Towards the Identification of Money Laundering in Bank Transactions using Recurrent Neural Networks

Isadora Xavier
Universidade Católica de Pernambuco
Recife, Brazil
isadora.00000844511@unicap.br

Edvaldo Veiga
Polícia Civil de Pernambuco
Recife, Brazil
edvaldosantos.veiga@policiacivil.pe.gov.br

Andrea Maria N C Ribeiro
Universidade Federal de Pernambuco
Recife, Brazil
andrea.marianogueira@ufpe.br

Rodrigo de Paula Monteiro
Universidade de Pernambuco
Recife, Brazil
rodrigo.monteiro@poli.br

Marcos Oliveira
University of Exeter
United Kingdom
M.A.Oliveira@exeter.ac.uk

Rodrigo Amaro
Universidade Federal Rural de Pernambuco
Recife, Brazil
rodrigo.amaro@ufrpe.br

Diego Pinheiro
Universidade de Pernambuco
Recife, Brazil
dmpfs@ecomp.poli.br

2024 IEEE Latin American Conference on Computational Intelligence (LA-CCI) | 979-8-3503-7457-5/24/$31.00 ©2024 IEEE | DOI: 10.1109/LA-CCI62337.2024.10814854

***Abstract*—Money laundering has a significant impact on both society and the economy, costing up to 2 trillion dollars of the global gross domestic product. Detecting money laundering involves analyzing banking data to identify its occurrence and its typologies, such as the typology I-d (i.e., fragmenting cash transfers to disguise the total value). The state-of-the-art detection methods can estimate the probability of classifying a client as suspicious based on financial reports, but they fail to identify specific typologies in a large volume of banking transactions and do not determine which transaction groups contribute to identifying a particular money laundering typology pattern. In this work, we propose automating the identification of money laundering typologies using banking transaction data represented as multiple time series. We use anonymized data provided from the Civil Police of the State of Pernambuco, Brazil, and trained recurrent neural network models to identify typology I-d in the data. The models include SimpleRNN, LSTM, and GRU. Our results show that LSTM and GRU models achieved the highest AUC-ROCs of 0.75 (0.74, 0.76). Our work demonstrates the potential for developing an automated tool to assist in identifying I-d typology, suggesting future work may focus on identifying other typologies and training models using different neural architectures.**

***Index Terms*—money laundering, banck transactions, recurrent neural networks, LSTM, GRU.**

I. Introduction

Money laundering inflicts significant damage on society, impacting between 2% and 5% of the Global Gross Domestic Product (GDP), which amounts to approximately 800 billion to 2 trillion dollars [1], [2]. Since the Vienna Convention (1988), the first global treaty to combat money laundering, numerous countries have joined forces to counter such practices. However, the challenges remain substantial, as laundering methods constantly evolve and become increasingly sophisticated, requiring criminal prosecution agencies to continually enhance and develop techniques to identify these patterns quickly [3].

To prevent and assist in suppressing money laundering, the Financial Action Task Force (FATF), established in 1989 during a G7 summit, recommended that signatory countries create Financial Intelligence Units (FIUs). These FIUs, which have national jurisdiction and operational autonomy, are tasked with receiving reports of activities identified as money laundering by gatekeepers. These reports are then forwarded to criminal prosecution agencies to assess the possibility of initiating a criminal investigation. This triangular interaction among gatekeepers, FIUs, and prosecution agencies form a system designed to prevent and suppress money laundering, aiming to stop illicitly sourced money from being integrated into the financial system as if it were legitimate. One typical activity reported to the Brazilian FIU—COAF, the Financial Activities Control Council—is the type I-d behavior, which is defined as the act of fragmenting deposits to hide the total value of the transaction, a practice known as smurfing.

State-of-the-art techniques in money laundering identification involve the use of statistical methods, data mining, and supervised and unsupervised machine learning methods [4], [5]. The use of deep neural networks has grown due to their capacity to analyze and model bank transactions as sequential data, enabling the detection of suspicious activities more efficiently than traditional methods [2]. Recurrent Neural Networks from SimpleRNN to Long Short-Term Memory(LSTM) and Gated Recurrent Units are suitable for learning temporal patterns in sequential data [6]. While some approaches focus at the individual level, another research direction has been to explore network-based approaches, accounting for collaborative group crime at the community level [4], [7], [8].

979-8-3503-7457-5/24/$31.00 ©2024 IEEE

However, there are still gaps in the study area that need to be addressed. One of these contexts is that, although there is a large amount of banking transaction data that exists in real situations, there are models that do not handle this high dimensionality [9]. Also, another challenge is the number of experiments carried out using synthetic data, which can present different results with real data [4], [10]–[12]. Moreover, major challenges include handling label bias [4], as well as the need for more public datasets [2].

The availability of bank transaction data from a group of suspects, obtained by lifting their bank secrecy, allow the civil police to automate the identification of money laundering typologies, thereby shortening the duration of the investigative process. Given that current approaches mainly use single bank transaction data instead of capturing bank transactions from all financial institutions involved and estimate a propensity for an individual to be reported to the Brazilian FIU instead of identifying a specific money laundering typology, they are mainly used by financial institutions and are not suited for adoption by the civil police for a money laundering investigation.

In this work, using anonymized bank transactions provided by the civil police of the state of Pernambuco, Brazil, we have developed recurrent neural network models with three different architectures to identify money laundering typology I-d. Our results demonstrated that recurrent neural network models trained with past bank transactions can effectively generalize the identification of money laundering typology I-d in the future transactions. A system automating the identification of money laundering typologies can help the civil police shorten the duration of the investigative process.

II. Methods

A. Data set

This work utilized an anonymized dataset provided by the Civil Police of the State of Pernambuco, containing two sets of files: Financial Intelligence Reports (produced by COAF from communications of gatekeepers) and bank transaction statements of the investigated individuals, provided by banks following proper judicial authorization. The bank transaction statements contain 138,000 anonymized transactions of individuals under investigation for suspected money laundering (see Table I).

The Financial Intelligence Report (RIF), produced by COAF, identifies indications of money laundering practices observed in specific bank accounts. It specifies the time frame and the type of money laundering typology identified in those accounts. Generally, financial investigations are preceded by a Financial Intelligence Report (RIF), while bank transaction statements cover a broader time frame than the RIF. We chose the typology I-d, defined as the fragmentation of deposits or other instruments of cash transfer to disguise the total amount of the transactions, by the Brazilian Central Bank Circular Letter 4.001/2020.

B. Data Preprocessing

All continuous data were preprocessed using the min-max scaler approach, such that each variable is mapped within the interval $[0, 1]$. Categorical data were preprocessed using the one-hot encoding approach such that each variable with k categories were transformed into k binary variables sparsely indicating the mutually exclusive presence of a category. Sequential data was extracted from the bank transactions according to the chosen parameters (see Table II) using the Keras package [13]. At the end, the I-d typology was automatically labeled in the bank statements according to the period indicated by the RIF.

C. Modeling Bank Transaction Sequences using Recurrent Neural Networks

The chosen architectures of recurrent neural networks, following widely adopted practices for sequence modeling [14], were the basic recurrent neural network (SimpleRNN), the recurrent neural network with long short-term memory cells (LSTM) [15], and the recurrent neural network with gated recurrent units (GRU) [16]. Each model was composed of either one recurrent layer with 10 units or two recurrent layers with 10 units each. At the end, each model had a final dense

TABLE I
Data description of the anonymized bank transactions provided by the civil police of the state of Pernambuco.

Variable	Description
Day	Day of the transaction.
Month	Month of the transaction.
Value	The money involved.
Balance Amount	The account balance of the person carrying out.
Transfer type	The types of transfers carried out can include withdrawals, deposits, or credit card transactions.
CNAB	A transaction standardization code classifying groups of similar transactions.
Balance Type	The nature of the balance, which can be either credit or debit.
I-d	The presence or absence of the I-d money laundering typology.

TABLE II
Parameter setting for generation of sequence data from bank transactions.

Parameter	Value
Sequence Length	$20, 30, 40, 50$
Stride	1
Batch size	32

TABLE III
NUMBER OF PARAMETERS OF THE RECURRENT NEURAL NETWORKS, BASED ON THE NUMBER OF LAYERS AND THE ARCHITECTURE

Model	Recurrent Layers	Number of Parameters
Simple RNN	1	571
	2	781
LSTM	1	2,251
	2	3,091
GRU	1	1,721
	2	2,381

layer with a sigmoid activation function with. The total number of parameters of each architecture is shown on Table III.

D. Training and Testing

A total of $138,305$ banking transactions were trained in a supervised manner with a recurrent neural network model using $69,153$ transactions, and subsequently tested with $69,152$ future transactions. One-third of the training transactions were used for validation. During training, we iterated over 30 epochs. An *EarlyStopping* callback was implemented with a patience of 3 epochs and monitoring *validation loss*, so the training process ceased early if no improvement was observed.

E. Statistical Analysis

The metric used in this work was the Receiver Operating Characteristic Area Under the Curve (AUC-ROC) using the package [17]. Confidence intervals were built using 30 bootstrap samples with replacements from the predictions of recurrent neural network models. ROC curves were built from the true positive rates and false positive rates using the package [17].

III. RESULTS

A. Descriptive Statistics

The anonymized data of bank transactions contains a rich diversity of distinct transactions (see Table IV). In total, it includes 188 different accounts with $138,305$ transactions of which almost $10,446$ (8%) are characterized as I-d typology.

TABLE IV
DESCRIPTIVE STATISTICS OF BANK TRANSACTION DATA.

Typology	I-d	None
Accounts	11	177
Individuals/Companies	6	35
Transactions	$10,446$	$127,859$

TABLE V
COMPARATIVE ANALYSIS OF RNN, LSTM, AND GRU IN DETECTING MONEY LAUNDERING TYPOLOGY I-D FROM BANK TRANSACTIONS. BOOTSTRAP 95% CONFIDENCE INTERVALS (CI) FOR THE AUC-ROC.

Model	Recurrent Layers	Sequence Length	95% CI AUC-ROC
Simple RNN	1	20	0.64 (0.63, 0.65)
		30	0.54 (0.54, 0.55)
		40	0.54 (0.53, 0.55)
		50	0.51 (0.49, 0.52)
	2	20	0.64 (0.63, 0.65)
		30	0.64 (0.63, 0.65)
		40	0.56 (0.55, 0.57)
		50	0.55 (0.54, 0.56)
LSTM	1	20	0.72 (0.71, 0.72)
		30	0.69 (0.68, 0.70)
		40	0.75 (0.74, 0.76)
		50	0.60 (0.59, 0.61)
	2	20	0.68 (0.67, 0.69)
		30	0.70 (0.70, 0.71)
		40	0.74 (0.73, 0.75)
		50	0.66 (0.65, 0.67)
GRU	1	20	0.72 (0.71, 0.72)
		30	0.75 (0.74, 0.76)
		40	0.73 (0.72, 0.73)
		50	0.72 (0.71, 0.73)
	2	20	0.74 (0.73, 0.74)
		30	0.70 (0.70, 0.71)
		40	0.71 (0.70, 0.71)
		50	0.73 (0.72, 0.74)

B. Handling Increased Sequence Lengths in Transactions

Our results show that recurrent network models with gated recurrent units (GRUs), across all sequence lengths from 20 to 50, consistently exhibited an AUC-ROC higher than 0.70 $(0.70, 0.71)$ (see Table V). In contrast, an increase in sequence length jeopardizes the performance of simple recurrent neural networks (SimpleRNN) 0.64 $(0.63, 0.65)$ to 0.51 $(0.49, 0.52)$. We find that LSTM is superior to SimpleRNN, but it appears to present some instability over greater sequence lengths. Our results reveal that achieving performance greater than a random classifier at lower regions of false positive rates (FPR) was challenging for all recurrent neural networks.

In general, we find that the models perform similarly most of the time, as shown by the ROC curves in Fig. 1. However, there are some intervals where the LSTM demonstrates superior performance.

The results indicate that increasing the sequence length or the number of neural layers does not necessarily lead to maximum performance. Consequently, the importance of optimizing hyperparameters becomes evident in the quest to develop better neural architectures, potentially leading to

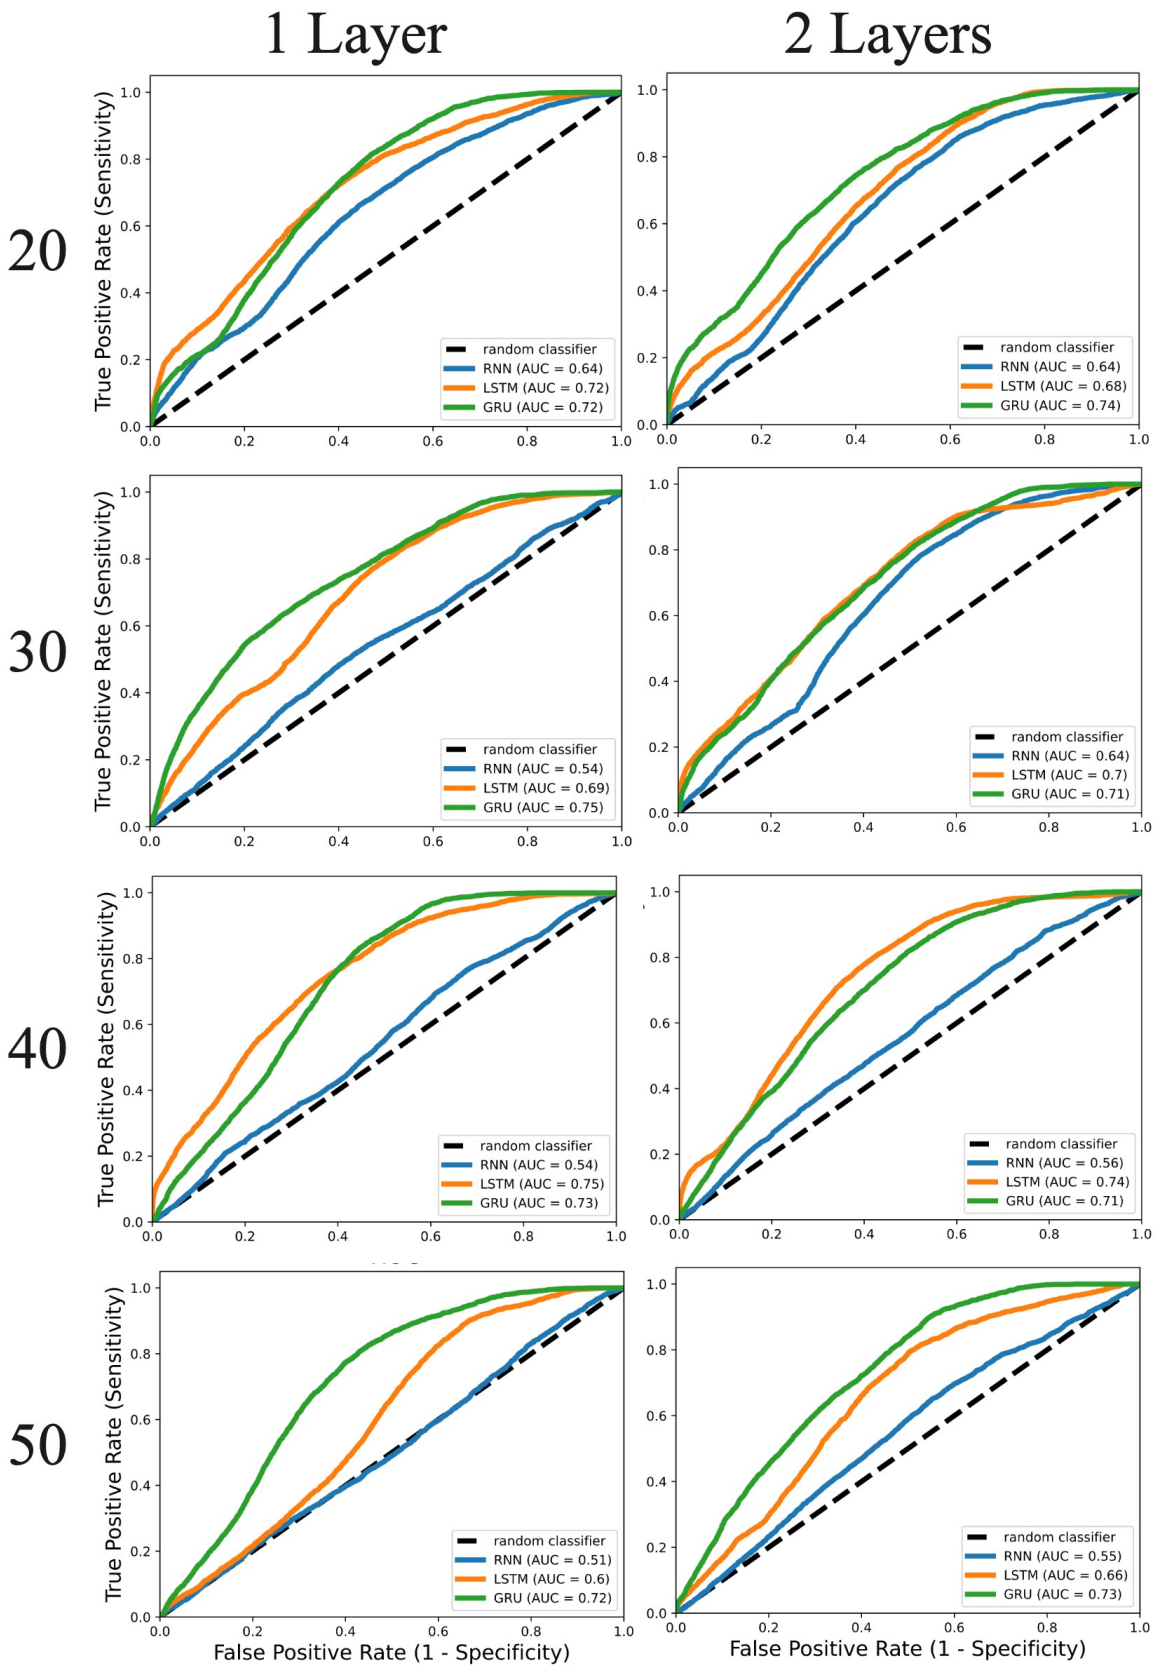

Fig. 1. Comparative analysis of RNN, LSTM, and GRU in detecting money laundering typology I-d from bank transactions through receiver operating characteristic (ROC) curves. Each column depicts results of the recurrent neural models with one (left) and two (rights) recurrent layers, and the rows correspond to sequence lengths of size 20, 30, 40, and 50.

higher-quality models.

IV. Conclusion

Detecting money laundering is crucial for protecting society and the economy, helping prevent illicit financial activities. Traditional detection methods estimate the likelihood of individuals being involved in money laundering but have struggled to pinpoint specific money laundering typologies within individual bank transactions. In this work, we developed different recurrent neural models to identify typology I-d, showing promising results on real-world data.

Our findings reveal a decrease in the performance of the simple recurrent neural network as sequence lengths increase. We believe this is related to the vanishing gradient problem. Although LSTM and GRU models appear to handle greater sequence lengths, the sequence length seems to be a significant hyperparameter that needs to be tuned. In general, deeper models can have superior data processing capability than their shallow counterparts, given that each additional layer provides the ability to increasingly express complex patterns. Nevertheless, more complex models require more data. Though the additional recurrent layer seemed to not increase the performance of our recurrent models, having more data can enable us to train deeper and superior models.

The automatic data labeling employed in our work may limit the algorithm's ability to accurately identify I-d transactions. Firstly, labeling the entire period indicated by the RIF as I-d may result in the incorrect labeling of transactions. Secondly, labeling data solely based on the information provided by the RIF may not capture all actual I-d transactions, as agents may fail to notify COAF of signs of I-d money laundering.

In future works, we will conduct a comprehensive neural architecture search, exploring a variety of hyperparameters, including the number of recurrent layers and the number of recurrent units at each layer. Other architectures will be considering, from bidirectional recurrent neural networks to transformers. Also, we will expand this work for the identification of money laundering typologies other than I-d typology.

Currently, a typical money laundering investigation demands approximately 1-2 months. Automating money laundering detection with machine learning-based approaches can reduce this time to 1-2 days. Accelerating fraud detection prevents significant financial losses by stopping unauthorized transactions before they are completed. Additionally, it can serve as a decision support system, helping analysts avoid investigating false positives using slower detection methods. Finally, faster fraud detection helps ensure compliance with national regulations, reducing the potential loss of human resources dedicated to preventing the misappropriation of public funds.

Acknowledgements

Diego Pinheiro and Isadora Xavier would like to thank the *Programa Institucional de Bolsas de Iniciação Científica* (FACEPE, Brazil) for financial support under grant number 1/2023.

References

[1] Z. Gao and M. Ye, "A framework for data mining-based anti-money laundering research," *Journal of Money Laundering Control*, vol. 10, no. 2, pp. 170–179, 2007.

[2] R. I. T. Jensen and A. Iosifidis, "Fighting Money Laundering With Statistics and Machine Learning," *IEEE Access*, vol. 11, pp. 8889–8903, 2023.

[3] Yufeng Kou, Chang-Tien Lu, S. Sirwongwattana, and Yo-Ping Huang, "Survey of fraud detection techniques," in *IEEE International Conference on Networking, Sensing and Control, 2004*, vol. 2, pp. 749–754, IEEE, 2004.

[4] S. Wang, P. Wang, B. Wu, Y. Zhu, W. Luo, and Y. Pan, "Structural entropy minimization combining graph representation for money laundering dentification," *International Journal of Machine Learning and Cybernetics*, Apr. 2024.

[5] A. A. S. Alsuwailem and A. K. J. Saudagar, "Anti-money laundering systems: a systematic literature review," *Journal of Money Laundering Control*, vol. 23, pp. 833–848, May 2020.

[6] T. Mikolov, A. Joulin, S. Chopra, M. Mathieu, and M. Ranzato, "Learning longer memory in recurrent neural networks," *arXiv preprint arXiv:1412.7753*, 2014.

[7] Z. Chai, Y. Yang, J. Dan, S. Tian, C. Meng, W. Wang, and Y. Sun, "Towards Learning to Discover Money Laundering Sub-network in Massive Transaction Network," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, pp. 14153–14160, June 2023.

[8] H. S. Assumpcao, F. Souza, and L. L. Campos, "Delator: Detecção automática de indícios de lavagem de dinheiro por redes neurais em grafos de transacões," *Brazilian Workshop on Social Network Analysis and Mining (BRASNAM)*, 2022.

[9] W. Hilal, S. A. Gadsden, and J. Yawney, "Financial fraud: a review of anomaly detection techniques and recent advances," *Expert systems With applications*, vol. 193, p. 116429, 2022.

[10] R. Van Belle, S. Mitrović, and J. De Weerdt, "Representation learning in graphs for credit card fraud detection," in *Mining Data for Financial Applications: 4th ECML PKDD Workshop, MIDAS 2019, Würzburg, Germany, September 16, 2019, Revised Selected Papers 4*, pp. 32–46, Springer, 2020.

[11] R. Van Belle, C. Van Damme, H. Tytgat, and J. De Weerdt, "Inductive graph representation learning for fraud detection," *Expert Systems with Applications*, vol. 193, p. 116463, 2022.

[12] S. Reddy, P. Poduval, A. V. S. Chauhan, M. Singh, S. Verma, K. Singh, and T. Bhowmik, "Tegraf: temporal and graph based fraudulent transaction detection framework," in *Proceedings of the Second ACM International Conference on AI in Finance*, pp. 1–8, 2021.

[13] F. Chollet *et al.*, "Keras." https://keras.io, 2015.

[14] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," *arXiv preprint arXiv:1412.3555*, 2014.

[15] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[16] K. Cho, B. Van Merriënboer, D. Bahdanau, and Y. Bengio, "On the properties of neural machine translation: Encoder-decoder approaches," *arXiv preprint arXiv:1409.1259*, 2014.

[17] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.