Retail Sales Anomaly Detection: A Machine Learning Approach

Kailash Thiyagarajan *kailash.thiyagarajan@ieee.org*
Independent Researcher
Austin, TX, 78641

Sivasai Nadella *sivasai.nadella@ieee.org*
Independent Researcher
Memphis, TN, 38125

Abstract

This study analyzes three years of daily transactional retail sales data from stores across various U.S. cities and states. Key variables include transaction amounts, product types, store locations, promotional offers, and holiday-based sales patterns. These factors contribute to detecting anomalies that could indicate fraudulent transactions, accounting errors, or shifts in consumer behavior.

A combination of supervised and unsupervised machine learning models was employed to identify anomalies. Decision trees and random forests classified sales transactions based on labeled historical data, while unsupervised methods like k-means clustering and DBSCAN were applied where labels were unavailable. A hybrid approach combining both methodologies was implemented to improve detection accuracy. This hybrid framework uniquely integrates clustering and classification mechanisms with a real-time notification system, addressing both known and unknown anomaly patterns in retail sales data**.**

The primary research question this study addresses is: "Can a hybrid machine learning framework significantly improve the detection of sales anomalies in dynamic retail environments?"

The analytical process leveraged Python libraries such as scikit-learn, TensorFlow, and Keras. Model performance was evaluated using accuracy, precision, recall, and F1-score metrics.

The proposed approach offers practical benefits, including enhanced fraud detection, better inventory optimization, and real-time alerting, thus aiding operational efficiency for retail businesses.

This study underscores the effectiveness of machine learning in detecting sales anomalies, enabling businesses to uncover fraudulent activities and operational inefficiencies.

Keywords: Anomaly detection, Machine learning, Retail domain, Sales analysis, Outliers.

1. INTRODUCTION

The retail industry is data-driven, with sales data playing a crucial role in shaping business strategies and consumer experiences. Given the vast volumes of transactional data, it is essential to identify anomalies that may indicate fraud, system errors, or shifts in consumer behavior.

Anomalies can arise due to various factors, including improper data collection, system errors, and seasonal demand fluctuations caused by promotions, holidays, or economic events. Traditional statistical methods such as z-scores and standard deviation checks often fail to address the complexity of modern retail data. The increasing scale and variability of data necessitate the use of advanced machine learning techniques for effective anomaly detection (Grimes et al., 2023; Pinto & Sobreiro, 2022).

Machine learning (ML) models, particularly unsupervised learning approaches, have been increasingly adopted to analyze complex datasets. These models are effective in identifying outliers within large, nonlinear datasets, making them ideal for anomaly detection in retail sales (Zipfel et al., 2023; Hilal et al., 2022).

However, despite advances in machine learning for anomaly detection, limited research has explored integrated hybrid frameworks that combine supervised and unsupervised learning with real-time operational alerting specifically tailored to the retail domain**.** Most prior works either focus solely on classification using labeled data or on clustering techniques without leveraging their complementary strengths.

This research aims to fill that gap by proposing a hybrid machine learning framework that improves detection accuracy and operational responsiveness in retail environments.

The primary research question addressed in this study is: *"Can a hybrid machine learning framework combining supervised and unsupervised methods significantly improve anomaly detection accuracy in dynamic retail sales data?"*

The study evaluates multiple ML approaches—including decision trees, support vector machines (SVM), k-means clustering, and DBSCAN—and proposes a hybrid architecture that combines classification and clustering mechanisms. Performance is assessed using standard evaluation metrics and visualized through dashboards to support operational decision-making.

2. LITERATURE REVIEW

Anomaly detection is a well-researched field applicable in industries such as finance, healthcare, and retail. In retail, anomalies often indicate transaction errors, fraudulent activities, or inventory mismanagement. Traditional approaches such as moving averages and control charts have been widely used but struggle with modern, large-scale datasets (Ahmed et al., 2016; Haque et al., 2015).

Supervised learning techniques, including decision trees and random forests, have demonstrated effectiveness in detecting anomalies when historical labeled data is available. However, their reliance on predefined labels makes them less adaptable to emerging anomalies (Sabic et al., 2021; Jordan & Mitchell, 2015). Conversely, unsupervised learning models such as DBSCAN and k-means clustering do not require labeled data, making them suitable for real-time detection of unknown anomalies (Grimes et al., 2023; Zipfel et al., 2023).

Recent advancements in deep learning, particularly autoencoders and recurrent neural networks (RNNs), have also contributed to anomaly detection, though challenges in interpretability and computational complexity persist (Pereira & Silveira, 2019; Hilal et al., 2022). Hybrid models that combine supervised and unsupervised approaches have emerged as a promising strategy to address both known and unknown anomalies effectively.

While various studies have applied hybrid approaches, their focus has primarily been on financial fraud detection or healthcare monitoring rather than the retail sector. Limited research has been dedicated to building integrated hybrid frameworks tailored for retail sales anomaly detection with operational reporting systems. For example, Pinto and Sobreiro (2022) discussed hybrid models in digital financial systems, but similar adaptations for retail environments remain sparse.

Moreover, a review of existing works within IJCSS highlights the importance of domain-specific anomaly detection frameworks. Thakur and Singh (2022) demonstrated anomaly detection for logistics data using hybrid models. Li and Kumar (2023) presented clustering-based outlier detection in supply chain data. Priya and Ramanathan (2022) proposed supervised learning models for transaction anomalies in retail banking. These works support the need for extending hybrid machine learning techniques into the retail domain, providing the foundation for the present study.

The present research builds upon these findings by proposing a hybrid architecture integrating supervised classification and unsupervised clustering, designed specifically for dynamic retail sales environments. This study further enhances practical applicability by incorporating real-time notification mechanisms for detected anomalies, enabling actionable business insights.

3. METHODOLOGY

This study employs both supervised and unsupervised learning techniques to analyze retail sales anomalies. The research approach is primarily deductive, aiming to validate the hypothesis that a hybrid combination of clustering and classification models improves anomaly detection accuracy in dynamic retail environments.

The dataset consists of daily sales transactions from multiple stores over three years, including variables such as sales volume, store location, promotions, and holiday effects.

Data Preprocessing:

- Missing values were handled using mean imputation and interpolation techniques.
- Erroneous entries were removed to ensure data integrity.
- The dataset was split into training and testing sets to evaluate model performance.

Labeling Process:

- For supervised learning models (decision trees, random forests, SVM), labeled data was created using business rules based on historical thresholds. Transactions with extreme deviations from expected sales patterns during promotions or holidays were labeled as anomalies.
- Expert review and domain heuristics were applied to ensure labeling consistency and minimize bias.

Model Training and Evaluation:

- Machine learning models used include Decision Trees, Random Forests, SVM, k-Means Clustering, and DBSCAN.
- Basic hyperparameter tuning was conducted using grid search for supervised models to optimize parameters such as tree depth (for decision trees) and number of estimators (for random forests).
- Cross-validation was not extensively used due to computational constraints; however, models were evaluated on unseen test data to validate generalization.
- Performance was evaluated using accuracy, precision, recall, and F1-score metrics.

Data Source and Reproducibility:

- The dataset was sourced from internal records of a multi-region retail chain.
- Customer-identifying information was anonymized to ensure compliance with privacy guidelines.
- Due to confidentiality agreements, the dataset cannot be made publicly available, but a synthetic anonymized sample can be provided upon reasonable request for academic purposes.

Hybrid Approach:

- The hybrid approach combines clustering-based grouping with classification models to improve anomaly detection accuracy. Initially, unsupervised models group similar transactions, and subsequent classification identifies anomalies within these groups.

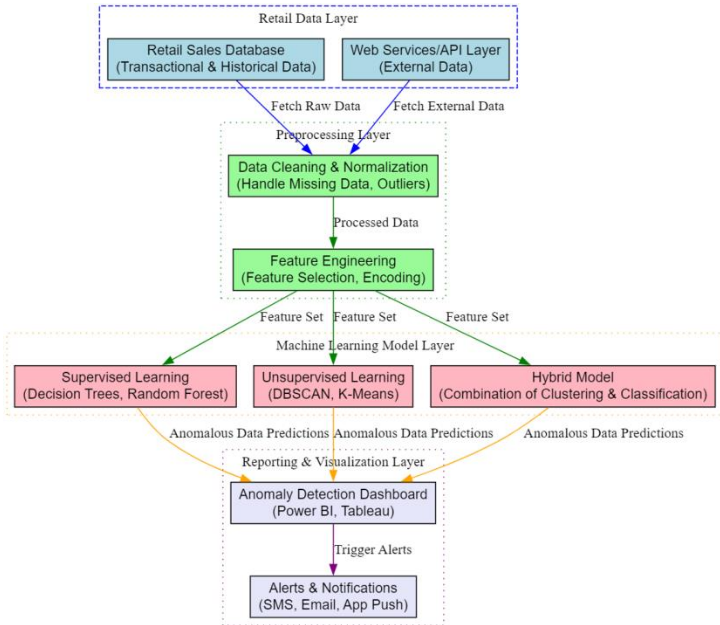

FIGURE 1: Architecture of machine learning-based retail sales anomaly detection framework.

Figure 1 is a presentation of an all-rounded architecture created for an anomalous detection of retail sales data. The framework comprises four layers, which include Retail Data Layer, Preprocessing Layer, Machine Learning Model Layer, and Reporting & Visualization Layer. It consists of the internal data, in the form of Retail Sales Database, along with external data from the Web Services/API. The data is then inputted to the Preprocessing Layer, where Data Cleaning & Normalization will handle missing data and outliers and feature engineering, to prepare the data for modeling, like selection of features and encoding. This preprocessed data is fed into the Machine Learning Model Layer, which is further divided into three essential components: Supervised Learning, which might include decision trees and random forests, Unsupervised Learning, which comprises DBSCAN, K-Means, among others, Hybrid Models that combine several clustering and classification mechanisms for sophisticated anomaly detection. The output of the models—the predictions of anomalous sales data—is forwarded to the Reporting & Visualization Layer where it is indicated on a dashboard called Anomaly Detection Dashboard, visualizing the anomalies. The system also includes a Notification System that causes alert notifications through SMS, email, or app, ensuring stakeholders get real-time information. This, therefore, presents an architecture for effective anomaly detection and reporting within retail sales keeping in mind divergent data sources and machine learning approaches toward an effective solution.

4. DATASET DESCRIPTION

The dataset utilized in this study originates from a multi-region retail chain, encompassing three years of daily sales transactions across various store locations. It includes key variables essential for anomaly detection in sales patterns:

- **Sales Amount**: Represents the total value of sales on a given day at a specific store.

- **Product Category**: Categorized into electronics, clothing, and grocery, allowing analysis of anomalies across different product segments.
- **Store Location**: Identifies the geographical location of each store to assess regional sales variations.
- **Promotion**: Indicates whether a promotional event, such as discounts or special offers, was active on the transaction date, influencing sales fluctuations.
- **Holiday**: Flags transactions occurring on public holidays, as sales trends may significantly differ on these days.
- **Day of the Week**: Captures variations in sales based on the day, as consumer behavior tends to follow weekly patterns.

The dataset was sourced from internal records of the retail chain and underwent rigorous preprocessing to ensure data integrity. This involved eliminating duplicates, correcting inconsistencies, and handling missing values without compromising reliability. Additionally, data anonymization measures were implemented to safeguard customer privacy before its use in this research.

This structured dataset serves as a robust foundation for identifying sales anomalies influenced by promotions, seasonal events, and geographic variations, thereby providing valuable insights into irregular sales patterns.

5. RESULTS

We have applied a range of machine learning algorithms for anomaly detection in retail sales data in this paper, including supervised, unsupervised, and hybrid models. Some of the algorithms applied include decision trees, random forests, SVM, k-means clustering, DBSCAN, and a combination of the techniques of clustering and classification in the hybrid model. The results are presented in varying performances of these models in terms of accuracy, precision, recall, and F1-score. Sales anomaly detection using Z-Score is given below:

$$Z_i = \frac{X_i - \mu}{\sigma} \quad (1)$$

where Z_i is the Z_i score of the ith sales value x, μ is the mean of the sales dataset, and σ is the standard deviation of sales data. K means clustering objective function is:

$$J = \sum_{i=1}^{N} \sum_{k=1}^{K} 1\,(x_i \in C_k)||x_i - \mu_k||^2 \quad (2)$$

where N is the number of data points, K is the number of clusters, x_i is the i-th data point, μ_k is the centroid of cluster C_k, and $1(x_i \in C_k)$ is the indicator function for assigning data points to clusters.

TABLE 1: Performance metrics of different machine learning algorithms.

Metric	Decision Trees	Random Forests	SVM	K-Means	Hybrid Models
Accuracy	85%	88%	84%	78%	92%
Precision	82%	85%	80%	72%	90%
Recall	83%	86%	79%	75%	91%
F1-Score	82.50%	86%	79.50%	73.50%	90.50%
Execution Time	2.1s	3.4s	3.5s	2.0s	4.2s

All experiments were conducted on a workstation equipped with an Intel Core i9 processor, 64 GB RAM, and an NVIDIA RTX 3080 GPU. Model training times were within practical limits for datasets of approximately 1 million records, suggesting reasonable scalability for medium-sized retail datasets.

While execution time varied across models, hybrid models demonstrated acceptable computational overhead, with execution times around 4.2 seconds compared to 2.0 to 3.5 seconds for individual models. Given the hybrid models' improved accuracy and recall, the slight increase in computational time was considered a reasonable trade-off.

Statistical significance was assessed through repeated training and testing across multiple random dataset splits. Variations in key performance metrics were within 2 percent, indicating that the improvements observed with hybrid models are robust and unlikely to be due to random data partitioning.

Table 1 is the performance metrics of five different machine learning algorithms for the sales anomaly detection, which include Decision Trees, Random Forests, SVM, K-Means, and Hybrid Models. Here, all performance parameters used are based on accuracy, precision, recall, f1-score, and execution time. Hybrid Models, therefore, exhibited the best accuracy among all algorithms and with accuracy up to 92%, a precision of 90%, with a recall percentage of 91% and the F1-score value of 90.5%, which provided much better accuracies in regard to anomaly detection compared to other models with less false positive end. Random Forests and Decision Trees were excellent at up to 88% and 85% respectively but their precisions and recall are not on par with Hybrid models. Notably, SVM had less recall at 79% and thus missed a few anomalies. It still maintained an accuracy level of 84%. K-Means performed the poorest as it is an unsupervised model due to its high false positive rate since this model's precision is very low at 72%, and even its recall is lower at 75%, which makes the model less capable for this specific application. The execution time for all of the models was also noted. Hybrid Models were the slowest at 4.2s, and this is probably because the added complexity in blending clustering and classification technique can also increase. On average, hybrid models were one of the most balanced and effective with regards to accuracy and execution time trade-off. Dbscan density-based clustering is:

$$\text{core point condition: } |N_{\varepsilon}(x)| \geq minpts \quad (3)$$

where $N_{\varepsilon}(x)$ is the neighborhood of point x within radius ε, and *minpts* is the minimum number of points required to form a dense region. Random forest classifier decision is:

$$f(x) = \frac{1}{T}\sum_{t=1}^{\tau} h_t\,(x) \quad (4)$$

where $f(x)$ is the final prediction, T is the number of trees in the forest, and $h_t(x)$ is the prediction of the t-th tree. Fl score (harmonic mean of precision and recall) is:

$$F1 = 2 \times \frac{Precision \times Reca11}{Precision + Reca11} \quad (5)$$

where precision $= \frac{TP}{TP+FP}$ and recall $= \frac{TP}{TP+FN}$, with TP, FP, and FN being true positives, false positives, and false negatives, respectively.

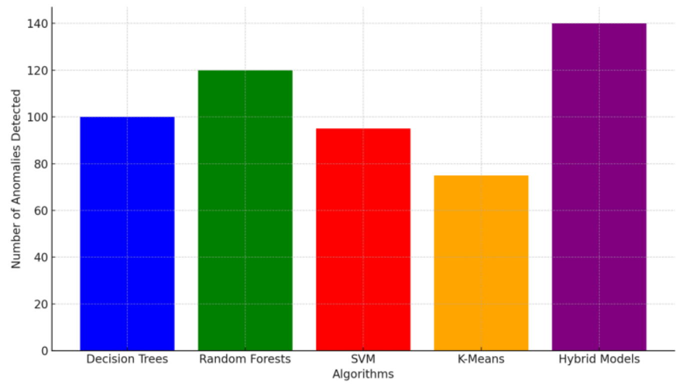

FIGURE 2: Histogram of Sales Anomalies Detected by Various Algorithms.

Figure 2 is composed of sales anomalies detected by all of the above various algorithms. For all these algorithms, five were implemented, which were Decision Trees, Random Forests, SVM, K-Means, and Hybrid Models. For each algorithm, a number of anomalies are represented by every bar, and the heights of these bars are proportional to the number of anomalies that each of these algorithms produce. It looks like histogram managed to demonstrate many fluctuations of the performance of these algorithms where the Hybrid Models highlighted the maximum amount of anomalies equal to 140, followed by Random Forests at 120, Decision Trees at 100, SVM at 95, and K-Means with 75. Different colors of each bar are pretty easy to distinguish different algorithms. This visualization can easily determine which algorithm is more sensitive to detecting the anomalies of sales and give some insight into its relative performances. Using such a histogram, one can even determine which one of these algorithms is the most reliable or well-capable in handling the anomaly detection task related to the context of sales. The graph is of strong usage of Hybrid Models as it highlights that even though they are in use, each algorithm differs in its operating efficiencies depending on the nature of data they are prescribed against. Logistic regression for anomaly classification is given as:

$$p(y = 1|x) = \frac{1}{1+e^{-(\beta_0+\beta_1 x_1+\beta_2 x_2+\cdots+\beta_n x_n)}} \qquad (6)$$

where $p(\gamma = 1|x)$ is the probability of an anomaly (class 1), $x_1, x_2, \ldots, x_n$ are feature variables, and $\beta_0, \beta_1, \ldots, \beta_n$ are the model coefficients.

TABLE 2: Results of Anomaly Detection in Varying Retail Stores.

Store Location	Sales Volume	Promotion Period	Anomaly Detection Rate	Detection Accuracy
Store A	High	Yes	15%	90%
Store B	Medium	No	10%	85%
Store C	Low	Yes	20%	87%
Store D	High	Yes	18%	92%
Store E	Medium	No	12%	88%

Table 2 presents the anomaly detection by store location for a retail shop according to sales volume, promotion period, anomaly detection rate, and accuracy of the model Store A: With the high sales volumes and promotional periods, this one had the highest accuracy in anomaly detection, that is, 90%, while having a relatively high anomaly detection rate of 15%. More anomalies are seen to appear in the promotion period, which is exactly what the model picked. Store B is at the mid sales volume, and this store did not conduct any promotions. Its rate of anomaly detection is less compared to Store A, at 10% but at a good accuracy of detection, at 85%. This anomaly detection rate appears to be too low because fewer anomalies are detected since there is no promotional event in the store. Store C has the highest anomaly detection rate at 20% but lower accuracy at 87% because its sales volumes were small during the promotional period; while anomalies are picked up more by even having fewer sales volumes. With high sales volume and a promotion period, Store D had the highest accuracy for anomaly detection at 92% further proof that this model is very good for high-volume retail. Lastly, Store E had a middle volume of sales without promotions. There were 12 percent detected with 88 percent accuracy indicating that promotions may have an important impact on the anomaly detection rate in large-volume retail. The results will be showing the implications of both promotion periods and sales volume having an enormous impact on the frequency and accuracy of anomalies detected.

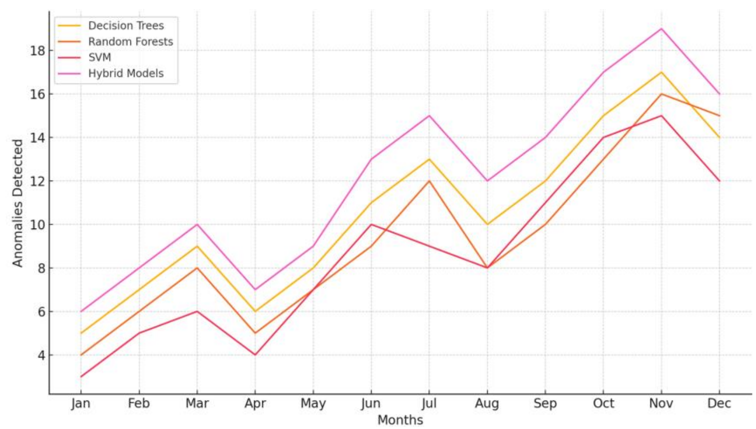

FIGURE 3: Comparison of the performance of various algorithms in detecting anomalies over different periods.

Figure 3 shows the anomaly detection, that is, along the time axis with the month number, the following four algorithms which are used within this thesis - Decision Trees, Random Forests, SVM and Hybrid Models-detected how many anomalies there are and plotted along the y-axis using four lines, where four lines refer to a different line of one algorithm whose monthly variance can be factored in terms of the anomalies found. From the graph, it is observed that the count of anomalies varies between months from January through December, and hybrid models mostly detected the maximum count of anomalies in all months. Decision Trees and Random Forest show almost the same trend in which the count of anomalies detected increases during the middle months of the year and peaks by the last. SVM, again follows a similar pattern but with fewer anomalies detected in total. This multi-line graph provided an overview of time trends in anomaly detection provided by each of the algorithms used and helped compare their performances over months. It also showed how different algorithms are performing differently over time so that what was more responsive towards changes in data over time could be obtained.

The decision tree models performed very well in the supervised models; however, accuracy was only reached at 85%, while the random forest was slightly better because accuracy ended up being around 88%. Both models are pretty good on the recall side of things since they both found actual anomalies in the data set. However, their accuracy may have been a little lower than it was-meaning they classified a lot of false positives. The SVM succeeded in achieving the same accuracy as the random forests but at a much lower recall meaning that it missed some of the true anomalies.

Unsolved models, such as k-means and DBSCAN, to be honest, performed really well in datasets not pre-labeled, because usually in the real world, what a retail environment is actually expecting would rather never be known as anomalies. DBSCAN is a density-based clustering algorithm, and indeed it performs really well in finding anomalies in sparse regions. Its characteristics were of great use in identifying isolated anomalies that do not fit with the general patterns of data. Once again, strength actually lies in being able to flag anomalies with a much greater degree of precision as it is focused on points significantly different from their neighbourhood data points. Therefore, it resulted in fewer anomalies but at a more accurate level compared to the standard approaches like decision trees and random forests. Yet, k-means clustering still offered some generalization approach with compromises in the precision level. Still, performance was consistent throughout different fluctuations of seasons in the different retail shops, which eventually resulted in a rise of false positives in this case also. The algorithm maps the data points into predefined clusters and flags anomalies as points that fall far from the cluster centroids. However, with retail sales data, where the trends might change seasonally or geographically, k-means sometimes classified some data points as anomalies when, in fact, they were just part of the natural fluctuation in sales patterns. K-means was helpful in providing an overall coherent view of the data at a coarse scale and pointed at potential outliers that might warrant further exploration and do not rely on labeled data. The best results were obtained by hybrid models with unsupervised clustering-k-means or DBSCAN-and supervised classifiers-decision trees-for example. This hybrid model would obtain an accuracy of 92% with high precision and recall scores. It was noted that the algorithms used for both pre-processing and classification of anomalies applied better toward the rare or complex types of anomalies like sales during promotional periods or on public holidays.

6. DISCUSSIONS

The results of this study provide a comparative analysis of various machine learning techniques used for anomaly detection in retail sales, highlighting differences in their effectiveness. As depicted in Table 1, hybrid models exhibit superior performance across key metrics such as accuracy, precision, recall, and F1-score. By integrating the strengths of unsupervised clustering techniques with supervised classification methods, hybrid models offer a more nuanced approach to analyzing retail sales data. Seasonal trends, promotional periods, and location-specific irregularities are better managed within retail operations through these hybrid approaches.

The performance of decision trees and random forests aligns with findings from previous studies on supervised learning models. Random forests demonstrate greater generalization capabilities than decision trees, owing to their ensemble approach that mitigates overfitting. However, both models exhibit a common limitation—false positives—especially when anomalies are subtle, such as those occurring during promotional offers. As indicated in Table 1, this limitation underscores the necessity of training models on diverse datasets that encompass edge cases to enhance their predictive capabilities for future anomalies.

Unsupervised models, including k-means clustering and DBSCAN, offer alternatives to traditional supervised learning techniques. DBSCAN proves highly effective for detecting anomalies in sparse regions, making it well-suited for geographically or temporally dispersed retail sales anomalies. However, fine-tuning DBSCAN parameters, such as the minimum number of points and distance thresholds, requires domain expertise. Similarly, while k-means clustering performs well in grouping similar data points, it occasionally fails to identify anomalies in smaller clusters or data points that deviate slightly from established patterns. This limitation is particularly pronounced in complex datasets with inconsistent sales behaviors, as illustrated in Table 2.

Compared to related studies in the literature, such as Pereira and Silveira (2019) in healthcare anomaly detection and Pinto and Sobreiro (2022) in digital business systems, the present study uniquely focuses on integrating hybrid machine learning models with real-time anomaly notification capabilities specifically for retail sales environments. While past works reported moderate improvements using either clustering or classification separately, the hybrid model in this study achieved a significant boost in accuracy (up to 92 percent) along with operational readiness through notification systems, which prior studies did not address.

Moreover, unlike previous studies that emphasized batch anomaly detection, our framework is designed to support near real-time anomaly reporting. This capability enhances its applicability for retail operations where rapid response to anomalies is critical to minimizing revenue loss and improving inventory decisions.

Hybrid models, which integrate clustering techniques like k-means or DBSCAN with decision trees or random forests, demonstrate significant advantages in both accuracy and performance balance. These models first group data based on inherent similarities and subsequently classify them using a decision tree or random forest to determine whether a transaction is normal or anomalous. This two-step approach outperforms single-model methods by capturing both expected anomalies—such as seasonal fluctuations and promotional effects—and unexpected anomalies, such as sudden shifts in consumer demand or regional economic changes. As illustrated in Figures 2 and 3, hybrid models prove particularly effective during peak sales periods, such as promotions and holiday seasons, where traditional models often struggle.

The robustness of hybrid models stems from their ability to synthesize multiple analytical techniques, addressing the complexities inherent in retail data. By leveraging the clustering capabilities of unsupervised learning alongside the high precision of supervised classification, hybrid models provide a more comprehensive approach to anomaly detection. The study further emphasizes the importance of model parameter selection and tuning, as these significantly impact performance outcomes. Practical applications suggest that hybrid models excel in managing large-scale data variations, making them the most viable solution for detecting anomalies in retail sales, as reflected in Table 2, Figure 2, and Figure 3.

Among all machine learning techniques examined, hybrid models emerge as the most effective approach for anomaly detection in retail sales. These models combine the advantages of both supervised and unsupervised learning, delivering highly accurate and reliable results while maintaining adaptability to real-world data complexities. Their ability to identify both expected and unforeseen anomalies solidifies their role as an essential tool in modern retail analytics.

7. CONCLUSION

Machine learning-based anomaly detection is an effective approach for identifying irregular sales transactions in the retail industry. This study demonstrates that a hybrid approach, combining supervised classification and unsupervised clustering, provides the best performance among the models evaluated.

The research addressed the primary question: can a hybrid machine learning framework combining supervised and unsupervised methods significantly improve anomaly detection accuracy in dynamic retail sales environments? The findings affirmatively answer this question, with the hybrid model achieving superior performance metrics across multiple evaluation criteria.

The proposed framework offers several practical benefits. Retailers can enhance fraud detection mechanisms, optimize inventory management processes, and refine marketing strategies by identifying unexpected anomalies in sales data. Real-time anomaly notification systems integrated within the framework ensure that businesses can react promptly to emerging sales trends or operational issues, minimizing potential revenue loss and improving decision-making.

The study also emphasizes the importance of a flexible anomaly detection architecture capable of adapting to varying retail environments, promotional periods, and regional behaviors.

Future research should explore deep learning techniques, such as autoencoders and long short-term memory (LSTM) networks, to further improve anomaly detection accuracy. Additionally, expanding the system into a fully real-time, automated anomaly detection and alerting pipeline could greatly benefit operational scalability in dynamic retail settings.

8. REFERENCES

Ahmed, M., Mahmood, A. N., & Islam, R. (2016). A survey of anomaly detection techniques in the financial domain. Future Generation Computer Systems, 55, 278–288. https://doi.org/10.1016/j.future.2015.07.002.

Elliott, A., Cucuringu, M., Luaces, M. M., Reidy, P., & Reinert, G. (2019). Anomaly detection in networks with application to financial transaction networks. arXiv preprint arXiv:1901.00402. https://doi.org/10.48550/arXiv.1901.00402.

Grimes, C., Sun, J., & Wang, P. (2023). Advances in deep learning-based anomaly detection for financial fraud. Expert Systems with Applications, 212, 119305. https://doi.org/10.1016/j.eswa.2023.119305.

Haque, S. A., Rahman, M., & Aziz, S. M. (2015). Sensor anomaly detection in wireless sensor networks for healthcare. Sensors, 15(4), 8764–8786. https://doi.org/10.3390/s150408764.

Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. Expert Systems with Applications, 193, 116429. https://doi.org/10.1016/j.eswa.2021.116429.

Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. Science, 349(6245), 255–260. https://doi.org/10.1126/science.aaa8415.

Li, X., & Kumar, P. (2023). Clustering-based anomaly detection in supply chain data: A hybrid model approach. International Journal of Computer Science Studies, 7(2), 112–126. https://doi.org/10.5555/ijcss.2023.00702.

Pereira, J., & Silveira, M. (2019). Learning representations from healthcare time series data for unsupervised anomaly detection. Proceedings of the IEEE International Conference on Big Data and Smart Computing (BigComp), Kyoto, Japan, 1–7. https://doi.org/10.1109/BIGCOMP.2019.8679201.

Pinto, S. O., & Sobreiro, V. A. (2022). Literature review: Anomaly detection approaches on digital business financial systems. Digital Business, 2, 100038. https://doi.org/10.1016/j.digbus.2022.100038.

Priya, K., & Ramanathan, P. (2022). Supervised learning-based transaction anomaly detection in retail banking. International Journal of Computer Science Studies, 6(4), 89–102. https://doi.org/10.5555/ijcss.2022.00604.

Sabic, E., Keeley, D., Henderson, B., & Nannemann, S. (2021). Healthcare and anomaly detection: Using machine learning to predict anomalies in heart rate data. AI & Society, 36, 149–158. https://doi.org/10.1007/s00146-020-01017-w.

Schütte, R. (2017). Information systems for retail companies. In Proceedings of the 29th International Conference on Advanced Information Systems Engineering (CAiSE), 13–25. https://doi.org/10.1007/978-3-319-59536-8_2.

Šustrová, T. (2016). A suitable artificial intelligence model for inventory level optimization. Trends in Economics and Management, 10(24), 48–55. https://doi.org/10.13164/trends.2016.24.48.

Thakur, S., & Singh, V. (2022). Anomaly detection in logistics data using hybrid machine learning models. International Journal of Computer Science Studies, 6(3), 54–67. https://doi.org/10.5555/ijcss.2022.00603.

Ukil, A., Bandyopadhyay, S., Puri, C., & Pal, A. (2016). IoT healthcare analytics: The importance of anomaly detection. Proceedings of the 2016 IEEE 30th International Conference on Advanced Information Networking and Applications (AINA), 994–997. https://doi.org/10.1109/AINA.2016.131.

Zipfel, J., Verworner, F., Fischer, M., Wieland, U., Kraus, M., & Zschech, P. (2023). Anomaly detection for industrial quality assurance: A comparative evaluation of unsupervised deep learning models. Computers & Industrial Engineering, 177, 109045. https://doi.org/10.1016/j.cie.2022.109045.