

AI-Driven Anomaly Detection Frameworks for Real-Time Payment System Compliance

Padma Vayuvegula^{1,*}

¹Department of Banking, University of the Cumberlands, Williamsburg, Kentucky, United States of America.
padmavayuvegula1@gmail.com¹

Abstract: Finance is moving at light speed towards Real-Time Payment Systems (RTPS), which, with as great a head start as they have over unimaginable convenience and speed, are beset with humongous compliance risk in the guise of hi-tech fraud and money laundering. The traditional rule-based monitoring architecture cannot keep pace with the speed and volume of such transactions; therefore, there is a historical need for intelligent analytical solutions. The current paper proposes a novel hybrid AI model for RPS anomaly detection. The model utilises the synergy of maximising the power of employing an Isolation Forest algorithm to effectively detect outliers and that of an Autoencoder neural network to learn non-linear, implicit features from transactional data. The data used was a synthetically generated dataset of 432 samples, constructed to include both regular and anomalous transactions. We developed and tested our environment using Python and libraries such as Scikit-learn for Isolation Forest and TensorFlow/Keras for Autoencoder. Our results indicate that the hybrid model presented in this paper achieves higher F1 Scores and AUC-ROC than standalone models, and significantly improves precision and efficiency in identifying fraudulent transactions. This paper presents an effective and scalable solution for banks to enhance their compliance processes and ensure the security of real-time payment systems.

Keywords: Anomaly Detection; Artificial Intelligence; Real-Time Payments; Financial Compliance; Machine Learning; Financial Transactions; Deep Learning; Transactional Data; Precision and Efficiency.

Received on: 03/12/2024, **Revised on:** 08/02/2025, **Accepted on:** 22/03/2025, **Published on:** 05/09/2025

Journal Homepage: <https://www.fmdbpublish.com/user/journals/details/FTSCL>

DOI: <https://doi.org/10.69888/FTSCL.2025.000430>

Cite as: P. Vayuvegula, “AI-Driven Anomaly Detection Frameworks for Real-Time Payment System Compliance,” *FMDB Transactions on Sustainable Computer Letters*, vol. 3, no. 3, pp. 150–157, 2025.

Copyright © 2025 P. Vayuvegula, licensed to Fernando Martins De Bulhão (FMDB) Publishing Company. This is an open access article distributed under [CC BY-NC-SA 4.0](#), which permits copying, remixing, redistributing, transformation, and building upon the material in any medium so long as the original work is properly cited.

1. Introduction

The literature on abnormality detection in financial transactions has grown significantly over the last two decades, from simple statistical methods to highly advanced deep learning techniques. It began with initial attempts that were largely statistical process control-based, where methods such as the Z-score or adjusted Z-score were used to identify outliers relative to the mean, as in Ahmad et al. [1]. In time series analysis, methods such as ARIMA (Autoregressive Integrated Moving Average) are used to make future trend predictions, with exceptions chosen as needed, as employed by Alghushairy et al. [2]. Such approaches, although computationally inexpensive, treated data as normally distributed and did not incorporate mechanisms to address seasonality and transaction non-stationarity, a limitation noted in Ali [3]. The inability to address multidimensionality

*Corresponding author.

made them outdated in modern systems, according to Ali et al. [4]. The next innovation came from machine learning approaches. Supervised machine learning models, such as Logistic Regression, Support Vector Machines (SVM), and Decision Trees, were employed whenever labelled data were available, achieving the best accuracy by learning decision boundaries from the available data, as in Arcos-García et al. [5]. Imbalanced labelled datasets posed challenges, and therefore, the application of unsupervised k-Means and DBSCAN algorithms to cluster similar transactions as outliers, as proposed in Banik et al. [6]. The unsupervised method avoided the use of labelled samples but was susceptible to distance measures and to high-dimensional data, a requirement suggested by Bashar and Nayak [7].

Ensemble and deep learning techniques were the way ahead in development. Model-based prediction union ensemble techniques proved to be effective. The Isolation Forest, developed specifically for anomaly detection, randomises the feature space to isolate outliers and performs well even in high-dimensional feature spaces [8]. Random Forests were applied in a hierarchical structure, with proximities between forests measured to detect anomalies, as in Boutaba et al. [9]. The focus then moved to more robust learning. Autoencoders, as a means of reconstructing in an unsupervised neural network, were designated as the main tool. They would reconstruct normal data, but abnormal data would result in beast-sized reconstruction errors, a process successfully utilised by Chen et al. [10]. RNNs and LSTMs were utilised to capture temporal dependencies and sequence transactions to predict future behaviour and trigger alerts on anomalies, as described in Choi et al. [11]. Although such models were resilient, they were not practical because they required humongous databases and humongous computational power, which were constraints for real-time use in the real world, as cited by Chen et al. [12]. This has led to hybrid models that combine the velocity of traditional approaches with the rich pattern abstraction of neural networks, which today represent the state of the art in efficient, real-time compliance monitoring.

2. Review of Literature

Anomaly detection in financial transactions has advanced significantly over the last two decades, evolving from simple statistical methods to more robust deep learning approaches. Some previous applications have extensively utilised statistical process control, employing techniques such as the Z-score or mZ-score to identify transactions outside the distribution's mean, as in Ahmad et al. [1]. For time-series data, ARIMA models have been used to forecast anticipated transaction patterns and to detect outliers as deviations, as in Alghushairy et al. [2]. These methods, although inexpensive to compute and easy to define, used normally distributed data and were unable to handle sophisticated multidimensional relationships among transactions, as highlighted by Ali [3]. They were extremely vulnerable to non-stationarity or seasonality in transaction patterns and were therefore inappropriate for modern dynamic payment systems, as highlighted by Ali et al. [4].

Classic machine learning algorithms represented the second wave of innovation. Supervised techniques, including Logistic Regression, Support Vector Machines (SVM), and Decision Trees, were employed, with examples of previous actual transactions and known frauds provided by Arcos-García et al. [5]. Models learned much more effectively by learning previous decision boundaries. Sparsity and skew in training data limited their application, however, since induced cases were an infinitesimally small fraction of all the transactions. This constraint facilitated the use of unsupervised learning, where the k-Means and DBSCAN algorithms classified similar transactions into a class and identified dissimilar transactions as suspicious, as stated by Banik et al. [6]. Such unsupervised learning algorithms were plagued by the drawback of never having been trained on labelled data. Still, they were highly sensitive to the distance function and performed very badly in high-dimensional space, a drawback noted by Bashar and Nayak [7].

Later, the development of ensemble methods leveraged the forecasting power of multiple models. Isolation Forest was successful because it was anomaly-based and automatically distinguished outliers by randomly partitioning the data space, as in Biswas and Samanta [8]. The ability to handle high-dimensional financial data with ease enabled easy integration into real-world systems. Together, Random Forest, being more classification-oriented, was even applied to anomaly detection based on proximity to points in the ensemble tree structure, for instance, by Boutaba et al. [9]. Ensemble algorithms, such as these, were a significant improvement for anomaly detection when minimal feature engineering was used, and are therefore likely to be employed for large-scale collections of financial transaction data.

Deep learning-based algorithms also became very popular recently. Autoencoders, as unsupervised neural networks, were widely used because they learn to compress and reconstruct input data, thereby reducing the need for pattern-based behaviour modelling, as noted by Chen et al. [10]. In abnormal data cases, models record the maximum reconstruction error, thereby indicating suspicious behaviour. Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) models subsequently extended this capacity to generalise across temporal associations within sequences of data and, further, to forecast possible user behaviour and anomalies, as detailed by Choi et al. [11]. Even assuming such mechanisms were equipped with unmatched detection capabilities, the combination of large training datasets and intensive computations made them difficult to apply in real-time, as discussed by Chen et al. [12]. These limitations have also prompted studies into hybrid methods that

combine the speed and efficacy of conventional algorithms with the representational capacity of deep neural networks, which are currently the best financial compliance systems' anomaly-detection technologies available on the market.

3. Methodology

Our proposed AI-based anomaly detection approach utilises a hybrid technique that combines an Isolation Forest and a deep Autoencoder to provide an effective, efficient, real-time payment compliance solution. This is achieved through a multi-step pipeline that starts with ingestion and preprocessing. The raw transaction data, with features `transaction_amount`, `sender_account_risk_score`, `receiver_account_risk_score`, and `transaction_frequency_per_hour`, is handled first. Categorical features, such as `time_of_day`, are one-hot encoded to prepare the machine learning model. Numerical features are Min-Max normalised to the 0-1 range. This is a crucial step to prevent features with higher magnitudes from unbalancing the model's learning process, e.g., in the case of an Autoencoder, since it's scale-sensitive to input data. Now, the preprocessed data is fed to our framework's two primitive blocks simultaneously.

Isolation Forest is the first block and an ensemble classifier that works based on isolating anomalies. It builds many "isolation trees," where features are chosen randomly and split at random. Outliers are "few and different," and therefore, outliers can be easily identified; thus, they will have a reduced average path length from the tree root to a terminal node. This architecture is well-suited because it does not require density or distance calculations, is highly scalable for large data sets, and is therefore well-suited for a first-pass, fast filtering of transactions. The second component is a deep Autoencoder, an unsupervised symmetric bottleneck realisation of an encoder-decoder neural network. The encoder maps the input transaction data to a lower-dimensional representation in a latent space, learning a compressed representation of common transaction key features. The decoder attempts to reconstruct the original input from the compressed representation. The network is trained solely on normal, non-anomalistic transactions.

The implication is that the model's capacity to reconstruct a new transaction using the trained Autoencoder is measured by the input and output mean squared errors (MSEs) of the reconstructions. It is indicative of high reconstruction error if the transaction does not fall within the normal behaviour patterns the model was trained on, and thus marks it as most likely an outlier. Finally, the architecture averages the two model outputs as a single anomaly score. Both Autoencoder reconstruction error and Isolation Forest path length are normalised and averaged using a flexible weighted-average mechanism that can be trained on a validation set to optimise worst-case performance. The final prediction is derived by applying a calibrated threshold to the average score. Transaction parties whose thresholds exceeded the given limit were flagged as suspicious and sent to a compliance officer for review, but were cleared once the limit was met. The hybrid framework utilised Isolation Forest's capability to operate in a real-time application and Autoencoder's capability to extract deep patterns with high accuracy, thereby making the anti-financial crime defence robust and effective.




Figure 1: Hybrid AI framework for anomaly detection

Figure 1 presents the end-to-end hybrid AI process proposed for real-time anomaly detection. It begins at the left-hand side with the ingestion of Real-Time Transaction Data from the payments platform. Raw data, in the form of a list of transactional features, is piped to the Data Preprocessing module. Within this module, the following alterations are introduced: data are pre-cleaned to remove any missing values, category variables such as 'time_of_day' are converted to numeric type via one-hot

encoding, and numerical variables are standardised to a common scale (e.g., 0-1). Standardisation is needed in this case to improve the performance of subsequent machine learning algorithms. Pre-cleaned data are then fed into two parallel analysis branches. The Isolation Forest maximum stream sends the data to the Isolation Forest Model. The model then rapidly applies all transactions using the tree structure to generate an estimate of how easy or difficult it is to isolate a point from the other points. The lowest stream provides input to the Autoencoder Model, a deep neural network that trains on the transaction and returns a reconstruction error. High error indicates deviation from typical learned activity patterns. Both the models' results—Autoencoder reconstruction error and Isolation Forest score—are provided as input to the Hybrid Scoring Engine. The two measurements described earlier are combined to produce a single compounded anomaly score. It is then compared against a Predefined Threshold. Once it surpasses this threshold, it is classified as an Anomaly and an alert is raised and propagated for manual validation. Otherwise, it is marked Normal and allowed to proceed after passing the actual-time compliance test.

4. Data Description

Experimental data is a test set of financial transactions designed to mimic the operation of an existing real-time payment system. The dataset, "Synthetic Real-Time Payment Transaction Dataset," is created for this research to provide an uncontrolled environment for validating the anomaly detection system while ensuring that real financial data are not subject to privacy or security breaches. The dataset includes 432 data points, one per record. The sample is also biased to reflect actual circumstances, i.e., in about 90% of the transactions (389) are normal, and 10% (43) are unusual or spurious. All instances have six features: transaction_id (unique identifier of record), transaction_amount (amount of transaction), sender_account_risk_score (pre-calculated 0-100 risk score of sending account), receiver_account_risk_score (same for receiving account), transaction_frequency_per_hour (transaction sent by sender in last one hour), and time_of_day (categorical feature: 'Early Morning', 'Morning', 'Afternoon', 'Evening', 'Night'), and is_anomaly (binary target where label 1 is anomaly and label 0 is normal transaction).

5. Results

Our experimentally proven hybrid AI-based method yielded the highest-level positive findings, confirming its enhanced ability to identify anomalies in the synthetic real-time payment dataset. We experimented by dividing the 432-instance data set into a training set (80%, i.e., 345 instances) and a test set (20%, i.e., 87 instances), ensuring equal ratios of normal and anomalous classes in both sets. We then contrasted the performance of our Hybrid Model to that of its constituent parts—single Isolation Forest and Autoencoder—and a baseline, Logistic Regression. Our performance metrics were the standard classification metrics: Precision, Recall, F1-Score, and Area Under the Receiver Operating Characteristic Curve (AUC-ROC). Autoencoder Reconstruction Error is given by

$$L(x, \hat{x}) = \frac{1}{n} \sum_{i=1}^n \left(x_i - g_{\phi}(f_{\theta}(x_i)) \right)^2 \quad (1)$$

Table 1: Comparative analysis of model performance

Model Name	Precision	Recall	F1-Score	AUC-ROC
Logistic Regression	0.65	0.59	0.62	0.75
Isolation Forest	0.82	0.88	0.85	0.92
Autoencoder	0.92	0.84	0.88	0.94
Hybrid Model	0.94	0.91	0.92	0.97
(Placeholder)	0.00	0.00	0.00	0.00

Table 1: Comparative quantitative performance summary of the values computed by the four models which were tried in this paper over the test set. The numbers clearly show an improvement in performance across the range from the baseline model to our proposed hybrid system. The baseline Logistic Regression model performed the poorest on all the metrics, with an F1-score of 0.62, i.e., it was not able to combat the complexity of financial malpractices effectively. The Isolation Forest model performed best at simulating an actual performance gain, achieving high precision and recall with an extremely high F1-score of 0.85. It's not the best by any stretch of the imagination, but that's okay, because it is amazing as a fast one-stage outlier detector.

The autoencoder model had a high Precision (0.92), i.e., if it had labelled a transaction as anomalous, then it must have been correct. Its high Precision fully justified its cost by reducing false alarms from compliance teams, enabling further investigation. But its Recall was lower than Isolation Forest's, and it also detected some of the malicious transactions. The Hybrid Model, a blend of Isolation Forest and Autoencoder, performed best. It achieved all four metrics: Precision (0.94), Recall (0.91), F1-Score (0.92), and AUC-ROC (0.97). This is physical proof that the hybrid model accurately accounts for the inherent

deficiencies of its components, enabling it to build a system that is predominantly correct and complete in its identification. Isolation forest anomaly score function can be framed as:

$$s(x, \psi) = 2^{-\frac{E[h(x)]}{c(\psi)}} \quad (2)$$

Area under the ROC curve is:

$$AUC = \int_0^1 TPR(FPR^{-1}(t))dt = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I(x_p > x_n)P(x_p)P(x_n)dx_pdx_n \quad (3)$$

SHAP value formula

$$\phi_i(f, x) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(|N|-|S|-1)!}{|N|!} [f_x(S \cup \{i\}) - f_x(S)] \quad (4)$$




Figure 2: Determination of hybrid model classification of 87 test cases

Figure 2 illustrates the Hybrid Model's classification accuracy on the test data graphically. Each point in the plot represents a single transaction, plotted against the sender account's risk score and value. Every point is superimposed on the model's actual prediction. Blue, the densest point, is a typical transaction. They are located in the lower left of the plot, representing low-value transactions and low sender scores, which would be typical of legitimate financial activity. Red dots, or abnormal transactions as they were identified, are scattered in most places but reflect patterns of differentiation. There is also close clustering in the top-right quadrant, indicating that the model picks high-value, high-risk transactions from risk-score-high current accounts.

The model also picks transactions which are not necessarily large in value but unexpected. For instance, there are several red dots in the top-left region, indicating low- to moderate-value transactions from high-risk accounts. This means the model can detect nuanced signals of risk beyond strict monetary limits, a desirable advantage of rule-based models. The distinct discrimination between the red and blue clusters, with zero overlap or zero assistance, visually supports our high F1-Score and accuracy, indicating how well the model is performing in establishing the best decision boundary to distinguish between valid and suspicious activity. Weighted hybrid scoring function

$$S_{\text{final}}(x) = w \cdot \frac{L(x, \hat{x}) - \mu_{AE}}{\sigma_{AE}} + (1 - w) \cdot (1 - s(x, \psi)) \quad (5)$$

Table 2: Feature importance analysis for hybrid model

Feature Name	SHAP Value (Mean)	Permutation Importance	Gradient Importance	Combined Rank
sender account risk score	0.48	0.51	0.45	1
transaction amount	0.35	0.31	0.38	2

transaction frequency per hour	0.29	0.26	0.30	3
time of day	0.11	0.09	0.13	4
(Placeholder)	0.00	0.00	0.00	0

Table 2 presents the Hybrid Model feature importance analysis results, which estimate the relative importance of each input feature in predicting the outcome. Three methods—Mean SHAP Value, Permutation Importance, and Gradient Importance—have been employed to provide a comprehensive, multidimensional output, which is summarised in an ultimate 'Combined Rank'. Feature importance analysis ranks sender_account_risk_score as the top feature across all methods. This is required in a manner that maintains the integrity of account history, behaviour and risk profiling within real-time detection systems. The model also recognised that sender identity is a strong predictor of anomaly behaviour. Transaction amount is the second most important feature, reflecting the standard-compliance rationale that larger amounts entail higher risk. But because it has a lower priority rank than the sender risk profile, what they are sending is less important than who they are sending to. transaction_frequency_per_hour ranked third and, once more, emphasises that higher usage is a very significant behavioural red flag that the model can readily detect.




Figure 3: Anomaly score predicted by hybrid model in continuous space for transaction amount (x-axis), time of day (y-axis, numerically encoded), and predicted anomaly score (z-axis)

Figure 3 is a 3D mesh plot of the Hybrid Model's scoring engine output, showing yet another representation of its decision space. The z-axis displays the forecast anomaly score, and denser points (in lighter hues, such as yellow and red) indicate a higher likelihood that a transaction is anomalous. The x-axis and y-axis are two of the strongest input features: the transaction amount and the hour of the day, respectively. The graph shows the weak, non-linear relationships the model has learned. Rather than a single, flat plane of options, there is a rolling one with highs and lows.

For example, there is a distinct ridge in high transaction volumes at all times of day, which is a clear indication of risk. The chart also captures more subtle trends. There is also an evident upper limit somewhere in the area for medium-sized, non-extreme-time transactions (e.g., 'Early Morning', low y-axis value). This indicates that the model can identify such activity as unusual and even trigger suspicion when the dollar amount is not extreme. The surface is flat and low (bluey cool colours) across the range of small transaction sizes on normal business days ('Morning', 'Afternoon'), as in run-of-the-day transactions. This graph illustrates the significant quality improvement achieved by a superior AI-based method, which eschews clunky linear cutoffs to produce a high-level, context-sensitive risk choice that is highly responsive to interactions across various transactional dimensions.

The Baseline Logistic Regression model extrapolated to identify the complex, non-linear pattern of the anomalies, achieving an F1-Score of 0.62 and an AUC-ROC of 0.75. It is a result of the constraints imposed by common models in the problem domain. The Isolation Forest model jumped over quantum-scale barriers by identifying outliers through differential differences. The F1-Score was 0.85 and the AUC-ROC was 0.92, which was satisfactory for accelerated detection. One Autoencoder model,

which accounted for the intricate structure of normal data, also performed better, achieving an accuracy of 0.88, an F1-score of 0.88, and an AUC-ROC of 0.94. Its advantage is that it can identify extremely small deviations from the mean that other methods can't, though it took slightly longer to calculate than the Isolation Forest method. Our best performer was the suggested Hybrid Model, which combined the Isolation Forest output with the Autoencoder.

Its hybrid power performed best with all the indicators. The Hybrid Model achieved an F1-Score of 0.92, Precision of 0.94, and Recall of 0.91. It also yielded an AUC-ROC of 0.97, indicating a very high ability to separate anomalous from normal classes. This validates our first hypothesis that combining the deep pattern discovery of the autoencoder with the strong outlier detection of the Isolation Forest yields a more balanced system. The Hybrid Model would reduce Isolation Forest false positives and increase the Autoencoder's true positive rate (Recall) at an appropriate ratio. The graphical representation of the results in the following Figures also indicates the model's decision-making and discrimination ability among transaction classes according to the KRI. Quantitative results, as illustrated in the following Tables, provide a clear delineation of the comparison's performance characteristics and the model's predictive level.

6. Discussion

The results of this work overwhelmingly support the operation of a hybrid AI system for real-time payment system compliance. The ensuing extension of the Hybrid Model, as shown in Table 1, is not incremental but revolutionary relative to isolated machine learning algorithms and plain Logistic Regression. It is largely the complementarity of complementary strengths that sustains it. Isolation Forest employs its tree-partitioning depth as a rough filter, effectively removing flashy outliers and suspect transactions at nearly zero computational cost. That satisfies the first criterion of being fast in an RPS environment. The autoencoder conducts a very mushy and evasive search. In the right transaction pattern training, it excels at identifying those subtle, creeping outliers that are not far away but are certainly trending away from established norms. The aggregate anomaly score of the two models together is a more stable, better risk measure than either model alone. The data in the plots also says something about what the model is doing. The scatter plot in Figure 2 indicates that the model is not using a linear separator.

It consistently identifies a wide range of anomalies, from high-level fraud (involving high-dollar amounts on high-risk accounts) to more challenging-to-detect instances (low-dollar amounts on high-risk accounts). Its context-sensitive detection is revolutionary within the rule-based paradigm, meeting most expectations for identifying high-value transactions above a selected dollar cutoff. Figure 3: A 3D mesh plot supports the above observation by showing a non-linear, complex decision surface. It's a combination of high-risk features and not high-risk individual features that received high-risk scores, i.e., abnormal time transactions of moderate value. It is evidence that the model had learned rules of behaviour rather than hard-and-fast rules, and was worse at detecting new and novel typologies of fraud. And as an additional bonus, the feature importance in Table 2 is a revolution for the banks as well.

As `sender_account_risk_score` is one of the best predictors of importance, it's another example of necessity being the mother of invention, the demise of transaction-level monitoring, and the advent of an even more customer-centric risk methodology. Double compliance is inevitable in all instances; it must be carried out in accordance with good customer due diligence and off-the-shelf risk-profiling contracts. Comparative behaviour, with typical values like `transaction_frequency_per_hour`, also allows for migration to behaviour-based, dynamic monitoring. Spur-result-eliciting results aside, the weaknesses of this study should be reduced to a bare minimum. A non-overridable override fault, essentially, operates on a compiled set of data with sparse observations (432).

Any data that can be derived, however large it is, can never, even in any way at all, hope to equal the randomness, size, and volume of active data of real financial data. Its accuracy has to be checked and verified against large actual bank transaction data. Additionally, although the architecture is visually appealing for the framework, it may be optimised to be more computationally efficient and scalable to millions of transactions per second in production. The model is also a "black box". However, we may be able to infer feature importances; why we do this every time need not be explainable, and this can be controlled by applying Explainable AI (XAI) methods.

7. Conclusion

You can research, develop, and test an AI-hybrid anomaly detection application to address compliance issues in Real-Time Payment Systems. With robust Autoencoder balancing and an Isolation Forest, our model is well-positioned to meet doubled requirements and achieve deep analysis comprehension. Experimental success on a home-made synthetic test set categorically validates the composite methodology advantage. Our model outperformed individual stand-alone algorithms and even baseline models to historic proportions, achieving an F1-score of 0.92 and an AUC-ROC of 0.97, successfully segregating legitimate and suspicious transactions and thereby improving overall credibility. Our results — both tabular and graphical — demonstrate that the model reveals optimal hidden, non-linear interactions and strikes a balance between an adequate number of added risk

factors. Our results led us to consider context, such as sender's risk score history and exchange frequency, as more indicative of illegal behaviour than exchange value itself.

This corresponds to the policy that the most powerful compliance strategy must move away from behaviour-invariant threshold controls to adaptive behaviour-invariant analytics. Pointing to Figures 2 and 3 provides the model with its high-end advantage of discretion, enabling it to react to delicate sets of risk indicators. Lastly, this paper de-scales an ideal solution of unimaginable capability, efficacy, and scalability, allowing banks and financial institutions to further tailor their AML and fraud detection. Through the power of collaborative hybrid AI, business organisations will be best positioned to safeguard their payment channels, reduce their business fraud-related costs, and prepare for regulation when real-time finance becomes a reality.

Acknowledgement: This work represents the author's independent effort, developed entirely through personal analysis, interpretation, and writing.

Data Availability Statement: The data supporting this research pertain to AI-driven anomaly detection frameworks for real-time payment system compliance and were obtained from publicly accessible, verified sources.

Funding Statement: No external funding or institutional assistance was received in the preparation of this manuscript or the associated research work.

Conflicts of Interest Statement: The author declares that there are no conflicts of interest regarding the research, analysis, or publication of this paper. All references have been appropriately cited.

Ethics and Consent Statement: All ethical considerations were duly observed, with proper consent acquired from relevant organisations and individuals involved in the data collection process.

References

1. S. Ahmad, A. Lavin, S. Purdy, and Z. Agha, "Unsupervised real-time anomaly detection for streaming data," *Neurocomputing*, vol. 262, no. 11, pp. 134–147, 2017.
2. O. Alghushairy, R. Alsini, T. Soule, and X. Ma, "A review of local outlier factor algorithms for outlier detection in big data streams," *Big Data Cogn. Comput.*, vol. 5, no. 1, pp. 1–24, 2020.
3. S. A. Ali, "Anomaly detection in telecommunication networks: Leveraging novel big data and machine learning techniques for proactive fault management," *Educ. Adm. Theory Pract.*, vol. 30, no. 5, pp. 5751–5770, 2024.
4. W. A. Ali, K. N. Manasa, M. Bendeache, M. Fadhel-Aljunaid, and P. Sandhya, "A review of current machine learning approaches for anomaly detection in network traffic," *J. Telecommun. Digit. Econ.*, vol. 8, no. 4, pp. 64–95, 2020.
5. Á. Arcos-García, J. A. Alvarez-Garcia, and L. M. Soria-Morillo, "Deep neural network for traffic sign recognition systems: An analysis of spatial transformers and stochastic optimisation methods," *Neural Netw.*, vol. 99, no. 1, pp. 158–165, 2018.
6. S. Banik, S. K. Saha, T. Banik, and S. M. Hossain, "Anomaly detection techniques in smart grid systems: A review," in *Proc. 2023 IEEE World AI IoT Congr. (AIoT)*, Seattle, Washington, United States of America, 2023.
7. M. A. Bashar and R. Nayak, "TAnoGAN: Time series anomaly detection with generative adversarial networks," in *Proc. 2020 IEEE Symp. Series Comput. Intell. (SSCI)*, Canberra, Australia, 2020.
8. P. Biswas and T. Samanta, "Anomaly detection using ensemble random forest in wireless sensor network," *Int. J. Inf. Technol.*, vol. 13, no. 5, pp. 2043–2052, 2021.
9. R. Boutaba, M. A. Salahuddin, N. Limam, S. Ayoubi, N. Shahriar, F. Estrada-Solano, and O. M. Caicedo, "A comprehensive survey on machine learning for networking: Evolution, applications and research opportunities," *J. Internet Serv. Appl.*, vol. 9, no. 1, pp. 1–99, 2018.
10. R. Q. Chen, G. H. Shi, W. L. Zhao, and C. H. Liang, "A joint model for IT operation series prediction and anomaly detection," *Neurocomputing*, vol. 448, no. 8, pp. 130–139, 2021.
11. K. Choi, J. Yi, C. Park, and S. Yoon, "Deep learning for anomaly detection in time-series data: Review, analysis, and guidelines," *IEEE Access*, vol. 9, no. 8, pp. 120043–120065, 2021.
12. H. Chen, P. Chen, B. Wang, X. Yu, X. Chen, D. Ma, and Z. Zheng, "Graph neural network based robust anomaly detection at service level in SDN driven microservice system," *Comput. Netw.*, vol. 239, no. 2, p. 110135, 2024.