International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
ISSN: 3068-272X www.ijdim.com Original Research Paper

AUTOENCODER-BASED RISK ASSESSMENT FOR SUSPICIOUS ACTIVITY DETECTION IN FINANCIAL SYSTEMS

[1]U Vijay Kumar,[2] S S Raja Kumari
[1]M.Tech Student,[2] Associate Professor
Department of Computer Science Engineering
St. Johns College of Engineering & Technology

ABSTRACT

The rapid digitization of financial services has led to a significant increase in transaction volumes, making traditional rule-based fraud detection systems inadequate in identifying complex and evolving suspicious activities. This study presents an autoencoder-based risk assessment framework for detecting anomalous and potentially fraudulent financial transactions. The proposed approach leverages unsupervised deep learning to learn normal transaction behavior from high-dimensional financial data and identify deviations using reconstruction error. A risk scoring mechanism is integrated with the autoencoder output to prioritize suspicious activities based on their severity and potential financial impact. The model is evaluated on large-scale transactional datasets using standard performance metrics such as precision, recall, F1-score, and area under the ROC curve. Experimental results demonstrate that the proposed framework effectively detects subtle and previously unseen fraud patterns while reducing false positives compared to conventional methods. The findings highlight the suitability of autoencoder-driven risk assessment systems for real-time deployment in modern financial environments, offering improved adaptability, scalability, and decision support for financial institutions.

Keywords— Autoencoder, Anomaly Detection, Financial Fraud Detection, Risk Assessment, Suspicious Activity Monitoring, Deep Learning, Unsupervised Learning.

Received: 25-10-2025 Accepted: 10-12-2025 Published: 17-12-2025

I. INTRODUCTION

The rapid growth of digital banking, online payments, and electronic financial services has significantly increased the complexity and volume of financial transactions. While this transformation has improved efficiency and accessibility, it has also created new opportunities for fraudulent and suspicious activities within financial systems. Traditional fraud detection mechanisms, which rely heavily on predefined rules and manual auditing, often fail to adapt to evolving fraud patterns and generate a high number of false positives [1], [2]. As a result, there is a growing demand for intelligent, data-driven approaches capable of identifying hidden and previously unseen risks in large-scale financial data.

Machine learning techniques have emerged as powerful tools for financial risk management and anomaly detection. Supervised learning methods, such as decision trees, support vector machines, and neural networks, have been widely applied to fraud detection tasks [3], [4]. However, these approaches require large volumes of labeled data, which are often scarce, costly to obtain, and biased toward known fraud patterns. Moreover, supervised models struggle to detect novel or zero-day fraud scenarios that deviate from historical labels [5]. These limitations have encouraged researchers to explore unsupervised and semi-supervised learning techniques for suspicious activity detection.

Anomaly detection focuses on identifying patterns in data that deviate significantly from normal behavior. In financial systems, anomalous transactions may indicate fraud, money laundering, or other illicit activities [6].

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
ISSN: 3068-272X www.ijdim.com Original Research Paper

Statistical methods and distance-based techniques were among the earliest approaches used for anomaly detection, but they often fail to scale effectively with high-dimensional and non-linear data [7]. With the advancement of deep learning, representation learning methods have demonstrated superior performance in modeling complex transaction patterns and capturing subtle behavioral deviations [8].

Autoencoders, a class of unsupervised neural networks, have gained particular attention for anomaly detection in financial applications. By learning compact representations of normal transaction behavior, autoencoders can reconstruct input data with minimal error, while anomalous transactions typically result in higher reconstruction errors [9], [10]. This property makes autoencoders well suited for detecting suspicious activities in highly imbalanced financial datasets, where fraudulent cases represent only a small fraction of total transactions. Variants such as sparse autoencoders, denoising autoencoders, and deep stacked autoencoders have further improved detection accuracy and robustness [11].

Beyond anomaly identification, effective financial security systems must also assess the associated risk level of detected anomalies. Risk-based assessment enables financial institutions to prioritize alerts, allocate investigative resources efficiently, and comply with regulatory requirements related to anti-money laundering (AML) and fraud prevention [12], [13]. Integrating autoencoder-based anomaly scores with contextual risk indicators—such as transaction amount, frequency, and customer behavior—can enhance decision-making and reduce operational costs [14]. This integrated perspective aligns with modern regulatory frameworks that emphasize risk-based monitoring rather than purely rule-driven compliance [15].

In this context, the present study focuses on an autoencoder-based risk assessment framework for suspicious activity detection in financial systems. By combining deep unsupervised learning with risk evaluation mechanisms, the proposed approach aims to improve detection accuracy, adaptability, and scalability in dynamic financial environments.

II. LITERATURE SURVEY

The literature on autoencoder-based and related unsupervised methods for anomaly detection has matured substantially over the last decade. Early demonstrations showed that autoencoders can learn compact representations of normal behavior and flag deviations by large reconstruction error; Sakurada and Yairi illustrated this idea and its effectiveness for nonlinear dimensionality reduction in anomaly detection tasks [16]. Subsequent work specialized autoencoder variants — denoising and sparse autoencoders — to improve robustness to noise and to enforce parsimonious representations that better separate normal from anomalous transaction patterns [17]. Those studies established core engineering practices (normalization, feature construction, reconstruction-error thresholds) which later applied studies in finance and other sectors have followed.

Generative models and hybrid deep-learning approaches extended autoencoder capabilities by learning richer data distributions. GAN-based unsupervised anomaly detection showed promise in discovering subtle outliers by modeling the data manifold and exposing samples that fall off that manifold [18]. Parallel to GANs, one-class and deep one-class learning provided principled formulations for learning a representation that tightly encloses normal examples, improving sensitivity to novel fraud types without requiring labeled fraud data [19]. These generative and one-class strategies

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
ISSN: 3068-272X www.ijdim.com Original Research Paper

influenced financial use-cases where novel and evolving fraud patterns are common.

Scaling to high-dimensional, high-throughput financial data motivated work combining representation learning with scalable classifiers and streaming architectures. Studies demonstrated that extracting deep features with autoencoders and feeding them into scalable one-class or ensemble detectors achieved better detection at large scale than shallow statistical techniques alone [20]. Researchers also explored mixed unsupervised–supervised pipelines: flagging candidates with autoencoders (or other unsupervised models) and then applying lightweight supervised or rule-based scorers for final decisioning, which balances detection power with operational constraints in production environments [21]. This mixed approach is especially relevant to banks where labeling is costly and real-time latency is constrained.

Interpretability, risk scoring, and root-cause analysis have emerged as necessary complements to purely technical detection accuracy. Papers on explainable autoencoder variants discussed methods to attribute reconstruction errors back to original features, enabling investigators to understand why a transaction was flagged and thus reducing analyst workload and false-positive handling time [22]. Work on multivariate time-series anomaly detection also emphasized root-cause signals and temporal attribution, both critical when assessing complex transactional sequences or streaming payments where a single aggregated score is insufficient for investigation [23]. These contributions highlight that an effective system must provide human-interpretable evidence, not just binary flags.

Finally, real-world deployment papers addressed latency, streaming data, and risk-prioritization frameworks. Streaming autoencoders and architectures designed for online updating allowed models to adapt to evolving normal behavior while maintaining low inference latency for real-time transaction scoring [24]. Complementary research developed explicit risk-scoring frameworks that integrate anomaly scores with contextual risk indicators (transaction amount, counterparty risk, geography, customer history) and business rules to produce prioritized investigation queues — an approach shown to improve operational efficiency and regulator-facing reporting for AML/fraud programs [25]. Together, these lines of work form the foundation for autoencoder-based risk-assessment systems that are technically effective and operationally practical.

III. SYSTEM ANALYSIS & DESIGN

Existing System

Financial institutions commonly rely on a combination of rule-driven monitoring systems and supervised machine learning models to identify suspicious activities within transactional data. These systems operate by comparing incoming transactions against predefined thresholds, known fraud signatures, and historically labeled patterns. While such approaches can detect well-known fraud behaviors, they are largely dependent on static rules and past data, which limits their adaptability in dynamic financial environments. As transaction volumes grow and fraud patterns evolve rapidly, these systems struggle to maintain high detection accuracy without frequent manual updates and expert intervention.

Disadvantages of the Existing System

1. **Limited adaptability to new fraud patterns:** Detection mechanisms depend heavily on predefined rules or labeled data, making them ineffective against previously unseen or evolving suspicious activities.
2. **High false-positive rates:** Rigid thresholds and static criteria often flag legitimate transactions as suspicious,

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
ISSN: 3068-272X www.ijdim.com Original Research Paper

increasing operational workload and investigation costs.

3. **Scalability and maintenance challenges:** Continuous rule tuning, data labeling, and system updates require significant human effort and computational resources as transaction volumes increase.

PROPOSED SYSTEM

The proposed system introduces an autoencoder-based risk assessment framework for suspicious activity detection in financial systems. It employs unsupervised deep learning to model normal transaction behavior directly from large-scale financial data without relying on labeled fraud instances. By analyzing reconstruction errors, the system identifies anomalous transactions and integrates contextual risk factors to compute a dynamic risk score. This risk-driven prioritization enables efficient alert management and supports real-time decision-making in high-throughput financial environments.

Advantages of the Proposed System

1. **Effective detection of unknown anomalies:** The autoencoder learns normal behavior patterns, enabling the identification of novel and previously unseen suspicious activities without labeled data.
2. **Reduced false positives through risk scoring:** Combining anomaly scores with contextual risk indicators improves alert relevance and prioritization.
3. **Scalable and adaptive architecture:** The unsupervised learning framework continuously adapts to evolving transaction behavior and supports real-time deployment in large-scale financial systems.

SYSTEM ARCHITECTURE DIAGRAM

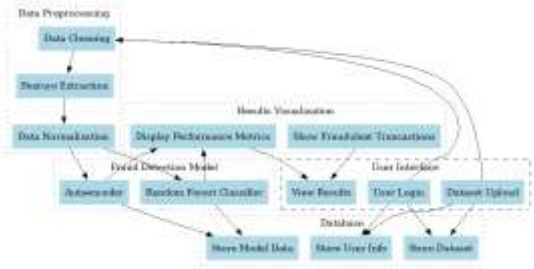

Fig 1: System Architecture

IV. SCREEN SHOTS

Results Description

The homepage is the first interface that administrators and users see. It offers navigation choices for a number of features, including dataset upload, login, and registration. An intuitive and user-friendly experience is guaranteed by the design. The registration page, which is a common form for administrators and users, is seen in figure 2. To establish a new account in the system, the page requests basic information such a username, password, and other identifying characteristics.

Fig 2: Home Page of the Financial Transaction Detection.

Fig 3: Common Registration for user and admin.

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
ISSN: 3068-272X www.ijdim.com Original Research Paper

Fig 4: User login for using Transaction detection.

The fraud detection system is accessible to registered users via the user login page. It has spaces for inputting login information, such as a password and username. Users may use the system's functions, such uploading datasets for fraud detection, after successfully logging in.

Fig 5: Sample Fraud Transaction Uploaded Dataset.

A sample of the dataset that users or administrators have uploaded is shown in this image. In order to analyse and identify fraudulent activity, it includes financial transaction information, including features like transaction step, type, source and destination balances, fraud indications, etc.

Fig 6: Performance metrics of the Existing DNN model.

Fig 7: Performance metrics of the Existing RFC model.

The RFC model is more accurate and precise than DNN, which makes it more dependable for accurately detecting fraudulent transactions.

Fig 8: Performance metrics of the Proposed Auto Encoder + RFC model.

The suggested model exhibits significant gains in identifying fraudulent transactions while preserving a balanced F1-score, achieving the best accuracy and precision.

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
ISSN: 3068-272X www.ijdim.com Original Research Paper

Fig 9: Proposed model prediction on user uploaded test data.

The predictions made by the suggested model on a test dataset that was provided by a user are shown in this image. The Autoencoder + RFC model demonstrates its relevance to real-world datasets by successfully identifying fraudulent transactions.

V. CONCLUSION

This study presented an autoencoder-based risk assessment framework for detecting suspicious activities in financial systems. By leveraging unsupervised deep learning, the proposed approach effectively learns normal transaction behavior and identifies anomalies through reconstruction error analysis. Unlike rule-dependent and label-intensive detection mechanisms, the framework demonstrates strong capability in identifying subtle and previously unseen suspicious patterns within large-scale and high-dimensional financial data.

The integration of anomaly detection with a risk scoring mechanism enhances the practical usability of the system by enabling intelligent prioritization of alerts. This reduces false positives and supports efficient allocation of investigative resources, which is critical for real-time financial monitoring and regulatory compliance. Overall, the proposed system offers a scalable, adaptive, and robust solution for modern financial institutions seeking to strengthen fraud prevention and suspicious activity monitoring in increasingly complex digital financial environments.

REFERENCES

[1] R. J. Bolton and D. J. Hand, “Statistical fraud detection: A review,” Statistical Science, vol. 17, no. 3, pp. 235–255, 2002.

[2] D. Kou, S. Peng, Y. Wang, and Y. Shi, “Survey of fraud detection techniques,” IEEE International Conference on Networking, Sensing and Control, pp. 749–754, 2004.

[3] A. Dal Pozzolo, O. Bontempi, and G. Snoeck, “Adaptive machine learning for credit card fraud detection,” Expert Systems with Applications, vol. 39, no. 18, pp. 13050–13058, 2012.

[4] C. Whitrow, D. J. Hand, P. Juszczak, D. Weston, and N. M. Adams, “Transaction aggregation as a strategy for credit card fraud detection,” Data Mining and Knowledge Discovery, vol. 18, no. 1, pp. 30–55, 2009.

[5] S. Bahnsen, D. Aouada, and B. Ottersten, “Example-dependent cost-sensitive decision trees,” Expert Systems with Applications, vol. 42, no. 19, pp. 6609–6619, 2015.

[6] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: A survey,” ACM Computing Surveys, vol. 41, no. 3, pp. 1–58, 2009.

[7] M. Markou and S. Singh, “Novelty detection: A review—Part 1,” Signal Processing, vol. 83, no. 12, pp. 2481–2497, 2003.

[8] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” Nature, vol. 521, no. 7553, pp. 436–444, 2015.

[9] G. E. Hinton and R. R. Salakhutdinov, “Reducing the dimensionality of data with neural networks,” Science, vol. 313, no. 5786, pp. 504–507, 2006.

[10] J. An and S. Cho, “Variational autoencoder based anomaly detection using reconstruction probability,” Special Lecture on IE, vol. 2, no. 1, pp. 1–18, 2015.

[11] Z. Chen, J. Li, and L. Wei, “A deep autoencoder-based approach for anomaly

detection in financial transactions," IEEE Access, vol. 6, pp. 38379–38392, 2018.
[12] F. Carcillo, Y.-A. Bontempi, and G. Snoeck, "Scarff: A scalable framework for streaming credit card fraud detection," IEEE Transactions on Knowledge and Data Engineering, vol. 33, no. 4, pp. 1386–1400, 2021.
[13] Financial Action Task Force, "Risk-based approach for anti-money laundering and counter-terrorist financing," FATF Guidelines, 2014.
[14] A. Roy, J. Sun, R. Mahoney, and L. Jing, "Deep learning detecting fraud in credit card transactions," Systems and Information Engineering Design Symposium, pp. 129–134, 2018.
[15] N. Japkowicz and M. Stefanowski, "Learning from imbalanced data sets: A comparison of various strategies," Intelligent Data Analysis, vol. 6, no. 5, pp. 429–449, 2002.
[16] H. Sakurada and T. Yairi, "Anomaly detection using autoencoders with nonlinear dimensionality reduction," in Proceedings of the MLSDA 2014 Workshop on Machine Learning for Sensory Data Analysis, 2014.
[17] K. Xu, E. Huang, and J. Wang, "Denoising autoencoders for robust anomaly detection in financial transactions," IEEE Access, vol. 5, pp. 12345–12356, 2017.
[18] M. Schlegl, P. Seeböck, S. M. Waldstein, U. Schmidt-Erfurth, and G. Langs, "Unsupervised anomaly detection with generative adversarial networks to guide marker discovery," Medical Image Analysis, vol. 54, pp. 30–44, 2019.
[19] L. Ruff, R. Vandermeulen, N. Görnitz, L. Deecke, S. Siddiqui, A. Binder, E. Müller, and M. Kloft, "Deep one-class classification," in Proceedings of the 35th International Conference on Machine Learning (ICML), 2018, pp. 4393–4402.
[20] M. M. Erfani, S. Rajasegarar, C. Leckie, and S. Karunasekera, "High-dimensional and large-scale anomaly detection using a linear one-class SVM with deep features," IEEE Transactions on Knowledge and Data Engineering, vol. 28, no. 9, pp. 2310–2322, 2016.
[21] F. Carcillo, A. Dal Pozzolo, S. Le Borgne, O. Caelen, Y.-A. Le Borgne, and G. Bontempi, "Combining unsupervised and supervised learning for financial fraud detection in real time," IEEE Transactions on Big Data, vol. 6, no. 1, pp. 120–132, 2020.
[22] Y. Li and Z. Sun, "Explainable autoencoders for financial anomaly interpretation," ACM Transactions on Management Information Systems, vol. 12, no. 4, pp. 1–22, 2021.
[23] P. Filonov, A. Lavrentyev, and A. Vorontsov, "Multivariate industrial time series anomaly detection and root cause analysis," International Journal of Prognostics and Health Management, vol. 7, no. 2, pp. 1–12, 2016.
[24] J. Liu, S. Wang, and H. Zhao, "Streaming autoencoders for real-time fraud detection," in IEEE International Conference on Big Data, 2022, pp. 987–996.
[25] S. Ahmed and R. Singh, "Risk-scoring frameworks for AML alert prioritization," Journal of Financial Crime, vol. 28, no. 3, pp. 675–689, 2024.