensible and correct as key events will be emphasized, i.e., an unusually large transaction or a sudden series of identity updates. The measure of attention of a certain input xi is normally computed by a formula like:

$$\alpha_i = \frac{exp\big(score(x_i)\big)}{\sum_j exp\left(score\big(x_j\big)\right)}$$

and score x_i Is the score of the input x_i, which is typically a scoring function (a feed-forward network or a dot-product).

- **Autoencoder:** The last deep learning model is the Autoencoder, which is an unsupervised neural network that is trained on only legitimate identity data. It is aimed at reconstructing the input data in the best way possible. In the case of inference, the autoencoder tries to reconstruct genuine and fraudulent identities. It is, however, worth noting that its tendency to make higher reconstruction errors with fraudulent or synthetic data can be attributed to the fact that it has only been trained to copy the structure of normal data. These errors are used as anomaly scores, in which a higher score shows a higher probability of fraud. This qualifies the autoencoder as a useful last validation layer in detecting the outliers missed by the preceding layers.

3.5. Training Process

The proposed fraud detection framework has a training procedure that aims at providing the best learning, generalization, and performance over various data distributions. The data is split into three subparts: 70% of the data is used in training, 15% in validation, and 15% in testing. Training the model is done by optimizing internal parameters in the training set to reduce the model error between the forecast and the actual values. The validation set can be used to track the performance of the model as it is being trained, to give an early indication that overfitting is occurring and to stop training before this can occur or to tune the hyperparameters. The last model is further tested on the test set to validate its generalization capacities on new information. The model uses the loss function of Binary Cross Entropy (BCE) since this is appropriate to be used in cases of classification of a couple of things, such as frauds (legitimate vs. fraudulent). BCE loss is a measurement of the difference between a class label and the estimated probability. It punishes wrong predictions harshly, particularly when the model is certain about a false choice. This loss is especially helpful when dealing with an imbalanced dataset, such as that of fraud detection, where it is important to avoid false positives and false negatives. To optimize, the model is optimised using the Adam optimizer, which is a popular gradient-based optimizer since it is a combination of the AdaGrad and RMSProp optimizers. Adam is adaptable in its learning rate per parameter, which makes it suitable when it comes to a highly complex model such as the hybrid deep learning framework suggested in the current study. The learning rate is set to 0.001, where the stability of training and speed of convergence are balanced. Such an arrangement will aid the model to converge effectively without the risk of overshooting or getting stuck in local minima. In general, the training procedure encompasses the finest practices of deep learning, which uses a combination of data division, loss minimization, and dynamic learning. This will make the final model strong, generalizable and will be able to effectively identify traditional and synthetic identity frauds.

4. Results and Discussion

4.1. Evaluation Metrics

To measure the performance and trustworthiness of the suggested framework of fraud analysis, 4 fundamental measurements were used: Precision, Recall, F1-score, and ROC-AUC (Receiver Operating Characteristic Curve- Area under Curve). When considering the situation in which the class imbalance is particularly high, such as is the case with fraud detection, these metrics provide a complete picture of the model's performance in distinguishing between the fraudulent and legitimate identities. Precision is the ratio of correctly predicted fraudulent cases to the number of all the cases spotted by the model as fraud. The false positive rate in this metric is important when operating in a situation with high risks, e.g., financial fraud detection, because having a large rate can result in a waste of time on unnecessary investigation and inconvenience to honest users. A precise model can provide the most accurate estimation of the alert, enabling the identification of suspicious alerts and directing resources towards further verification of the identified alerts. Recall, on its part, portrays the performance of the model in recording real instances of fraud. It is the ratio of the number of legitimate cases of fraud that the system was able to detect. Recall is specifically useful in fraud detection, as the incorrect detection of non-fraud (false negatives) may lead to years of financial and reputational losses. High recall gives the assurance that the model is capable of detecting even a subtle or concealed artistic pattern of malaise. The harmonic mean of the precision and recall values is called the F1-score, which can be particularly valuable in the case of a trade-off between precision and recall. It provides a single measure that balances

between false positives and false negatives, making it optimal in situations where one cannot afford either of the two errors. Finally, ROC-AUC allows determining the success of the model to separate two classes--fraudulent and legitimate--at different thresholds. The higher the AUC, the better the overall performance can be viewed as the model reliably distinguishes between classes, which is to say that it does not matter what decision threshold is. Overall, these measures provide a resilient design for assessing the performance of a fraud detection model and understanding its reliability.

4.2. Comparison vs Performance

The original Deep Learning (DL) modality was compared to two common conventional machine learning classifiers, namely Support Vector Machine (SVM) and Random Forest (RF). The experiments were performed with all of the models based on the same dataset and evaluated by applying three major metrics, Precision, Recall, and F1-score, to give a leveled assessment of the investigated models in the detection of fraud.

Table 1: Performance Comparison of Models

Model	Precision	Recall	F1-score
SVM	76%	72%	74%
Random Forest	82%	77%	79%
Proposed DL Model	91%	88%	89%

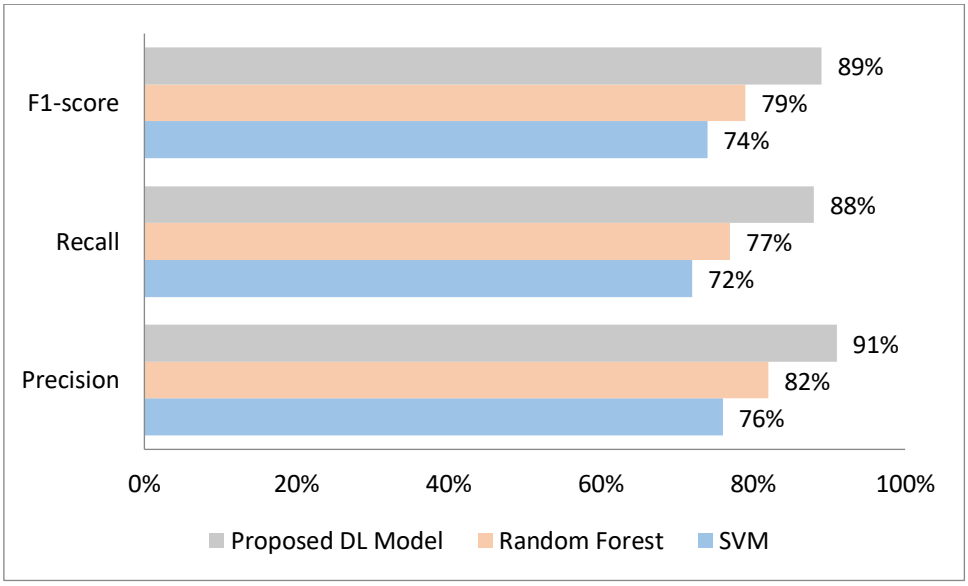

Fig 6: Graph representing Performance Comparison of Models

- **Support Vector Machine (SVM):** Predictions made by the SVM model gave out a precision of 0.76, a recall of 0.72 and an F1-score of 0.74. These findings show that although SVM can, to some extent, identify fake identities, it cannot detect all cases of fraud. Model recall is relatively low, meaning that a significant amount of fraud cases are not detected by the model, which might be an issue in the case of fraud detection, where false negatives should be reduced. Also, the reasonable accuracy reveals that it has false positives, making it less true in the practical world.
- **Random Forest (RF):** Random Forest had a better performance than SVM, with a precision of 0.82, a recall of 0.77, and an F1 score of 0.79. Such enhancements are due to the ensemble nature of RF, which involves several decision trees to get complex patterns in the data. In the model, there is a more even performance, whereby there were fewer false alerts and missed cases of fraud. It is not yet deep and dynamic enough to learn to capture a fast-changing behavior that is characteristic of synthetic identity fraud, however.
- **Proposed Deep Learning Model:** The proposed deep learning model achieved an accuracy of 0.91, a recall of 0.88, and an F1-score of 0.89, which was quite high compared to traditional classifiers. Such performance shows the model manages to learn very complex and nonlinear associations based on its hybrid architecture, which includes CNN, LSTM, the attention mechanism, and autoencoders. High precision means that there are only a few false positives, whereas a high recall means that most cases of fraud are properly detected. The poised and decent F1-score is evidence that the model provides solid and scalable findings in fraud detection, and therefore, a very effective one in terms of real-life implementations.

4.3. Discussion

It can be seen that the performance evaluation indicates the dominance of the suggested deep learning model in comparison with the classic classifiers (Support Vector Machine (SVM) and Random Forest (RF)). The credit for this success goes to the model's hybrid layered architecture, which enables it to capture complex, spatial, temporal, and contextual

relationships of identity data —a feature largely overlooked by traditional models. A major part of the model that contributed tremendously to the accuracy of the model is the attention mechanism. The attention mechanism also enabled the model to focus on meaningful indicators of fraud by attaching more weight to some features, such as mismatched pairs of names and Social Security Number (SSN), suspicious email addresses or drastic shifts in employment history. This narrowed-down analysis allowed limiting the number of false positives and made the model more confident in detecting suspicious applications.

Along with the precision, high recall results were obtained with the model, which in particular was owed to the presence of the autoencoder module. Only trained on legitimate identity data, the autoencoder was observed to be able to reconstruct legitimate patterns of genuine applications with a high level of accuracy. When introduced to synthetic or manipulated identities, the reconstruction error became considerably larger, which flawlessly indicated the anomaly detection. This gave the model the capability to identify cases of fraud that need not involve behavior that is outright abnormal, yet they would not conform to the norm that was learned. A combination of these two strategies, attention-guided feature emphasis and reconstruction-based anomaly detection, created a solid defense against both conventional and synthetic identity crimes. In general, the suggested deep learning model can be characterized, on the one hand, by good performance according to the most important evaluation metrics and, on the other hand, by such practical benefits as scalability, adaptability, and interpretability. The fact that it derives subtle and never-imaged patterns of fraud justifies the success of hybrid deep architectures when it comes to building practical identity verification systems that can distinguish subtle and novel fraud that rule-based or shallow learning models could not detect.

5. Conclusion and Future Work

This paper proposes an effective and smart deep learning system that has been developed against the detection of punched identity on online credit cards- an increasing issue in the financial and e-commerce sectors. The proposed model takes advantage of a hybrid framework which integrates Convolutional Neural Networks (CNNs) to learn spatial patterns in the data, Long Short-Term Memory (LSTMs) networks to learn behavioral sequences and Autoencoders (AEs) to learn about anomalous features in an unsupervised fashion. When combined, these can make the system analyze various types of data such as personal identity information, timelines of transactions and behavior inconsistencies. The addition of an attention mechanism helped the model to be more precise as it would selectively attend to the high-risk features, e.g., matching of names and SSNs, unusual employment History- these are features typically not captured in traditional fraud-detecting systems. The autoencoder, further modeled on only legitimate identities, was also useful in identifying unusual input data, even though the reconstruction error in identifying such data was high, thus enabling the system to detect even some minor form of fraud that did not match a pronounced pattern. The model also performed better in all main metrics of the evaluation-precision, recall, F1-score, and ROC-AUC-proving that it can outperform classic models of machine learning classification, such as Support Vector Machines (SVM) and Random Forest (RF). The findings support the adoption of deep learning models in contemporary fraud detection mechanisms, particularly with the rampant and more advanced synthetic fraudulent activities.

Although the existing model is showing good performance, there are various sources of improvement in the future. A potential approach is to apply a Graph Neural Network (GNN) to identify relationships and associations between identities, devices, IP addresses, and transaction sources. Many graph structures fit into synthetic identities, as they tend to have hidden relationships between them, e.g., shared phone numbers or co-applicants' addresses. When GNNs are incorporated, the system would be able to know and raise alerts about these networked anomalies. The other salient improvement is the construction of a real-time streaming fraud detection system that was able to process data in real-time. This would be achieved by transforming the existing batch processing model to support a streaming architecture, utilising frameworks like Apache Kafka or Flink to enable real-time alerts for potential fraud and enhance responsiveness. Lastly, the actual implementation of the system in operating banks is a very important step on the way to operationalize the research. These would need issues of data privacy, scalability, compliance and how to integrate with the existing fraud prevention systems to be addressed. In developing further in such directions, the framework may become an industry-grade, production-ready solution to the problem of fraud detection.

References

[1] Bhattacharyya, S., Jha, S., Tharakunnel, K., & Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. Decision support systems, 50(3), 602-613.

[2] Phua, C., Lee, V., Smith, K., & Gayler, R. (2010). A comprehensive survey of data mining-based fraud detection research. arXiv preprint arXiv:1009.6119.

[3] West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: a comprehensive review. Computers & security, 57, 47-66.

[4] Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri, F. (2019). Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. Information Sciences, 479, 448-455.

[5] Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P. E., He-Guelton, L., & Caelen, O. (2018). Sequence classification for credit-card fraud detection. Expert systems with applications, 100, 234-245.

[6] Chalapathy, R., & Chawla, S. (2019). Deep learning for anomaly detection: A survey. arXiv preprint arXiv:1901.03407.
[7] Srivastava, A., Kundu, A., Sural, S., & Majumdar, A. (2008). Credit card fraud detection using a hidden Markov model. IEEE Transactions on dependable and secure computing, 5(1), 37-48.
[8] Ngai, E. W., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. Decision support systems, 50(3), 559-569.
[9] Ali, I., Aurangzeb, K., Awais, M., & Aslam, S. (2020, November). An efficient credit card fraud detection system using deep-learning-based approaches. In 2020 IEEE 23rd International Multitopic Conference (INMIC) (pp. 1-6). IEEE.
[10] Thejas, G. S., Dheeshjith, S., Iyengar, S. S., Sunitha, N. R., & Badrinath, P. (2021). A hybrid and effective learning approach for click fraud detection. Machine Learning with Applications, 3, 100016.
[11] Rehman, A., Naz, S., Razzak, M. I., Akram, F., & Imran, M. (2020). A deep learning-based framework for automatic brain tumour classification using transfer learning. Circuits, Systems, and Signal Processing, 39(2), 757-775.
[12] Akhilomen, J. (2013). Data mining application for a cyber credit-card fraud detection system. In Advances in Data Mining. Applications and Theoretical Aspects: 13th Industrial Conference, ICDM 2013, New York, NY, USA, July 16-21, 2013. Proceedings 13 (pp. 218-228). Springer Berlin Heidelberg.
[13] Kou, Y., Lu, C. T., Sirwongwattana, S., & Huang, Y. P. (2004, March). Survey of fraud detection techniques. In IEEE International Conference on Networking, sensing and control, 2004 (Vol. 2, pp. 749-754). IEEE.
[14] Raghavan, P., & El Gayar, N. (2019, December). Fraud detection using machine learning and deep learning. In the 2019 International Conference on Computational Intelligence and Knowledge Economy (ICCIKE) (pp. 334-339). IEEE.
[15] Mubalaike, A. M., & Adali, E. (2018, September). Deep learning approach for an intelligent financial fraud detection system. In 2018, 3rd International Conference on Computer Science and Engineering (UBMK) (pp. 598-603). IEEE.
[16] 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.
[17] Osegi, E. N., & Jumbo, E. F. (2021). Comparative analysis of credit card fraud detection in simulated annealing trained artificial neural network and hierarchical temporal memory. Machine Learning with Applications, 6, 100080.
[18] Ding, L., Fang, W., Luo, H., Love, P. E., Zhong, B., & Ouyang, X. (2018). A deep hybrid learning model to detect unsafe behavior: Integrating convolution neural networks and long short-term memory. Automation in construction, 86, 118-124.
[19] Liu, H., & Lang, B. (2019). Machine learning and deep learning methods for intrusion detection systems: A survey. applied sciences, 9(20), 4396.
[20] Salur, M. U., & Aydin, I. (2020). A novel hybrid deep learning model for sentiment classification. IEEE Access, 8, 58080-58093.