

A Hybrid Machine Learning Framework for Cross-Domain Fraud Detection Using Multimodal Data from Healthcare and Financial Systems

James Oliver Mateo

AI/ML Software Developer, USA

ABSTRACTSS

Fraud detection remains a critical challenge across domains such as healthcare and finance, where fraudulent behaviors often differ in complexity, structure, and data representation. This paper proposes a hybrid machine learning framework that leverages multimodal data integration to improve cross-domain fraud detection capabilities. The approach combines decision trees, neural networks, and ensemble techniques to capture non-linear and latent patterns in both structured data (e.g., transaction records, billing codes) and unstructured data (e.g., clinical notes, textual claims).

By unifying heterogeneous data sources from healthcare and financial systems, the model demonstrates enhanced adaptability to evolving fraud trends and context-specific anomalies. The hybrid architecture supports deeper pattern recognition across domains and enables improved generalization in environments where data variability and concept drift are prevalent.

Experimental results validate the effectiveness of the proposed approach, showing significant improvements in detection accuracy and robustness compared to baseline single-domain models. This research highlights the potential of cross-domain machine learning and multimodal integration in developing next-generation fraud detection systems.

KEYWORDS: Fraud detection, Cross-domain learning, Multimodal data, Hybrid machine learning, Healthcare analytics, Financial fraud.

CITATION: James Oliver Mateo. (2024). A Hybrid Machine Learning Framework for Cross-Domain Fraud Detection Using Multimodal Data from Healthcare and Financial Systems. *International Journal of Computing Science and Systems (IJCSS)*, 5(1), 8–15.

1. Introduction

Fraud poses a persistent and costly challenge in both healthcare and financial industries. According to global reports, fraud losses are projected to exceed \$5 trillion annually, with over 10% of healthcare spending lost due to fraudulent activities. The disparate nature of these domains—structured numeric data in finance vs. complex clinical notes and claims in healthcare—requires innovative, cross-domain solutions.

This study proposes a hybrid machine learning (ML) framework that combines supervised and unsupervised models for fraud detection. By leveraging multimodal data (textual, numerical, categorical) and integrating features across both sectors, this approach aims to detect anomalies and fraudulent patterns more effectively. The proposed system includes data preprocessing, multimodal fusion, model training, and performance evaluation.



2. Literature Review

Several researchers have addressed fraud detection using various machine learning techniques, although cross-domain and multimodal integration is still emerging.

1. **Ngai et al. (2011)** explored common ML models such as decision trees, neural networks, and SVMs for financial fraud detection. Their review showed the promise of ensemble models in complex fraud scenarios.
2. **Ravisankar et al. (2011)** introduced a hybrid approach combining logistic regression and decision trees to detect fraudulent financial statements.

3. **Bauder et al. (2017)** investigated text mining for unstructured healthcare data to detect fraud. The study found natural language processing (NLP) essential in identifying false insurance claims.
4. **Binh et al. (2020)** utilized autoencoders and clustering algorithms to detect financial anomalies without labels, showing that unsupervised methods can effectively flag unknown fraud patterns.
5. **Karami et al. (2023)** proposed an interpretable fraud detection model using SHAP (SHapley Additive exPlanations), ensuring transparency and trust in black-box models.

3. Proposed Methodology

The proposed system introduces a **hybrid machine learning architecture** that merges advanced modeling with multimodal data integration to detect fraudulent behavior across healthcare and financial domains. Given the distinct nature of data in each field—structured transactions in finance and semi-structured claims or unstructured medical texts in healthcare—the system is designed to normalize, unify, and analyze diverse inputs using both traditional and deep learning models.

At its core, the methodology comprises **four interconnected modules**:

- **Data Preprocessing,**
- **Feature Engineering and Fusion,**
- **Hybrid Modeling,** and
- **Evaluation Metrics.**

Each of these components is strategically built to handle the **heterogeneous and cross-domain challenges** involved in fraud detection.

3.1 Data Preprocessing

The preprocessing stage involves transforming raw data into a consistent, clean format suitable for machine learning.

For structured data like **transaction logs** or **billing records**, this includes handling missing values, removing outliers, normalizing numerical columns, and encoding categorical variables (e.g., merchant type, provider type).

For unstructured data such as **medical notes** or **diagnostic summaries**, we use natural language processing (NLP) techniques including tokenization, lemmatization, and stop-word removal. Tools such as **NLTK** and **SpaCy** are employed to create clean text sequences.

3.2 Feature Engineering and Fusion

In this phase, we extract meaningful features from both textual and numeric sources to create a **unified feature space**.

TF-IDF (Term Frequency-Inverse Document Frequency) is applied to textual data like clinical notes to transform them into sparse vectors.

Embeddings such as Word2Vec or BERT can also be introduced to encode richer semantic relationships when needed.

For tabular data, domain-specific features are engineered—for example, “average billing per visit,” “time between claims,” or “transaction frequency deviation.”

Once domain-specific features are extracted, the data streams are **merged using early fusion** techniques into a single vector, which ensures joint learning during model training.

Table 1: System Components and Roles

Model Component	Role	Technology Used
Preprocessing	Clean & standardize data	Pandas, NLTK
Feature Fusion	Merge text & tabular features	TF-IDF, Word Embeddings
Classifier 1	Rule-based detection	Decision Tree (CART)
Classifier 2	Anomaly detection	Autoencoder (Keras/PyTorch)
Classifier 3	Deep pattern recognition	LSTM / Dense Neural Net

4. Data Description and Integration Strategy

The effectiveness of any machine learning-based fraud detection system is highly dependent on the quality and diversity of its input data. In this study, we utilize a **multimodal dataset** composed of de-identified records from two domains—**healthcare insurance claims** and **financial transactions**. Each data source contains structurally and semantically distinct types of features, presenting significant integration challenges.

4.1 Data Sources

Healthcare Insurance Data

This dataset includes:

ICD Codes: Standardized codes representing diseases and medical conditions. **Diagnosis and Procedure Notes:** Semi-structured or free-text fields capturing physician notes and patient care summaries.

Prescription Logs: Structured fields with drug codes, dosage information, and refill counts. **Billing Metadata:** Including provider ID, claim submission date, and payment amounts.

These records often contain redundancies, domain-specific acronyms, and temporal inconsistencies.

Financial Transaction Data

This dataset includes:

Transaction Volume and Frequency: Total amounts, transaction counts, and temporal sequences. **Geolocation Coordinates:** Location of merchants and ATM use, mapped via reverse geocoding. **Merchant Categories:** Classified based on type (e.g., groceries, entertainment, electronics). **Account History and Risk Scores:** Behavioral flags such as rapid purchases, high refunds, etc. These attributes are highly structured but exhibit temporal volatility and hidden correlations.

5. Results and Discussion

The hybrid model proposed in this study demonstrates superior performance in detecting fraudulent activities across healthcare and financial domains by effectively leveraging multimodal data and combining diverse machine learning paradigms. To validate the system, we conducted a series of experiments using preprocessed and integrated data described in Section 4. The models were trained on 80% of the dataset and evaluated on a holdout test set (20%) using standard classification metrics.

5.1 Performance Metrics and Evaluation Criteria

The following metrics were used to assess model performance:

- **Accuracy:** The proportion of correct predictions among the total instances evaluated.
- **Precision:** The fraction of correctly identified frauds among all instances flagged as fraudulent.
- **Recall:** The ability of the model to detect actual fraudulent cases from all frauds present.
- **F1-Score:** The harmonic mean of precision and recall, providing a balanced evaluation of the model's detection capability.

5.2 Model Comparison

Four models were evaluated: three baseline models—**Decision Tree**, **Autoencoder**, and **LSTM**—and the **proposed Hybrid Ensemble Model**. Each model was assessed on the same dataset and test set.

Table 2: Model Performance Comparison

Model	Accuracy	Precision	Recall	F1 Score
Decision Tree	85.2%	0.80	0.76	0.78
Autoencoder	88.3%	0.85	0.82	0.83
LSTM	90.1%	0.88	0.87	0.875
Hybrid (Ours)	94.6%	0.93	0.89	0.91

5.3 Interpretation of Results

The **Decision Tree model**, while interpretable and fast, lacked the depth to detect complex fraud patterns, leading to a modest F1-score of 0.78. It performed well on obvious anomalies but missed subtle multi-feature interactions.

The **Autoencoder**, an unsupervised model, performed better by identifying deviations from typical patterns. It was particularly useful in highlighting novel or zero-day fraud behaviors not seen in the training data. However, it had a slightly higher false positive rate due to its sensitivity to data distribution.

The **LSTM model**, designed for temporal sequence analysis, provided even better performance—accurately capturing fraudulent trends based on event ordering and frequency, common in financial transactions and claims history. It reached an F1-score of 0.875, balancing detection power with contextual learning.

6. Conclusion

This study presented a hybrid machine learning model that integrates healthcare and financial data to improve fraud detection. By utilizing multimodal data and combining diverse learning paradigms, our system demonstrates improved accuracy and generalization. Future work will focus on real-time fraud detection and the inclusion of reinforcement learning for adaptive defense.

Conflicts of Interest: The authors declare no conflicts of interest.

Publisher's Note: The views and statements presented in this article are solely those of the authors and do not represent the positions of the publisher, the editors, or the reviewers.

References

- [1] Ngai, E.W.T., Hu, Y., Wong, Y.H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection. *Decision Support Systems*, 50(3), 559–569.
- [2] Gundaboina, A. (2022). Quantum Computing and Cloud Security: Future-Proofing Healthcare Data Protection. *International Journal for Multidisciplinary Research*, 4(4), 1–12. <https://doi.org/10.36948/ijfmr.2022.v04i04.61014>
- [3] Ravisankar, P., & Ravi, V. (2011). Financial statement fraud detection using data mining techniques. *European Journal of Operational Research*, 203(2), 591–605.
- [4] Bauder, R.A., Khoshgoftaar, T.M., & Bullard, L. (2017). A survey of Medicare data preprocessing and machine learning techniques for fraud detection. *Health Information Science and Systems*, 5(1), 5. 0026-4
- [5] Binh, T.H., Pham, N.T., & Le, Q.N. (2020). Fraud detection in financial data using unsupervised learning. *Journal of Finance and Data Science*, 6(1), 12–26.
- [6] Gundaboina, A. (2023). Securing Non-Human Identities (NHIs) in Cloud-Native Healthcare Systems. *International Journal of Innovative Research in Engineering & Multidisciplinary Physical Sciences*, 11(5), 1–12. <https://doi.org/10.37082/IJIRMPS.v11.i5.232621>
- [7] Karami, A., & Gupta, N. (2023). Explainable AI in fraud detection: A SHAP-based evaluation. *Expert Systems with Applications*, 202(1), 117531.
- [8] West, J., and M. Bhattacharya. "Intelligent Financial Fraud Detection: A Comprehensive Review." *Computers & Security*, vol. 57, 2016, pp. 47–66.
- [9] Gundaboina, A. (2023). Securing IoT Devices in Healthcare: Endpoint Protection for Patient Monitoring Systems. *International Journal of Innovative Research in Engineering & Multidisciplinary Physical Sciences*, 11(5), 1–11. <https://doi.org/10.37082/IJIRMPS.v11.i5.232623>
- [10] Ding, Y., and H. Peng. "Multimodal Machine Learning for Healthcare: Recent Advances and Applications." *IEEE Reviews in Biomedical Engineering*, 2022.
- [11] Buczak, A. L., and E. Guven. "A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection." *IEEE Communications Surveys & Tutorials*, vol. 18, no. 2, 2016, pp. 1153–1176.

- [12] Gundaboina, A. (2023). Data Loss Prevention in Healthcare: Advanced Strategies for Protecting PHI in Cloud Environments. *Journal of Artificial Intelligence, Machine Learning & Data Science*, 1(2), 3045–3051. <https://doi.org/10.51219/JAIMLD/anjan-gundaboina/628>
- [13] Jurgovsky, J., et al. "Sequence Classification for Credit-Card Fraud Detection." *Expert Systems with Applications*, vol. 100, 2018, pp. 234–245.
- [14] Ghosh, S., and D. L. Reilly. "Credit Card Fraud Detection with a Neural-Network." *Proceedings of the 27th Hawaii International Conference on System Sciences*, 1994, pp. 621–630.
- [15] Gundaboina, A. (2023). Endpoint Detection and Response (EDR) in Healthcare: Mitigating Threats on Critical Devices. *Artificial Intelligence, Machine Learning & Data Science*, 1(2), 3107–3114. <https://doi.org/10.51219/JAIMLD/ramesh-potla/637>
- [16] Zhou, Y., and K. Kapoor. "Cross-Domain Transfer Learning in Healthcare Fraud Detection: A Comparative Analysis." *Journal of Medical Systems*, vol. 45, 2021, article 79.
- [17] Chandola, V., A. Banerjee, and V. Kumar. "Anomaly Detection: A Survey." *ACM Computing Surveys*, vol. 41, no. 3, 2009, article 15.