UNIVERSIDAD SAN FRANCISCO DE QUITO USFQ

Colegio de Posgrados

A Two-Stage Unsupervised Learning Approach for Anomaly Detection in Transactions of Cooperative Financial Institutions

Proyecto de Titulación

Christian Sebastian Hernández Mosquera

Felipe Grijalva, Ph.D.

Director de Trabajo de Titulación

Trabajo de titulación de posgrado presentado como requisito para la obtención del título de Magíster en Inteligencia Artificial

Quito, 02 de diciembre de 2024

UNIVERSIDAD SAN FRANCISCO DE QUITO USFQ
COLEGIO DE POSGRADOS

HOJA DE APROBACIÓN DE TRABAJO DE TITULACIÓN

A Two-Stage Unsupervised Learning Approach for Anomaly Detection in Transactions of Cooperative Financial Institutions

Christian Sebastian Hernández Mosquera

Nombre del Director del Programa:	Felipe Grijalva
Título académico:	Ph.D. en Ingeniería Eléctrica
Director del programa de:	Inteligencia Artificial
Nombre del Decano del colegio Académico:	Eduardo Alba
Título académico:	Doctor en Ciencias Matemáticas
Decano del Colegio:	Ciencias e Ingenierías
Nombre del Decano del Colegio de Posgrados:	Dario Niebieskikwiat
Título académico:	Doctor en Física

Quito, diciembre 2024

© DERECHOS DE AUTOR

Nombre del estudiante:	Christian Sebastian Hernández Mosquera
Código de estudiante:	00338710
C.I.:	1720408499
Lugar y fecha:	Quito, 2 de diciembre de 2024

ACLARACIÓN PARA PUBLICACIÓN

Nota: El presente trabajo, en su totalidad o cualquiera de sus partes, no debe ser considerado como una publicación, incluso a pesar de estar disponible sin restricciones a través de un repositorio institucional. Esta declaración se alinea con las prácticas y recomendaciones presentadas por el Committee on Publication Ethics COPE descritas por Barbour et al. (2017) Discussion document on best practice for issues around theses publishing, disponible en http://bit.ly/COPETheses.

UNPUBLISHED DOCUMENT

Note: The following graduation project is available through Universidad San Francisco de Quito USFQ institutional repository. Nonetheless, this project – in whole or in part – should not be considered a publication. This statement follows the recommendations presented by the Committee on Publication Ethics COPE described by Barbour et al. (2017) Discussion document on best practice for issues around theses publishing available on http://bit.ly/COPETheses.

DEDICATORIA

A mis padres, por su amor incondicional, gracias por ser mi mayor ejemplo de perseverancia y dedicación.

A mi hermana, por su apoyo constante y por compartir alegrías y retos conmigo.

Y a Andrés, el amor de mi vida, por tu paciencia infinita, tu ánimo constante y por caminar conmigo en este viaje.

Este logro no habría sido posible sin ustedes. Con todo mi amor, esta tesis está dedicada a ustedes.

AGRADECIMIENTOS

Quiero expresar mi más profundo agradecimiento a **Coonecta Red Transaccional Cooperativa** por su invaluable apoyo al proveer los datos necesarios para este trabajo. Su colaboración fue esencial para llevar a cabo esta investigación.

Finalmente, extiendo mi gratitud a mis mentores y profesores por su guía académica y a mis compañeros de la maestría por sus aportes y discusiones enriquecedoras.

RESUMEN

La detección de anomalías en transacciones financieras es un desafío crítico para mitigar los riesgos asociados con actividades fraudulentas. Este estudio propone un marco robusto adaptado al sector financiero cooperativo, utilizando técnicas de aprendizaje no supervisado para identificar comportamientos anómalos en los datos transaccionales. Se empleó Hierarchical Clustering para segmentar a los usuarios en clústeres de comportamiento distintivos, revelando patrones indicativos de actividad normal y sospechosa. Posteriormente, se implementó un modelo de One-Class Support Vector Machine (SVM) para la detección de anomalías, logrando puntuaciones de Area Under the Curve (AUC) de hasta 0.90 en clústeres con patrones bien definidos. La visualización mediante t-SNE destaca un solapamiento mínimo entre transacciones normales y anómalas en estos casos, demostrando la eficacia del modelo. Sin embargo, los clústeres con mayor variabilidad mostraron un desempeño reducido, reflejando los desafíos que plantean los datos transaccionales diversos. Esta investigación enfatiza la importancia de estrategias específicas para la detección de anomalías y proporciona perspectivas para mejorar la prevención del fraude en sistemas financieros, particularmente en sectores subrepresentados como el cooperativo. El trabajo futuro explorará cómo mejorar la escalabilidad y la interpretabilidad para fortalecer la aplicabilidad en el mundo real.

Palabras clave: Anomaly Detection, Financial Transactions, Fraud Detection, Cooperative Financial Sector, Unsupervised Learning, Hierarchical Clustering, One-Class Support Vector Machine (SVM), t-SNE Visualization, Fraud Prevention.

ABSTRACT

Anomaly detection in financial transactions is a critical challenge in mitigating risks associated with fraudulent activities. This study proposes a robust framework tailored to the cooperative financial sector, utilizing unsupervised learning techniques to identify anomalous behavior within transaction data. Hierarchical clustering, was employed to segment users into distinct behavioral clusters, revealing patterns indicative of normal and suspicious activity. A One-Class Support Vector Machine (SVM) was then implemented for anomaly detection, achieving Area Under the Curve (AUC) scores up to 0.90 in clusters with well-defined patterns. The t-SNE visualization highlights minimal overlap between normal and anomalous transactions in these cases, demonstrating the model's efficacy. However, clusters with greater variability exhibited decreased performance, reflecting the challenges posed by diverse transactional data. This research emphasizes the importance of domain-specific anomaly detection strategies and provides insights for enhancing fraud prevention in financial systems, particularly in underrepresented sectors like cooperative sector. Future work will explore improving scalability and interpretability to strengthen real-world applicability.

Key words: Anomaly Detection, Financial Transactions, Fraud Detection, Cooperative Financial Sector, Unsupervised Learning, Hierarchical Clustering, One-Class Support Vector Machine (SVM), t-SNE Visualization, Fraud Prevention.

TABLA DE CONTENIDO

ÍNDICE DE TABLAS

ÍNDICE DE FIGURAS

A Two-Stage Unsupervised Learning Approach for Anomaly Detection in Transactions of Cooperative Financial Institutions

Christian Hernández Felipe Grijalva

***Abstract*—Anomaly detection in financial transactions is a critical challenge in mitigating risks associated with fraudulent activities. This study proposes a robust framework tailored to the cooperative financial sector, utilizing unsupervised learning techniques to identify anomalous behavior within transaction data. Hierarchical clustering, was employed to segment users into distinct behavioral clusters, revealing patterns indicative of normal and suspicious activity. A One-Class Support Vector Machine (SVM) was then implemented for anomaly detection, achieving Area Under the Curve (AUC) scores up to 0.90 in clusters with well-defined patterns. The t-SNE visualization highlights minimal overlap between normal and anomalous transactions in these cases, demonstrating the model's efficacy. However, clusters with greater variability exhibited decreased performance, reflecting the challenges posed by diverse transactional data. This research emphasizes the importance of domain-specific anomaly detection strategies and provides insights for enhancing fraud prevention in financial systems, particularly in underrepresented sectors like cooperative sector. Future work will explore improving scalability and interpretability to strengthen real-world applicability.**

***Index Terms*—Anomaly Detection, Financial Transactions, Fraud Detection, Cooperative Financial Sector, Unsupervised Learning, Hierarchical Clustering, One-Class Support Vector Machine (SVM), t-SNE Visualization, Fraud Prevention.**

I. Introduction

IN an increasingly interconnected world, digital financial transactions have become an essential pillar of the global economy. Since the early 2000s, with the democratization of digital payment methods through debit and credit cards, user behavior has evolved steadily. However, the COVID-19 pandemic brought a dramatic shift in consumer behavior, rapidly promoting the use of online transactions and the adoption of digital technologies for payments and transfers. While these transactions are significantly more convenient and, in many cases, safer, they are not exempt from the risk of increasingly sophisticated fraud attempts.

Fraud in digital transactions encompasses everything from the physical theft of cards to complex social engineering strategies designed to illicitly obtain debit and credit card information, resulting in unauthorized charges made by someone other than the cardholder. To mitigate these risks, various advanced fraud detection and prevention mechanisms have been developed, including rule-based tools that allow financial institutions to manage risk effectively. While these solutions have proven effective in the traditional financial system, they do not always adapt well to the unique characteristics of other segments, such as the popular and solidarity-based financial system, which includes savings and credit cooperatives and microfinance banks.

The cooperative financial system in Ecuador, composed of cooperatives and microfinance banks, faces unique challenges in fraud risk management. Currently, there are no fraud management tools specifically built with data from the cooperative sector, creating a significant gap. Although conventional financial system tools can provide a foundation, they do not fully match the reality and needs of these institutions.

This thesis project proposes a tailored solution based on transactional information from cooperatives within the COONECTA RED TRANSACCIONAL network. The implementation of this approach will enable the addition of a customized control layer that complements existing safeguards established by card franchises, thereby adding an extra layer of fraud protection. This solution aims to minimize friction in the end-user experience while strengthening security.

The main objective of this project is to classify the customers of the COONECTA network and build anomaly detectors for each customer class

using a two-phase unsupervised learning approach. This method, based on user segmentation and anomaly detection, has the potential to serve as an innovative tool for fraud risk management in a more granular and precise manner, optimizing protection and efficiency.

II. Prior Works

Fraud detection in financial transactions has been an active field of research due to its importance in preserving the integrity of economic systems and reducing significant monetary losses. The methodologies used range from traditional approaches to advanced machine learning techniques, each with its advantages and limitations.

Several studies have addressed fraud mitigation in financial transactions through anomaly detection using unsupervised learning methods, which allow unusual transactions to be identified without explicit fraud labels. For example, the work in [1] proposes a relevant approach using hierarchical clustering with K-means and GMM to isolate anomalous transactions in a method called Clustering Tree. This method organizes suspicious transactions in specific nodes as the feature segmentation deepens. This technique has proven effective, particularly for imbalanced datasets, as is often the case with fraud data.

On the other hand, in [2], ensemble learning strategies are implemented to optimize performance in both overlapping and imbalanced data. This study proposes a hybrid model that combines Elliptic Envelope, Isolation Forest, Local Outlier Factor, and One-Class SVM, consolidating the results in an ensemble classifier, which significantly improves fraud detection accuracy.

Other studies have used a combination of supervised and unsupervised techniques to address the same problem. In [3], a hybrid approach is presented that employs unsupervised clustering methods to generate outlier features, which are then transferred to a supervised model, such as Random Forest. This approach has proven to improve accuracy by detecting both known and new fraud patterns.

Neural networks and attention-based methods have also been explored to capture transaction-specific features. In [4], an unsupervised model called UAAD-FDNet (Unsupervised Attentional Anomaly Detection Network for Fraud Detection) is used, which combines autoencoders and attention networks to reconstruct normal transactions and detect anomalies more accurately in high-dimensional data.

the subspace-based learning model in [5] focuses on one-class classification through dimensionality reduction, optimizing pattern analysis in large volumes of transactional data. This approach combines algorithms like One-Class SVM with dimensionality reduction techniques, facilitating the detection of anomalies in high-dimensional data with greater efficiency.

The reviewed studies indicate that unsupervised methods and ensemble learning approaches have shown better performance compared to traditional methods. Although these approaches tend to reduce false positives (FPR), their reduction of false negatives (FNR) is limited, which is critical in fraud detection, as false negatives (undetected fraudulent transactions) have a greater impact on customer experience and financial security. Therefore, the correct approach would be to find a balance, prioritizing FNR minimization [1], [2].

Additionally, [4] proved to be more effective than traditional models such as Isolation Forest and can handle high-dimensional data; however, it requires a larger volume of data for training and is more complex due to the use of GANs.

A common limitation of these studies is the dependence on synthetic data or data augmentation techniques to mitigate class imbalance. This limitation arises because well-regulated financial systems tend to experience few fraud cases relative to the total volume of transactions. Furthermore, supervised learning approaches are difficult to implement in this context due to the lack of labeled fraud information [1], [2].

Some studies suggest the inclusion of reinforcement learning to improve model adaptability and integration of various data sources, which could enhance robustness and generalization capability [2], [4].

Finally, the work of Hejazi and Singh [6] presents an approach similar to this study using One-Class SVM (OC-SVM) for anomaly detection in financial

transactions. However, this approach is based on a synthetic dataset, limiting its applicability in a real-world setting.

This study implements a rarely used approach among the reviewed studies and is conducted under favorable data conditions. It uses a real dataset provided by Ecuador's private financial sector, specifically the cooperative financial sector, including transactions from at least 50 financial entities in this sector.

Additionally, two instances of unsupervised learning are employed: an initial hierarchical clustering strategy for customer segmentation and an OC-SVM (One-Class Support Vector Machine) model. This approach enables adaptive construction of anomaly detectors specific to each customer type (cluster). This combination facilitates a more granular detection adjusted to the reality of the customer.

III. Materials and Methods

A. Dataset Description

The dataset used in this project was provided by COONECTA RED TRANSACCIONAL, a company specialized in giving financial services for cooperatives and microcredit banks in Ecuador. This network encompasses approximately 70 financial institutions operating under various services, including ATM networks, international remittance reception, and purchase networks powered by one of the world's most recognized card franchises.

For this project, the dataset focuses exclusively on transactions carried out with cards issued by financial entities associated with COONECTA RED TRANSACCIONAL. These transactions cover a one-year period and are classified as "card-not-present" transactions. This means that the card was not physically used at a Point of Sale (POS) terminal. Instead, users either manually entered their card details on online shopping platforms or utilized the "credential-on-file" method. In this approach, the e-commerce platform or application stores the card information as a token, enabling seamless future transactions or recurring charges, such as those for streaming services or subscription-based platforms.

Each record in the dataset includes multiple features critical for analysis, such as:

Unique customer identifiers. Transaction details: date, time, amount, and the region where the purchase was made. Merchant Category Codes (MCC), which specify the type of merchant where the transaction occurred. Fraudulent transactions in this dataset were identified using the risk management system provided by the card franchise. This system operates with a set of customizable rules configured by issuing entities, acting as filters to authorize or decline transactions. In many cases, the system can reject a suspicious transaction even before it is processed by the bank's core authorization system. Additionally, the system uses a model that calculates risk scores based on previously detected fraud patterns and predefined configurations. It is important to note that the features of fraudulent transaction records are identical to those of legitimate transactions, ensuring uniformity in analysis.

The dataset comprises a total of transactions distributed as follows:

- **97% non-fraudulent transactions:** These represent legitimate operations carried out by cardholders with no indications of irregularity.
- **3% fraudulent transactions:** These include transactions reported by cardholders as unauthorized purchases. Such reports were handled through COONECTA's customer support center and recorded in the risk management system. Additionally, this category includes transactions that were automatically declined by the fraud detection system or flagged based on recommendations tied to:
  - The location of the purchase origin.
  - The frequency of transactions within a short period.
  - The risk level associated with the merchant type (MCC).

This imbalanced class distribution (97% vs. 3%) highlights the inherent challenge of class imbalance in such datasets. This issue is particularly relevant for designing and evaluating anomaly detection models, as any solution must minimize false positives while maintaining the ability to accurately detect actual fraud cases.

B. Methodology

The project's pipeline follows a two-stage structure based on unsupervised learning. As illustrated in the block diagram Fig 1, this workflow ensures that anomalies are detected accurately within each group of existing customers. The design leverages transactional features and behavioral patterns to model significant differences between normal and anomalous transactions.

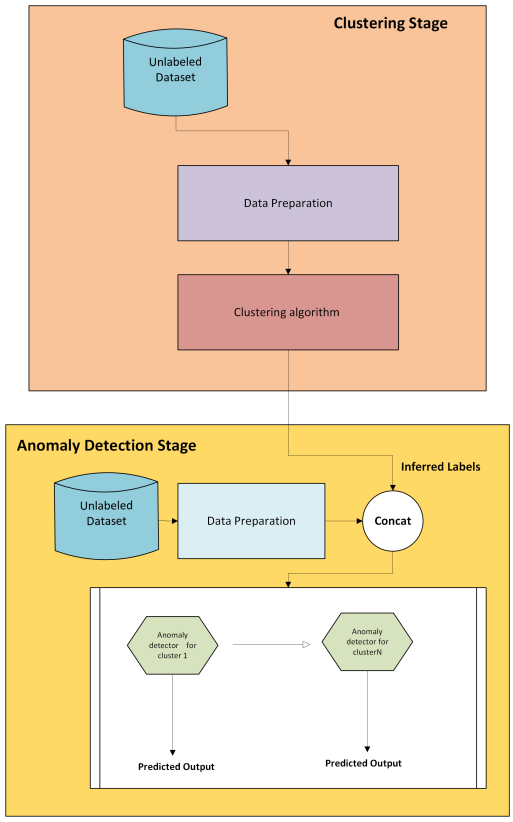

Figure 1. Project Pipeline

1) Clustering Stage: The problem of anomaly detection in financial transactions is initially approached as a clustering problem, owing to the multidimensional and contextual nature of transactional data. Unlike typical time-series problems, which focus on temporal dependencies, this case involves identifying natural groups in transactional behavior to establish normal patterns and detect significant deviations.

In this stage, transactional data was used to determine the number of customer types (clusters). The underlying assumption is that financial institutions share cardholders with similar behaviors. The 70 participating financial entities vary in size and complexity, ranging from small cooperatives to larger institutions. By pooling data from all institutions, smaller cooperatives benefit from shared insights and patterns, despite generating lower transaction volumes.

The goal of this stage is to assign each card ID to a specific cluster, regardless of its associated financial institution. Transaction data was preprocessed by grouping transactions by quarter and segmenting them by merchant type (MCC). This approach allowed the generation of representative features such as total transaction volume, usage frequency, average value, and the proportion of successful transactions, forming a multidimensional space for analysis.

Various algorithms were evaluated to determine the most suitable clustering methodology. Initially, Gaussian Mixture Models (GMM) were applied. This probabilistic approach assumes that data originates from a mixture of Gaussian distributions, allowing for the identification of clusters with soft boundaries. However, in this project, GMM failed to produce interpretable clusters aligned with the analysis objectives.

Next, Hierarchical Clustering was implemented, which constructs a hierarchy of groups based on similarity or distance measures. This method provided a clearer representation of the relationships between observations. The final choice of Hierarchical Clustering was based on its ability to logically represent differences between clusters, maximizing intracluster cohesion and intercluster separation using the Silhouette index.

C. Anomaly Detection Stage

Using the clusters identified in the previous stage, an anomaly detector was built with One-Class Support Vector Machine (One-Class SVM). This model is ideal for identifying patterns deviating from normal behavior, particularly in scenarios where anomalous data is scarce or unlabeled. One-Class SVM models a decision boundary using only normal data, identifying observations significantly deviating from this standard.

The model configuration included aggregated features such as transaction frequency, average amounts. Using the RBF kernel, One-Class SVM effectively modeled nonlinear decision boundaries in high-dimensional spaces.

1) Experimental Configuration: The One-Class SVM model was configured and optimized through a grid search process, evaluating combinations of hyperparameters to determine the optimal configuration. Table I shows the ranges of values tested for each hyperparameter:

Table I
Hyperparameter Configuration for One-Class SVM

Hyperparameter	Values Tested
ν (nu)	0.000005, 0.00005, 0.0005, 0.005, 0.01
γ (gamma)	0.000001, 0.00001, 0.0001, 0.001, 0.01
Kernel	Radial Basis Function (RBF)

The selected kernel, RBF, transforms the data into a higher-dimensional space, enabling greater flexibility in separating anomalies. A total of 25 configurations were evaluated for each cluster.

The model's performance was assessed using the following metrics:

The performance of the One-Class SVM model was assessed using the following metrics:

- **F1-Score:** Combines precision and recall into a single value, providing a balanced measure of the model's ability to handle false positives and false negatives.
- **Precision:** Measures the proportion of transactions classified as anomalies that were genuinely fraudulent.
- **Recall (Sensitivity):** Indicates the proportion of actual fraudulent transactions correctly identified by the model.
- **Area Under the ROC Curve (AUC-ROC):** Evaluates the model's overall capability to distinguish between normal and anomalous transactions.

These metrics provided a comprehensive evaluation of the model's performance, enabling the identification of the most effective configurations.

IV. Results

This section presents the key findings derived from the implementation of the proposed model for anomaly detection in financial transactions. The results are organized into two main stages, aligned with the pipeline previously described: customer clustering to identify homogeneous groups and anomaly detection within these groups.

In the clustering stage, transactional patterns among customers from various financial entities were analyzed to group them into representative profiles based on features such as usage frequency and average transaction amounts. These clusters provided a robust foundation for customizing anomaly detectors according to the specific patterns of each group.

In the second stage, a One-Class Support Vector Machine (SVM) model was implemented for each cluster, optimized to maximize anomaly detection while minimizing false positives. The model's performance was evaluated using key metrics such as F1-Score, precision, recall, and the area under the ROC curve (AUC-ROC), ensuring reliable and reproducible analysis.

The findings confirm the effectiveness of the proposed approach, highlighting its applicability in improving fraud management within the cooperative financial sector. Moreover, these results emphasize the importance of adaptive strategies in anomaly detection systems for sectors with heterogeneous data and varying usage frequencies.

A. Clustering Stage

To determine the optimal number of clusters, the Silhouette index was used. This metric, widely employed in clustering analysis, measures the quality of the clusters formed. The Silhouette index evaluates two main aspects: Intracluster cohesion wich is how similar the observations within the same cluster are and intercluster separation is how different the observations in one cluster are from those in other clusters.

It is important to note that this metric was applied using hierarchical clustering. The index was calculated and compared across varying numbers of target clusters to identify the most suitable grouping.

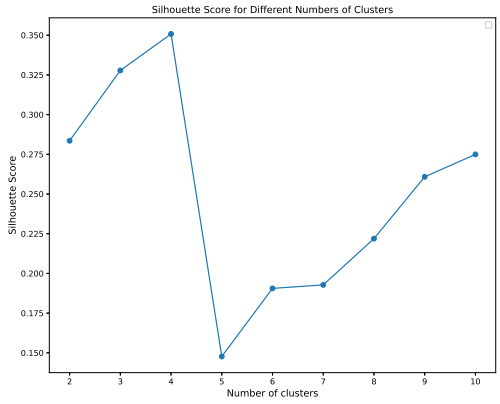

Figure 2. Silhouette Score

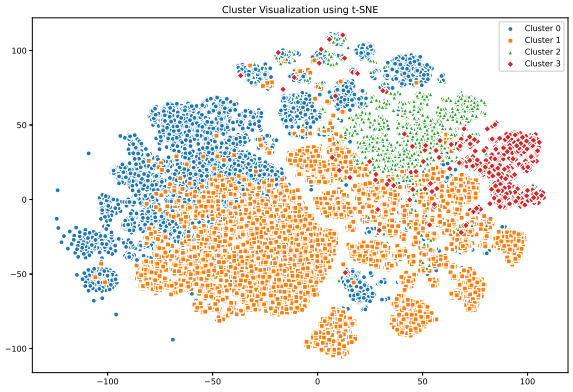

Figure 3. Cluster Visualization using t-SNE

In this case, the Silhouette score achieved its maximum value with 4 clusters, as shown in Fig 2. This result suggests that dividing the transactional data into four groups provides the most appropriate segmentation, maximizing homogeneity within each cluster and ensuring clear differentiation between clusters.

Once the optimal number of clusters was established, hierarchical clustering was reapplied, generating a segmentation that reflects specific transactional behaviors. To visualize the separation between customer types or clusters, t-SNE (t-Distributed Stochastic Neighbor Embedding) was employed. This dimensionality reduction technique projects high-dimensional data into two- or three-dimensional spaces, enabling a graphical representation of the clusters.

t-SNE is a non-linear algorithm designed to preserve the local structure of the data, meaning that points close in the original space remain close in the projection. This is particularly useful in this application, given the large number of features used for class division. It is important to note that t-SNE serves as a visualization guide and should not be used as a standalone clustering technique. Instead, it helps validate clusters visually and identify specific patterns.

In Fig 3, the transactional data are grouped into four clearly distinct clusters, represented by different colors and shapes. This visualization validates the separation achieved through hierarchical clustering and offers an intuitive understanding of the transactional patterns present in each group.

Cluster 0: The second-largest group, representing around 30% of total transactions, is characterized by frequent, low-value purchases in retail, entertainment and recreation, and agriculture and food.

Cluster 1: The largest group, accounting for approximately 45% of total transactions. Customers in this cluster focus on planned purchases, primarily in retail, entertainment and recreation.

Cluster 2: A smaller group, contributing 15% of total transactions, Spending in this cluster is more diversified, spanning retail, transportation/logistics, and entertainment/recreation.

Cluster 3: The smallest group, comprising less than 10% of total transactions, Transactions in this cluster are concentrated in retail.

The segmentation is critical for the next stage of the pipeline: anomaly detection. Each cluster represents a unique context of transactional behavior, enabling models to be trained and tailored to the specific characteristics of each group. This ensures greater accuracy and relevance in detecting deviations, optimizing the ability to identify potential fraud.

B. Anomaly Detector

The primary goal is to identify transactions that deviate significantly from normal behavior. Given the high dimensionality and complexity of financial transactions, a rule-based system would fail to adequately model such behaviors and might struggle to detect subtle or novel anomalies. As outlined in the methodology section, this project

implements an anomaly detection model based on One-Class SVM.

This algorithm is particularly well-suited for this application, as it is designed to learn a boundary around the normal data points in the feature space. This allows it to flag transactions that fall outside this boundary as potential anomalies, even in the absence of a comprehensive labeled dataset for fraudulent transactions.

The objective of this phase is to create a robust model capable of minimizing both false negatives (fraudulent transactions incorrectly classified as normal) and false positives (normal transactions misclassified as anomalies). False negatives are the most critical type of error, as they directly impact users and their finances, allowing fraudulent transactions to go unnoticed due to a model that is not strict enough. On the other hand, false positives, while important, primarily affect the user experience. In real-world applications, this second type of error can be mitigated by introducing measures such as challenges or OTPs (one-time passwords), allowing users to authenticate and proceed with their transactions without excessive friction.

The results of the anomaly detection phase are presented for each cluster using two key visualizations: a t-SNE plot and a ROC curve. The t-SNE plots provide a graphical representation of the separation between normal and anomalous transactions, while the ROC curves evaluate the model's performance in distinguishing between these categories.

For Cluster 0 in Fig 4, the model performed well, with a t-SNE plot showing clear separation between normal and anomalous transactions and a high AUC of 0.89. This indicates strong detection capabilities, likely due to the homogeneity of the cluster's transactional patterns. However, a few anomalies mixed within normal transactions suggest minor room for improvement.

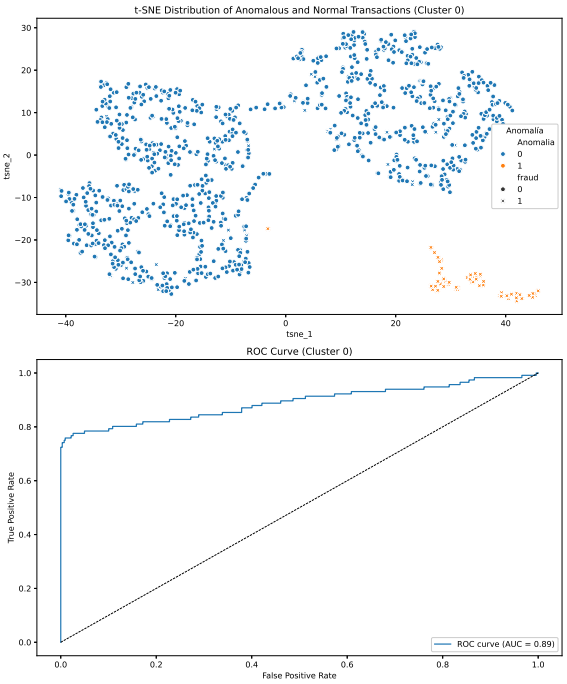

Figure 4. Tsne Transaction distribution and ROC Curve for cluster 0

In Cluster 1 in Fig 5, the t-SNE plot still shows separation, but with more overlap compared to Cluster 0. The model achieves a slightly lower AUC of 0.86, reflecting increased variability or complexity in this group. While still effective, the model's performance is marginally less consistent than in Cluster 0 or Cluster 2.

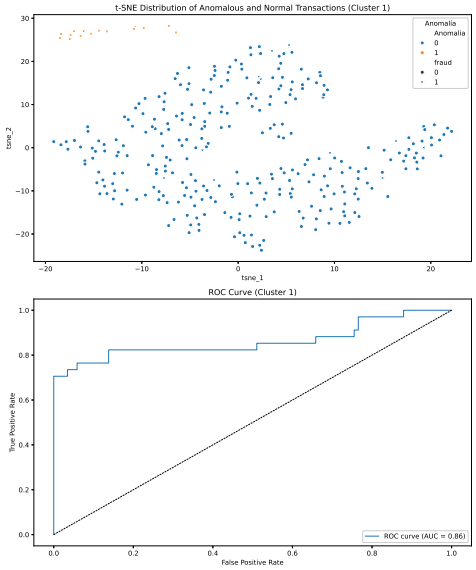

Figure 5. Tsne Transaction distribution and ROC Curve for cluster 1

Cluster 2 in 6 achieved the best overall results, with an AUC of 0.90 and a t-SNE plot that shows

distinct and minimal overlap between normal and anomalous transactions. This indicates that the model is particularly effective for clusters with well-defined patterns, outperforming Clusters 0 and 1 in anomaly detection.

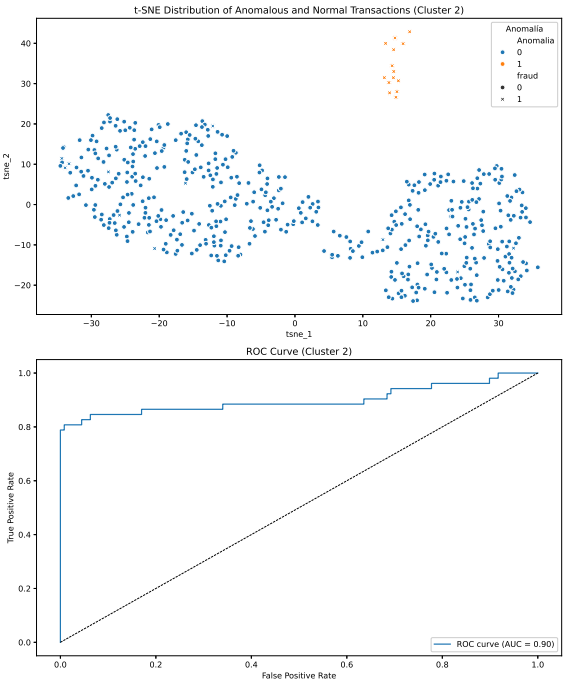

Figure 6. Tsne Transaction distribution and ROC Curve for cluster 2

Finally, Cluster 3 in 7 exhibited the weakest performance, with an AUC of 0.64 and significant overlap between normal and anomalous transactions in the t-SNE plot. This cluster's small size and higher transactional variability likely contributed to the reduced accuracy.

Results of the anomaly detection phase reveal both the strengths and limitations of the proposed methodology, shaped significantly by the inherent characteristics of the dataset and the processes used for labeling and data management. While the model demonstrated strong performance in clusters with homogeneous and well-defined patterns, the variability in results highlights the challenges posed by the data's origin and the potential sources of error in the labeling and fraud management processes.

The dataset used in this study primarily relied on customer service operators to manually label transactions as fraudulent or normal. While this approach is common in financial fraud detection, it introduces the potential for human error in

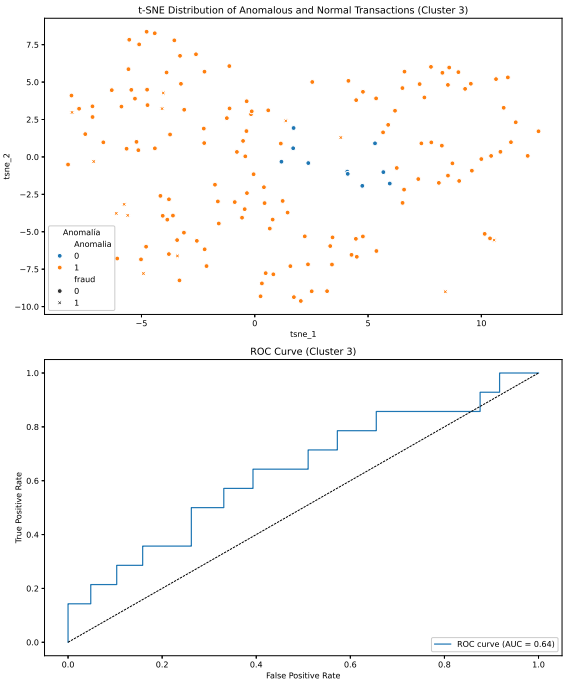

Figure 7. Tsne Transaction distribution and ROC Curve for cluster 3

the labeling process. Operators may misclassify transactions due to a lack of detailed information, inconsistencies in interpretation, or time constraints. This creates noise in the data, which can affect the model's ability to learn accurate patterns and could partly explain the overlap between normal and anomalous transactions observed in some clusters, particularly Cluster 3.

Additionally, the dataset included transactions flagged by automated fraud management tools. Some of these tools improperly used reserved fraud response codes for rule-based decisions unrelated to actual fraudulent behavior. This misclassification introduced further inaccuracies in the dataset, blurring the line between true anomalies and legitimate transactions. As a result, certain financial entities whose data contained significant irregularities in their fraud management practices were excluded from the study to preserve the integrity of the analysis. However, this necessary exclusion reduced the overall size of the dataset, particularly for smaller clusters like Cluster 3, further limiting the model's ability to generalize effectively for these groups.

The model performed well in clusters with a higher volume of transactions and more consistent patterns, such as Clusters 0 and 2, where clear behav-

ioral trends allowed the model to establish robust boundaries for normal and anomalous behavior. In contrast, the lower performance observed in Cluster 3 reflects the compounding effects of data sparsity, higher variability in transactional behavior, and potential labeling inconsistencies.

V. Conclusion

The study underscores the critical role of data quality in anomaly detection. The reliance on manually labeled data introduces the risk of human error, while inconsistencies in the application of fraud response codes by some entities added further noise to the dataset. These issues not only affected the model's performance but also led to the exclusion of certain entities from the study, reducing the dataset's size and diversity. Addressing these challenges through standardized labeling processes, improved fraud management practices, and enhanced feature engineering could further refine the model and improve its applicability.

Despite these challenges, the results demonstrate the scalability and adaptability of the proposed approach. The clustering phase ensured that the model accounted for the diverse transactional behaviors across entities, while the anomaly detection phase proved effective in identifying deviations from normal behavior, particularly in well-defined clusters. The two-phase methodology offers significant potential for financial institutions to strengthen their fraud detection capabilities, reduce false positives and negatives, and enhance customer trust.

Looking forward, future iterations of this approach should focus on mitigating the impact of noisy and sparse data. Incorporating advanced data augmentation techniques, developing semi-automated labeling systems would address these limitations and further improve model reliability. Additionally, applying this methodology to other financial contexts or regions could validate its generalizability and provide insights into the adaptability of the proposed framework.

In conclusion, this study provides a foundation for a more adaptive, accurate, and scalable anomaly detection system tailored to the cooperative financial sector. By addressing data quality and operational challenges, this approach has the potential to significantly enhance fraud management, ensuring a secure and seamless experience for customers while supporting the operational needs of financial institutions.

References

[1] Y. Zhang, G. Liu, L. Zheng, and C. Yan, "A hierarchical clustering strategy of processing class imbalance and its application in fraud detection," in *2019 IEEE 21st International Conference on High Performance Computing and Communications; IEEE 17th International Conference on Smart City; IEEE 5th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*. IEEE, Aug. 2019, p. 1810–1816. [Online]. Available: http://dx.doi.org/10.1109/HPCC/SmartCity/DSS.2019.00249

[2] M. A. Islam, M. A. Uddin, S. Aryal, and G. Stea, "An ensemble learning approach for anomaly detection in credit card data with imbalanced and overlapped classes," *Journal of Information Security and Applications*, vol. 78, p. 103618, Nov. 2023. [Online]. Available: http://dx.doi.org/10.1016/j.jisa.2023.103618

[3] F. Carcillo, Y.-A. Le Borgne, O. Caelen, Y. Kessaci, F. Oblé, and G. Bontempi, "Combining unsupervised and supervised learning in credit card fraud detection," *Information Sciences*, vol. 557, p. 317–331, May 2021. [Online]. Available: http://dx.doi.org/10.1016/j.ins.2019.05.042

[4] S. Jiang, R. Dong, J. Wang, and M. Xia, "Credit card fraud detection based on unsupervised attentional anomaly detection network," *Systems*, vol. 11, no. 6, p. 305, Jun. 2023. [Online]. Available: http://dx.doi.org/10.3390/systems11060305

[5] Z. Zaffar, F. Sohrab, J. Kanniainen, and M. Gabbouj, "Credit card fraud detection with subspace learning-based one-class classification," 2023. [Online]. Available: https://arxiv.org/abs/2309.14880

[6] M. Hejazi and Y. P. Singh, "One-class support vector machines approach to anomaly detection," *Applied Artificial Intelligence*, vol. 27, no. 5, p. 351–366, May 2013. [Online]. Available: http://dx.doi.org/10.1080/08839514.2013.785791

[7] M. R. MOHAIMIN, M. Sumsuzoha, M. A. H. Pabel, and F. Nasrullah, "Detecting financial fraud using anomaly detection techniques: A comparative study of machine learning algorithms," *Journal of Computer Science and Technology Studies*, vol. 6, no. 3, p. 01–14, Jun. 2024. [Online]. Available: http://dx.doi.org/10.32996/jcsts.2024.6.3.1

[8] C. Phua, V. Lee, K. Smith, and R. Gayler, "A comprehensive survey of data mining-based fraud detection research," 2010. [Online]. Available: https://arxiv.org/abs/1009.6119

[9] Y. Chen, J. Qian, and V. Saligrama, "A new one-class svm for anomaly detection," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, May 2013, p. 3567–3571. [Online]. Available: http://dx.doi.org/10.1109/ICASSP.2013.6638322

[10] P. Vanini, S. Rossi, E. Zvizdic, and T. Domenig, "Online payment fraud: from anomaly detection to risk management," *Financial Innovation*, vol. 9, no. 1, Mar. 2023. [Online]. Available: http://dx.doi.org/10.1186/s40854-023-00470-w

[11] R. Chalapathy, A. K. Menon, and S. Chawla, "Anomaly detection using one-class neural networks," 2018. [Online]. Available: https://arxiv.org/abs/1802.06360

[12] M. Amer, M. Goldstein, and S. Abdennadher, "Enhancing one-class support vector machines for unsupervised anomaly detection," in *Proceedings of the ACM SIGKDD Workshop on Outlier Detection and Description*, ser. KDD' 13. ACM, Aug. 2013, p. 8–15. [Online]. Available: http://dx.doi.org/10.1145/2500853.2500857

[13] 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, May 2022. [Online]. Available: http://dx.doi.org/10.1016/j.eswa.2021.116429

[14] T. Pourhabibi, K.-L. Ong, B. H. Kam, and Y. L. Boo, "Fraud detection: A systematic literature review of graph-based anomaly detection approaches," *Decision Support Systems*, vol. 133, p. 113303, Jun. 2020. [Online]. Available: http://dx.doi.org/10.1016/j.dss.2020.113303

[15] J. Perols, "Financial statement fraud detection: An analysis of statistical and machine learning algorithms," *AUDITING: A Journal of Practice amp; Theory*, vol. 30, no. 2, p. 19–50, May 2011. [Online]. Available: http://dx.doi.org/10.2308/ajpt-50009

[16] A. Ali, S. Abd Razak, S. H. Othman, T. A. E. Eisa, A. Al-Dhaqm, M. Nasser, T. Elhassan, H. Elshafie, and A. Saif, "Financial fraud detection based on machine learning: A systematic literature review," *Applied Sciences*, vol. 12, no. 19, p. 9637, Sep. 2022. [Online]. Available: http://dx.doi.org/10.3390/app12199637

[17] N. Shirodkar, P. Mandrekar, R. S. Mandrekar, R. Sakhalkar, K. Chaman Kumar, and S. Aswale, "Credit card fraud detection techniques – a survey," in *2020 International Conference on Emerging Trends in Information Technology and Engineering (ic-ETITE)*, 2020, pp. 1–7.

[18] A. Bakumenko and A. Elragal, "Detecting anomalies in financial data using machine learning algorithms," *Systems*, vol. 10, no. 5, p. 130, Aug. 2022. [Online]. Available: http://dx.doi.org/10.3390/systems10050130

[19] M. R. Hasan, M. S. Gazi, and N. Gurung, "Explainable ai in credit card fraud detection: Interpretable models and transparent decision-making for enhanced trust and compliance in the usa," *Journal of Computer Science and Technology Studies*, vol. 6, no. 2, p. 01–12, Apr. 2024. [Online]. Available: http://dx.doi.org/10.32996/jcsts.2024.6.2.1

[20] M. Islam, "Artificial intelligence exploring its applications across industries," *Journal of Artificial Intelligence General science (JAIGS) ISSN:3006-4023*, vol. 2, no. 1, p. 20–24, Feb. 2024. [Online]. Available: http://dx.doi.org/10.60087/jaigs.v2i1.42