

Sequential Pattern Recognition and Entity Linkage in Payment Ecosystems: AI-Driven Computational Frameworks for Financial Transaction Behavioural Analysis

Dr. Agata Grabowska, Associate Professor of Computer Science, Wrocław University of Science and Technology, Poland

1. Introduction

The banking system has been overwhelmed by data since its origins. Historical transaction records cover millions of items from varied categories and dimensions. From deposits to loans, clients' transactions have to run through a myriad of controls aiming to ensure compliance with regulations and risk assessment purposes. Historically, humans were the first and exclusive group to analyze these records, but the rapid growth of electronic and intelligent systems has allowed the delivery of increasingly complex analytical and decision-making capabilities and the execution of tasks originally designed for human beings. From the introduction of a sales and cash management system developed in 1964 to artificial intelligence enabled capabilities, the range of tasks was drastically expanded, supporting the branches' transactions and reporting services, balance analysis, and credit risk profile assessment and segmentation, in order to develop more targeted products.

In an environment dominated by IT development and reducing the cost burden of AI technologies, transaction processing has moved from early daily manual processes to automated processes covered by 24/7 high-performance service level agreements. Nevertheless, the hidden and unexploited electronic traces of these transactions partially mask valuable information that currently does not leave the corporate repository walls. The omnipresence of client historical transaction records in banking supports AI system implementation as a valuable source of information to develop solutions that increase value for banks and clients, on a variety of aspects, either through process optimization, increased risk assessment, or the development of new revenue sources, as we will

discuss in this chapter. Such advantages are also expanded to the creation of more tailored credit products, since a historical record of transactions, together with access to social network resources, a deposit and transactional bank, or payment association can reduce asymmetric information between banks and applicants, leading to fair and personalized rates.

1.1. Background and Significance

At the core of modern finance, regardless of the palette of participants or nature of financial instruments, lies a myriad of transactions. They lie at the very core of the creation of value and influence the decision-making taking place across all levels in every entity. From an employment contract in an individual micro-enterprise to complex interdependent chains of credit underpinning the activities of the largest corporations in the world, through the exchange of contracts facilitated by intertwined networks of intermediaries, the financial transaction cycle is the heartbeat that makes the market societies tick. While the early philosophers on the nature of money made it clear that both money and credit had to be interoperable to settle their debts, thus—by extension—separate entities representing binary-coded states of a transmission mechanism that is shared between them, it was only after the digital revolution of the late 20th century that the behavior of actors manifested in transactions themselves became the focus of not just regulation, but also automated oversight using artificial intelligence methods.

In earlier stages of computational civilization, that conversation revolved almost exclusively around money laundering. However, shortly under two decades into the new century, it begins to be clear that other risks are lurking, not just within the shadows, but within every single transaction. Addressing these risks again calls for the development of new tools. We investigate how the tasks of ensuring the robustness of markets from the perspective of capital markets supervisors can be approached through the usage of AI in the form of end-to-end deep learning in this work. Industries are in the process of digitalizing many of the tasks connected with the implementation and oversight of filters which "real-world logic" dictates via a variety of what one might group as "traditional" filters, ascertaining symmetric dynamics between similar classes of financial transactions. Oftentimes, automatic, expert-free, and scalable AI approaches performing at multi-layer filtering of routing-related transactions on an end-to-end basis are eagerly sought by both the market participants and the regulators alike. Every

regulated entity is implementing these filters. Market supervisors also routinely implement joint filters and control the effectiveness of those implemented by the regulated entities, seeking to ensure harmonious operation of the financial fabric in time and space. While all that is being done and has been performed with various degrees of success, a new set of risks is emerging within the transactions that have already passed through the filtering blocks as identified above. These risks, usually emanating from side channels of vulnerabilities within the ecosystem, often require other means to be mitigated. Deep learning is the current top technology performing in this sphere.

1.2. Research Objectives

The underlying research pays close attention to developing an artificial intelligence-driven paradigm that accomplishes continuous automatic financial transaction history analysis. This addresses three different challenges associated with real-world financial transaction data: effective data encoding, modeling of the behavior of future standalone transactions, and meaningful segmentation of transactions within each account. Building on modern AI research, adversarial learning and sequence-to-sequence modeling were used to solve the data encoding problem, while recurrent neural networks are leveraged as an efficient means to model complex transactional patterns. Additionally, self-attention mechanism-based refreshable sequential methods were developed to break down historical continuous transaction sequences to reflect on incremental recurrently revisited topics with minimal assumptions. Based on the proposed methods, a transaction pattern recognition system is developed to assist commercial institutions in risk control, marketing, and customer relationship management.

2. Theoretical Framework

A significant body of research, theories, and computational methods have been developed to study AI model development and understanding. The integration of finance and AI is a multilayered, multidisciplinary science. AI is the automation of rule-driven analytics in finance. The core technology underpinning AI is machine learning. Machine learning develops algorithms to progressively improve predictions. There are five core principles in developing a machine learning model for AI. To begin with, data is integral to the machine learning algorithm's learning, and it is the nutrient of AI, so a clean set of data is absolutely central to developing an effective AI model. Next, developing a model based on data that is finite in size allows the algorithm to learn from

the training set, and the resulting model is primarily based on the data it has been trained on.

Third, the AI optimization model is, at its core, a search algorithm that searches a set of possible solutions or features, in contrast to a rule-based programmed algorithm. Machine learning is fundamentally a probabilistic algorithm; the statistical measure of uncertainty adds rightful context. Finally, machine learning, such as neural networks with internal mechanisms, is much more complex as they are effectively 'black boxes' that make it possible for nonprogrammed decisions. This approach makes it difficult to make fine-tuned adjustments to machine learning algorithms at critical decision points. In summary, when developing an AI model, it is crucial to understand the theoretical framework, which is based on the predictability of the model. An integrated review of the financial literature and the heart of AI, such as algorithms, data processing, and predictive analytics, is essential. There are many machine learning approaches; presently, the most utilized in many financial applications. A review of the finance literature has provided a comprehensive understanding of hypotheses' directionality and popularity. This framework is to be used to understand how and what the expected impacts of AI on finance are.

2.1. Machine Learning in Financial Transactions

Today, it is hard to find transaction data applications that do not involve machine learning. Developers and researchers can choose from a wide range of machine learning techniques and algorithms since many of them have shown great potential in handling transaction data. There have been in-depth studies on support vector machines for the supervised learning paradigm, which estimates a mapping function that correlates input data with corresponding output labels; deep learning models, particularly long short-term memory and convolutional neural networks, for applying complex nonlinear mappings to transaction data to learn from regularity and incipient financial activities as well as prediction tasks; and unsupervised clustering-based methods, such as k-means, for identifying groups based on specific features in the transaction dataset.

In a broader categorization, these models belong to one of two main types: supervised and unsupervised learning. Supervised learning trains a model using input data along with the desired or expected output data, where the model learning algorithm iteratively devises a hypothesis function that correlates the input data to a set of output

labels to minimize the prediction error. Regarded as a forecasting system, for instance, one can leverage supervised models to learn regular trends of input-output data on a sound footing. The developed supervised models have prospered in extensive finance applications, covering sentiment analysis, predicting volatility, evaluating portfolios, and estimating returns. Unsupervised learning, in contrast, is a learning framework that does not require known output labels for input data. Typical unsupervised learning behavior is minimizing the difference between input data attributes or features by grouping them according to similarity or mutual dependency. Known as clustering, it has found a range of applications in financial analytics, such as customer profiling based on transaction data.

Using machine learning methods to analyze transaction data provides several benefits in terms of predictive capabilities. Theoretically, machine learning is capable of capturing complex interdependence between input and output data, as opposed to traditional quantitative analytics that relies heavily on simplifying financial models. Machine learning does not generally suffer from the "noise" problem often viewed in econometric models or financial pricing methods, so it is less sensitive to underlying assumptions about the data or financial conditions. Likewise, machine learning models are often superior in pattern recognition, which can be beneficial in assessing subtle financial indicators and activities, or in daily consumer profiling and market basket analysis in retail. Despite the merits, the adoption of machine learning methods is not without challenges. In particular, their performance is sensitive to low-quality data. They require extensive training in the presence of vast amounts of data to accurately learn implicit patterns. Also, most machine learning models tend to be complex and difficult to interpret, making them less robust for financial institutions, particularly those that are subject to stringent regulations.

3. Data Collection and Preprocessing

In the financial technology market, many diverse sources of financial transaction data exist from which these datasets could be leveraged. The primary consideration is to resolve the differences between models applied to them and assess their effect on model performance. Financial transactions are natural entities for which large volumes of data are available, and collecting multifaceted data is discussed as well from numerous sources. Despite the complexities of financial data, the approach described should

nonetheless aid in resolving the mismatch between multiple reference data sources to create a master reference dataset.

Data preprocessing refers to a data mining step where data collected is converted into a utilizable format. Within a business context, it may represent a time at which deals are formalized with the initial focus placed on securing the best possible data representation. Collecting data from different sources, in the context of finance, usually results in large, complex, and disparate datasets. Throughout this discussion, typical financial benchmarks are employed. Typically, financial data is cleaned, normalized, and transformed in preparation for data analysis. Data cleaning involves tasks such as detecting errors and inconsistencies, dealing with missing data, and removing spikiness from the data. Parameter transformations, such as log and root transformations, are also commonly used to transform a dataset when it is asymmetric or skewed. In the case of financial transaction data, parameter transformation would be natural as one would expect that a Pareto distribution would arise in this dataset. Finally, feature engineering is employed to identify a set of key summary measures identifying patterns within the data and is useful for identifying clusters for any anomalies or outliers in the transactions. All aspects of this preprocessing stage play a role in improving the performance of an analysis.

3.1. Sources of Financial Transaction Data

Financial transaction data can be obtained from various sources. Banks are typically where the most detailed transaction data can be found, in particular, data on direct debits and standing orders. In addition to banks and credit card companies, other online payment platforms, such as electronic wallets, also have transaction data that can be relevant for financial decisions, for example, for fraud detection in online shops. Finally, many transactions or post-transaction data are also stored in merchant databases and payment systems, such as especially stock trading systems and webshops. These vast merchant and online systems often contain the data points necessary to make credit decisions and have the advantage of being already digital and not needing any physical transformation. They can also save the historical product-buying behavior of a customer, which might be an important factor in order to assess sustainable investments. The data sources vary substantially in data quality, data privacy issues, accessibility, and cost of data acquisition. In general, the more directly related a data source is to finance and

transactions, the higher the quality of the financial transaction data. Furthermore, the data collected from these sources are most relevant in terms of financial decision-making. For example, identifying a purchase in a webshop that sells electronics affirms the transaction that was agreed upon in the contract. The comprehensive data needed for training an AI/ML model, especially for matching multiple user-generated messages to a single transaction, is typically found in the databases of contract parties. Thus, at stake are not just potential biases in the predictive performance of AI models when leveraging the data, but also questions pertaining to fairness. The emergence of numerous data protection laws and regulations also restricts the ability of researchers to access the transaction data with consent. Given these complexities, it is important to strategize with respect to data collection quality to ensure the ability to effectively train AI and machine learning systems across the key tasks.

3.2. Data Cleaning and Feature Engineering

One of the most challenging, yet very important, procedures in data analysis is data preparation, a process that is believed to take up the majority of an analyst's time. Specifically for financial transactions, which are prone to generating numerous inconsistencies and inaccuracies, data cleaning is of utmost importance. These inconsistencies in a financial database are mostly brought about by error-prone data entry, system glitches, broken transmission links, improper software installation, and user errors. However, in a financial transaction dataset, especially one related to real operations, simply removing it loses the proper pattern and behavior present in them. Thus, data imputation techniques may be used in order to substitute the missing values of transactions.

Covariance-based graphics can be simple and powerful tools in visualizing and distilling large feature sets into simpler, more workable charts. Feature engineering, on the other hand, is a very crucial part of a predictive model. It can significantly impact the outcome of a predictive model. If it is done correctly, it can enhance model performance, thus improving the model's predictive capabilities. In addition, it is a process of selecting and transforming independent variables when we create new features out of existing features to train a machine learning algorithm. Some techniques and methods can be used in feature engineering. A variety of methods can be used in feature reduction. However, as much as it is crucial to identify variables that are important for a model,

one has to consider balancing the model's complexity and interpretability. Feature selection, when effective, will simplify the task of the learning algorithm by creating algorithms that are much easier to understand. Cleaning and preparing the dataset before developing a machine learning model is a critical area in AI-driven financial transaction analysis.

4. Model Development

The sample model development would focus on how to create two machine learning models to analyze financial transactions. The algorithms chosen for the models depend on the nature of the data to be analyzed, as well as what sort of transactional behavior or financial event the user wishes to identify. Group 1 models are meant to be used when the exact target is unlabelled or unknown. This type of model is useful when one does not have data on specific instances of fraud but wishes to identify this or something similar in a large dataset. Techniques used in these types of models include the family of clustering or unsupervised learning techniques, some examples of which are K-Means, DBSCAN, Gaussian Mixture Models, and Self-Organizing Maps. Supervised learning models, as represented by Group 2, simply create a mathematical function or equation that we can use to model a relationship between various features within a typical financial transaction.

No significant cross-prediction is learned unless they are paired – this means that where a transaction is labeled fraud, the models are learning to predict that transaction as fraud or not; various other attributes remain neutral. Although the algorithms differ in how they take the data, both represent gradient boosting models. GBM is a technique comprised of a series of decision trees trained on errors from the previous trees; thus, the overall prediction was an aggregate of all of the trees. The first GBM uses a mixture of deep and shallow trees grown from randomly selected built-in tree splits. This is very useful for model deep learning. Once a model is trained, the next step is to evaluate its performance. To understand how well a model predicts fraud in financial transaction data, institute a performance metrics framework in comparison with the other models described in the next section. Model performance will be measured with accuracy, precision, and recall. This evaluation illustrates that the configured Z score has received an average overall maximum true negative rate predictive value of 3.7 percent above the other configurations.

4.1. Supervised Learning Algorithms

The hallmark of supervised learning is that models are trained on labeled datasets, in the course of which they capture associations between input data and output values. As a result, the above models can uncover and predict hidden patterns in the unlabeled test data. Supervised learning algorithms are extensively used in the area of financial data analysis. Considering the intrinsic characteristics of financial data, a variety of supervised learning algorithms have been implemented for data modeling and prediction. Each of the aforementioned algorithms has its strengths and weaknesses regarding such financial data features as non-linearity, multi-modality, and outliers. The challenges of overfitting data and underfitting in machine learning should be handled with appropriate strategies for model validation and feature selection. A poor feature choice can lead to spurious correlation, decreased prediction accuracy, or increased overhead in computational expense.

The consequential output of machine learning models enhances the quality of decisions ranging from customer loans and fraud detection to investment management. The potential applications of supervised learning to problems in financial economics and econometrics have initiated considerable interest in this research area. Vast amounts of historical data are available for analyses of financial market data and economic time series. Apart from the estimation of trading strategies, the analysis of financial time series can provide means for risk management, quality of investment opportunities, and measures of the diversification of investment portfolios. In addition to financial market data, supervised learning algorithms can be exploited to predict real economics, such as stock prices, house prices, sales values, etc. Overall, the principal motivation for the use of supervised learning algorithms is their potential to enhance the decision-making process in finance, as well as issues in business, economics, and management.

4.2. Unsupervised Learning Algorithms

Unsupervised learning refers to creating models that analyze input data based on their intrinsic structure and significant features. This type of learning is very relevant in managing new data since it does not require labels, which is very useful for financial institutions, as algorithms can explore a vast amount of data to find patterns not yet discovered. The data patterns described in the transactions consist of individual behavior profiles, most used routes, places where the card has been used the most,

explicit user preferences, etc. The above can be applied using a variety of clustering methods, ranging from variance and distribution analysis to quantitative and/or qualitative information about the transactions. There are different categories where this kind of clustering applies: partitional clustering, hierarchical clustering, declustering, and directional clustering, etc. Each of them offers different ways to divide the information that allows for generating subgroups and analysis from the case itself.

For example, k-means and hierarchical clustering could be useful when the number of customers increases and needs to be segmented into small groups in order to improve marketing strategies and the identification of the most valuable customers for the organization. In addition to subgroups, anomaly detection methods can also be applied since it is crucial to avoid fraudulent transactions or even money laundering activities. This kind of approach encompasses some challenges when interpreting and validating the clusters through the areas or the marketing departments. This is important work that can be done in future research that combines supervised and unsupervised models to validate and strengthen the results. Unsupervised learning, with clustering algorithms, has a relevant application for financial transaction data. This type of algorithm, without using labels, can help produce new financial datasets and open up new research areas that can be used alongside supervised learning.

5. Performance Evaluation

Performance evaluation of the developed AI-driven models is necessary to ensure the model's ability to make valid predictions. For classification tasks, various metrics such as accuracy, precision, recall, and F1 score are used to reveal different aspects of model prediction. Precision measures the fraction of actual positive results among those positive results predicted by the model, whereas recall measures the fraction of the total amount of relevant instances that were actually retrieved. An F1 score is the harmonic mean of precision and recall, and tells us the accuracy rate regardless of class distribution. In the context of financial transactions, the strongest selling point about precision, recall, and F1 score is that these indicators allow for making robust back- or straight-through processing decisions. As financial transactions could be one of the main revenue sources for any exchange or banking industry, reliability and precision are important characteristics.

For an accurate and unbiased evaluation, validating the AI-based model through cross-validation, rather than using train/test, is important to demonstrate the model's global performance. Although cross-validation techniques consume a huge amount of data, they show results that are more stable with a low variance. Lastly, a critical performance assessment of different models using different performance yardsticks is necessary to know which model performs better across classification problem domains. A joint analysis of the accuracy, precision, recall, and F1 scores of different models provides a comprehensive assessment report on the performance of existing violations in order to rely on learning methods. The performance yardsticks and their joint investigation help us to precisely understand the strengths and weaknesses of existing classifiers and choose a suitable AI-based learning method for financial transaction analysis. In summary, decision makers and industrial experts could dwell upon the performance evaluation of different models to analyze the most suitable classifier.

5.1. Metrics for Model Evaluation

This subsection introduces specific performance metrics to evaluate machine learning models for financial transaction analysis. In addition to predicting a transaction's class label, the underlying machine learning model should also generate per-instance probabilities, as this output is required for fraud prevention systems. Model evaluation provides point estimates that do not necessarily denote the true quality of the model. A per-instance fraud probability is useful for decision-making, such as setting the operating point of a system or resolving alerts by humans. In practice, misclassification costs exist when using thresholds on fraud probability estimates to convert them into class predictions. Without explicit modeling of misclassification losses, each threshold can be considered a separate problem. However, comparing models directly is difficult in the absence of good criteria or scoring mechanisms.

By explicitly modeling the misclassification costs, one can mitigate some of the limitations associated with other metrics. The advantages of subsetting methods are clear, and as a lowest common denominator, using the confusion matrix aggregated from all thresholds can be a simple yet effective approach. In this subsection, we decode several well-known performance metrics for model evaluation. Moreover, we discuss potential pitfalls and inappropriate interpretations of these metrics on imbalanced and unclean datasets. It is crucial to choose the right performance metrics as per the

modeling objective, and we propose practical steps for choosing the right metric in real-life situations. The suggested metrics selection is demonstrated through an example related to fraud detection in the finance sector. Model evaluation is crucial for developing successful machine learning applications, especially in the financial industry.

5.2. Comparative Analysis of Models

The final step of the experimental evaluation involved a comparative analysis of the ten machine learning models or algorithms. This was done to understand their capability to solve the problem of financial transaction analysis. Our aim was to approach this from a scenario identification lens rather than a blanket comparison. In general, the more advanced machine or deep learning models were not advantaged by the inclusion of univariate LSTM networks. Nevertheless, they performed well in modeling the task. If the focus is on prediction, LightGBM or a single LSTM was found to be the most performant and should be favored due to lower specificity at trading positive instances. RF also demonstrated strong performance in the base scenario.

In comparing the six basic algorithms, given the real-world utility lay in identifying transactions, a shallower cut-off was considered, of 80% specificity. RF outperformed the other basic algorithms in their out-of-the-box performance in this setting, within the short and base scenarios. All algorithms exhibited worsened performance overall in the complexified, market-affected scenarios and models, underscoring the lack of generalizability of the generated base models and advocating for events-based updating in a real-world setting. Regarding complexity and interpretability, the most complex basic algorithm, in terms of complexity, was the LightGBM and the simplest was OLS or a logistic regression. It is not clear whether the gain in out-of-the-box predictive performance is greater than the handicap in interpretability and complexity. This will depend on the relevance of different metrics to the context in which this research is deployed. Overall, for simple models, there was a large trade-off between the AUC-ROC and BV-TNR of the models. We were unable to include the LSTM cells in the most complex algorithms because of the excessive time required for training. The RF and LightGBM showed improved outcomes, with higher model complexity, in the L3 and L5 scenarios, and an improved long-run predictive capability over the GIF.

6. Future Direction

In the research over the past two decades, many AI-based financial transaction analysis models have been proposed, achieving some success. A few AI techniques have already been integrated or have the potential to further improve the performance of the existing transaction analysis models. For example, deep learning-based attention mechanisms can be employed to filter noise and generate informative data for machine learning models. In the future, more deep learning techniques can be introduced, especially for big data analysis in finance. The emergence of blockchain technologies will continuously increase the demand for related tools and services. In such cases, neural network-generated embedding techniques potentially provide new opportunities for sentiment analysis or act as predictors of stock prices. Furthermore, the rapid development of the futures market presents new methodologies and applications for financial transaction analysis. In addition, we argue that with the evolving financial regulations, one important direction for future research should be the ethics or legal aspects of applying AI in finance.

Today, AI techniques are popular in finance-related problems, for instance, blockchain-based AIs and trading strategies. In this changing environment, continuous research and innovation are needed to take the lead. Interdisciplinary collaboration between technology and finance can bring more insights into, and a better understanding of, market changes and patterns by leveraging each other's knowledge and techniques. For example, a household finance company should not only hire experts in finance but also experts in customer requirements and AI technologies. Some customer systems can optimize the best credit policies by leveraging state-of-the-art machine learning models. Academia, including finance, economic behavior, and technology backgrounds, is in a unique position to seize research and educational opportunities in this area. In conclusion, significant opportunities exist in leveraging AI models for financial transaction analysis. With improvements in AI algorithms, the transactional data from big financial data are becoming increasingly complex to model and predict. The challenges include using big data from both the volume and the variety of the data.

7. Conclusion

In conclusion, this study demonstrated the power of AI technologies when applied to the field of financial transactions. It showcased how conventional financial practices

could be transformed to perform a detailed and accurate analysis of transaction activities. The significance of these results in both the academic and business spheres can be inferred given the interest in fluctuations and ever-changing financial transactions. Nevertheless, the difficulty of the problem being considered may result in solutions that are less robust than one might expect. To counter this, with a touch of caution, researchers and developers of such models must constantly evaluate the validity of their proposed methods through real-world challenges. This study also encountered several bottlenecks that elucidate the inefficiency of previous approaches. The continual increase in the size of transactional datasets highly affects the capabilities of traditional machine learning models. This results in the need for adaptable AI algorithms that respond promptly to changes. The ethical use of these tools is also an issue that needs to be considered. Hence, this field is constantly in a state of innovation and will remain so for the foreseeable future. The core technology and methodology in this study remain in its early stages. Our understanding of these developments is also relatively basic. As such, the financial sector will continue to change in line with the growth and findings of researchers in this area.