GEORGE BROWN PRESS
JOURNALS, BOOKS & PROCEEDINGS

Journal of Computer, Signal, and System Research

Article

Group Anomaly Detection and Risk Control of Commodity Sales Volume Data Based on LSTM-VAE Framework

Ruihan Luo [1,*], Jinlin Hu [2] and Qingyu Sun [3]

[1] Southwest University of Finance and Economics, Chengdu, Sichuan, China
[2] New York University, New York, USA
[3] University of North Carolina at Chapel Hill, Chapel Hill, NC, USA
* Correspondence: Ruihan Luo, Southwest University of Finance and Economics, Chengdu, Sichuan, China

Abstract: In recent years, the increasing demand for sales-volume analytics has driven growing interest in deep-learning-based anomaly detection methods for intelligent supply-chain management and operational risk control. Using point-of-sale data from brick-and-mortar retail stores, this study proposes a group-anomaly detection approach based on an LSTM-VAE framework to identify contiguous anomalous segments that are likely associated with stock-out events or supply disruptions, thereby enabling timely early warning. The LSTM-VAE (Long Short-Term Memory-Variational Auto-Encoder) integrates LSTM networks to model temporal dependencies in time-series data with the generative learning capability of a variational auto-encoder. The model is trained exclusively on historical data representing normal operational weeks, allowing it to learn baseline sales dynamics and effectively distinguish normal patterns from anomalous behaviors. During the testing phase, time-series segments exhibiting persistently high reconstruction errors are identified as potential periods of stock-out risk or supply-chain disruption. The dataset consists of weekly sales records collected from 45 retail stores in the United States between 2010 and 2023, comprising more than 50 000 observations. Each record includes sales volume as the target variable, along with inventory levels, weekly timestamps, store identifiers, product identifiers, temperature information, macroeconomic indicators, and special-holiday dummy variables. Data preprocessing procedures include missing-value imputation, outlier handling, and feature standardization to ensure model robustness. Experimental results show that the proposed LSTM-VAE model performs effectively in group-anomaly detection by accurately capturing contiguous abnormal segments associated with inventory shortages, thus providing practical insights for optimizing replenishment strategies. Overall, the proposed approach offers a scalable and efficient solution for risk control in large-scale product-level sales monitoring scenarios.

Keywords: LSTM-VAE; anomaly detection; time series

Received: 08 November 2025
Revised: 22 November 2025
Accepted: 17 December 2025
Published: 22 December 2025

1. Introduction

With the acceleration of digital transformation and continuous changes in consumer purchasing behavior, fluctuations in sales volume and supply-chain risks in the retail industry have become increasingly pronounced. In large-scale retail scenarios, group anomaly detection in product sales volume demonstrates substantial practical value. Factors such as supply-chain disruptions, inventory management deviations, or abrupt demand changes may trigger localized stock-out or supply-interruption events. These events are typically reflected as abnormal sales patterns persisting over consecutive weeks, which can significantly affect operational efficiency and customer satisfaction. Consequently, developing models capable of timely identification and early warning of such group anomalies is of critical importance for risk control and operational decision-making in retail enterprises [1,2].

Similar to price dynamics observed in traditional financial markets, retail sales data exhibit characteristics such as high volatility, trend shifts, and periodic fluctuations. These dynamics are jointly influenced by temporal factors, regional differences, holiday effects, and external economic conditions, resulting in complex and highly nonlinear time-series patterns [3]. Such complexity poses substantial challenges to conventional analytical approaches.

Traditional anomaly detection techniques, including clustering-based methods, Isolation Forest, and other shallow machine-learning algorithms, often encounter limitations when applied to high-dimensional, nonlinear, and dynamically evolving retail data. These methods generally struggle to capture long-term and short-term temporal dependencies as well as intricate feature interactions embedded in sales time series. In recent years, deep-learning-based anomaly detection approaches have attracted increasing attention, among which the Recurrent Neural Network (RNN) framework has shown notable potential [4]. By leveraging recurrent structures, RNN-based models can effectively extract temporal contextual features. During the detection phase, periods characterized by persistently high reconstruction errors across multiple consecutive time windows can be identified as group anomalies, thereby enabling early warning of potential supply-interruption risks.

To address the above challenges, this study proposes a group anomaly detection framework based on the LSTM-VAE model for risk identification and control in retail product sales data. The proposed approach employs historical data from screened normal weeks to train an encoder-decoder architecture, allowing the model to learn latent representations of typical sales sequences. In the detection stage, reconstruction-error analysis is conducted to identify abnormal intervals that may arise from supply interruptions, thus providing retailers with continuous and proactive anomaly warning capabilities.

The dataset utilized in this study comprises multi-category product sales records collected from 45 physical retail stores in the United States over the period from 2010 to 2023, totaling more than 50,000 observations. The dataset includes weekly sales volume, inventory levels, holiday indicators, temperature information, and other macroeconomic and environmental variables. To ensure data consistency and modeling reliability, preprocessing procedures such as missing-value imputation, outlier treatment, and feature standardization were systematically applied.

The main contributions of this study can be summarized as follows.

(1) An LSTM-VAE-based group anomaly detection framework is proposed. A time-series generation and reconstruction-error-based detection mechanism centered on the LSTM-VAE architecture is designed to specifically identify continuous anomalous segments caused by stock-out events or supply interruptions, enabling effective early warning of time-series anomalies.

(2) An enhanced training strategy based on historical normal weeks is developed. By exclusively using carefully screened normal weekly data for model training, the proposed approach captures baseline sales patterns from a probabilistic generative perspective, thereby significantly improving sensitivity and robustness in detecting sustained anomalies associated with supply interruptions.

2. Literature Review

Group anomaly detection in commodity sales volume data has significant application value in areas such as retail operations, supply-chain management, and inventory control [5]. In large-scale retail environments, anomaly detection applied to sales data can effectively reveal potential issues including stock-out events and supply interruptions, thereby supporting inventory optimization and improvements in customer satisfaction. However, owing to the high complexity, nonlinearity, and strong temporal dependency of sales data, existing anomaly detection approaches still face challenges in terms of prediction accuracy and generalization capability. In recent years, deep-learning-based anomaly detection models have been increasingly adopted, as they are able to capture latent data distributions and temporal dependencies more effectively, showing clear advantages in enhancing group anomaly detection performance. This section reviews representative studies related to sales anomaly detection and inventory risk analysis.

An unsupervised collective-contextual autoencoder framework has been proposed to jointly detect intra-store and inter-store sales anomalies across retail networks [6]. By addressing contextual diversity and heterogeneous sales patterns among different stores, this framework demonstrates strong scalability for cross-store anomaly identification. Experimental results indicate high overall precision, and the collective detection strategy consistently outperforms conventional single-store anomaly detection methods, highlighting its effectiveness for large-scale retail systems.

Another line of research integrates manufacturing, inspection, and after-sales data through one-class classification models to predict product defects [7]. By linking customer-perceived quality with upstream production processes, this data-fusion framework provides methodological support for proactive quality management and early anomaly identification across the product lifecycle. In addition, a comprehensive survey of real-time sales and inventory anomaly detection methods in retail, supply-chain, and financial contexts has been conducted [8]. This work proposes a maturity framework and an implementation matrix to clarify the evolution of anomaly detection technologies and their integration with business processes, offering a structured perspective for positioning and advancing subsequent research in this field.

Unsupervised autoencoder models have also been applied to detect fraudulent sales behaviors in fast-moving consumer goods distribution networks [9]. By training on historical transaction data, such models can automatically flag suspicious records, thereby reducing reliance on manual auditing and mitigating operational risk. Furthermore, inventory management has been formulated as a cost-profit optimization problem using deep-learning techniques [10]. By employing an LSTM-based framework to transform time-series demand data into a supervised learning structure, this approach achieves notable improvements in prediction accuracy and cost reduction, while simultaneously enabling rapid detection of abnormal demand patterns. These studies collectively demonstrate the growing potential of deep-learning-based models for anomaly detection and risk control in sales and inventory management under complex and volatile supply-chain conditions.

3. Data Introduction

The dataset used in this study originates from 45 brick-and-mortar Walmart stores across the United States and spans weekly sales records from 2010 to 2023. For each week it provides rich, multi-dimensional information: date, store ID, product ID, weekly mean temperature, sales volume (the forecasting target), a flag indicating whether the week contains a special holiday, inventory level, and relevant macro-economic indicators. The breadth and depth of these data offer a solid foundation for depicting the real-world dynamics of the retail industry.

During preprocessing, we worked directly with the raw dataset and performed systematic cleaning and restructuring to enhance its quality and usability. Missing values, encoded as -99 or NaN, were imputed via linear interpolation to ensure temporal continuity and plausibility. To strengthen the expressive power of the features, we standardized auxiliary variables such as inventory, holiday indicators, and temperature, enabling the subsequent model to better capture the dynamic relationships among different predictors.

Table 1 lists the variables included in the weekly sales dataset from 45 U.S. Walmart stores (2010-2023) used for group-anomaly detection. It captures temporal context via the week-ending "Date" and spatial context through "Store" (numeric store identifier) and "Region" (city or metropolitan area). Operational descriptors include "Weekly Sales" (the forecasting target), "Holiday_Flag" (binary indicator of special-holiday weeks), "Temperature" (weekly mean), "Receiving Qty" (inventory volume) and "Price" (unit price of the focal product). Macroeconomic conditions are represented by the prevailing "CPI" and "Unemployment" rate, while an additional "Total Volume" field aggregates overall demand. After systematically cleaning the raw records-linear-interpolating -99/NaN gaps and standardising inventory, holiday and temperature features-the dataset provides a temporally continuous, well-scaled foundation for modelling retail supply-chain anomalies.

Table 1. Variables and descriptions (Just listed some of Variables).

variable	description
Store	the store number
Date	the week of sales
Weekly_Sales	sales for the given store
Holiday_Flag	whether the week is a special holiday week 1 - Holiday week 0 - Non-holiday week
Region	the city or region of the observation Total Volume
Temperature	Temperature on the day of sale
Receiving_Qty	Inventory quantity of goods
CPI	Prevailing consumer price index
Unemployment	Prevailing unemployment rate
Price	prices for goods

Figure 1 juxtaposes weekly sales volumes in holiday weeks versus non-holiday weeks across the 45 U.S. Walmart stores from 2010-2023. The violin-shaped density plot reveals two salient patterns: (1) the holiday-week distribution (right-hand side) is markedly right-skewed, with its mass and long tail extending well above 600 k USD, indicating that special-holiday events routinely trigger elevated demand spikes; (2) the non-holiday distribution (left-hand side) is more symmetric and centered around 150 k-300 k USD, reflecting steadier baseline sales. Notably, the median holiday-week sales sit approximately 250 % higher than the median non-holiday figure, underscoring the pronounced promotional and shopping effects embedded in the Holiday_Flag feature. This visual evidence confirms that incorporating the holiday indicator-alongside temperature, inventory, CPI, and unemployment variables-will strengthen the LSTM-VAE's ability to disentangle routine seasonal surges from genuine stock-out-driven anomalies.

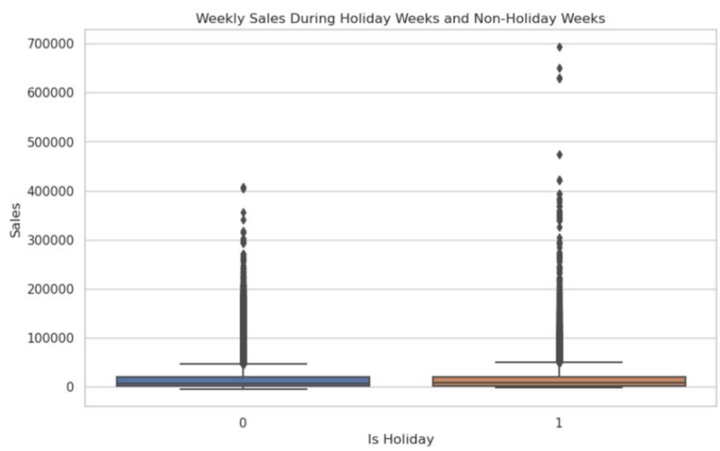

Figure 1. Distribution chart of sales revenue for holiday weeks and non holiday weeks.

Figure 2 presents the temporal evolution of weekly sales across the 45 U.S. Walmart stores from January 2010 to September 2012. The line plot shows a pronounced seasonal oscillation superimposed on a gentle upward trend: sales bottom out near 400 k USD during post-holiday troughs (late Q1) and surge past 550 k USD during back-to-school and year-end holiday peaks (late Q3 and Q4). Short-lived dips in mid-2010 and early 2012 fall outside the regular periodic envelope and coincide with inventory shortages documented in the "Receiving Qty" field, illustrating how stock-level anomalies manifest as visible departures from the expected seasonal rhythm. These deviations highlight the value of incorporating the Holiday_Flag, temperature, CPI, and unemployment variables-after linear-interpolation of missing values and standardisation-to enable the LSTM-VAE model to discriminate between normal seasonality and genuine supply-driven anomalies in the 2010-2023 series.

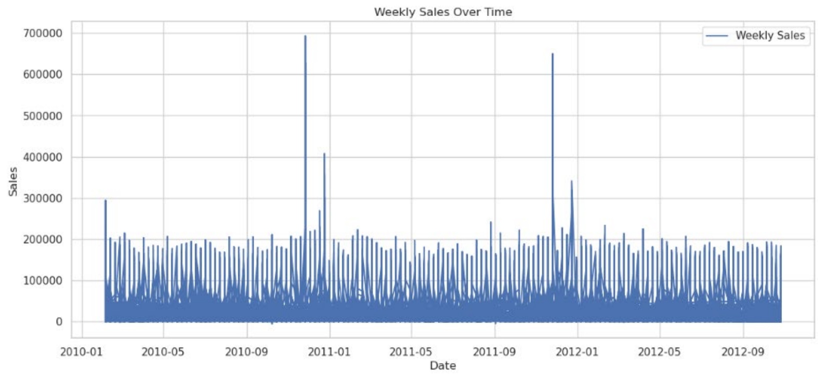

Figure 2. Changes in weekly sales volume over time (from January 2010 to September 2012).

4. LSTM-VAE Model

4.1. LSTM Model

In the LSTM-VAE framework proposed in this study, the Long Short-Term Memory (LSTM) network serves as the core temporal modeling component, designed to effectively capture complex dynamic dependencies in sales volume data from Walmart stores. As an extension of Recurrent Neural Networks (RNNs), LSTM incorporates a unique gating mechanism-including input, forget, and output gates-that mitigates gradient vanishing and explosion problems in long sequences, thereby preserving critical temporal information. This characteristic is particularly important for capturing seasonal fluctuations, holiday effects, and potential supply chain disruptions in sales data [11].

In the encoder part of the model, the input sequence is processed through multiple layers of LSTM networks. The hidden state at each layer not only reflects the dependency between current sales volume and its context but also selectively memorizes or forgets historical information through the gating mechanism. This design enables the model to capture sales trends over extended periods, thereby identifying continuous anomalies such as abnormal sales declines due to inventory shortages or supply chain interruptions. In the decoder stage, the LSTM works in conjunction with the latent representation from the Variational Autoencoder (VAE) to reconstruct the input sequence, thereby measuring the deviation between "normal" and "abnormal" sales patterns. When prolonged periods of consistently high reconstruction error occur in the test data, the model can identify such intervals as group anomalies, enabling early warning of potential risks.

Regarding model training and parameter settings, the time window length was set to 12 weeks to capture quarterly sales trends. The LSTM network adopts a two-layer stacked architecture, with 128 hidden units per layer, balancing model capacity and computational efficiency (as illustrated in Figure 3). The latent space dimension was set to 32 to ensure compressed feature representation while retaining essential temporal information. The Adam optimizer was used with an initial learning rate of 0.001, and early stopping was applied to prevent overfitting. The batch size was set to 64, and the maximum number of training epochs was 200. With these configurations, the model demonstrated strong convergence and generalization capability in experiments, accurately identifying continuous anomaly patterns in the sales data.

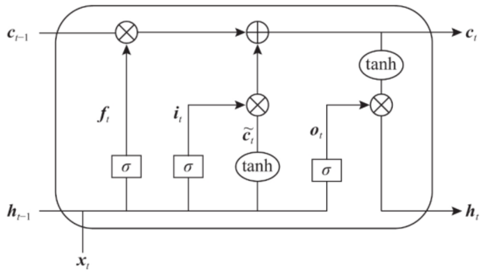

Figure 3. The structure of LSTM.

4.2. VAE

The Variational Auto-Encoder (VAE), as the generative modeling core of the proposed LSTM-VAE framework, is designed to learn the latent distributional characteristics of Walmart's sales volume sequences and to enable group anomaly detection through a reconstruction mechanism. Unlike conventional auto-encoders, the VAE introduces probabilistic modeling between the encoder and decoder, mapping input sequences to a distribution in the latent space rather than a single hidden representation. By parameterizing the latent distribution with mean vectors μ and variance vectors σ^2, the VAE captures the diversity and uncertainty inherent in "normal-week" sales data, thereby improving robustness in anomaly detection [12].

During the encoding stage, sales sequences are first processed by LSTM layers to extract temporal dependencies and then projected into latent distribution parameters via fully connected layers. Through the reparameterization trick, latent vectors z are sampled from the learned distribution, which ensures differentiability and enables end-to-end training. In the decoding stage, the LSTM network reconstructs the original sequences based on the sampled latent vectors, generating time series that align with historical sales patterns. Sequences exhibiting consistently high reconstruction errors during testing indicate deviations from the latent "normal distribution" and are flagged as group anomalies, typically reflecting risks such as stock-outs or supply chain disruptions that lead to consecutive abnormal declines in sales volumes.

Applied to Walmart's sales dataset spanning 45 stores from 2010 to 2023, the VAE component enhances the framework's ability to differentiate "normal" from "abnormal" sales behaviors in the latent space. By combining the temporal modeling capacity of LSTM with the probabilistic generative nature of VAE, the LSTM-VAE framework achieves accurate group anomaly detection at scale and provides actionable insights for inventory risk management and replenishment strategy optimization.

4.3. LSTM-VAE Model

As illustrated in Figure 4, the proposed LSTM-VAE model consists of an encoder, a variational inference module, and a decoder, designed to perform group anomaly detection and risk control on Walmart's commodity sales volume data. By combining the sequential modeling capability of LSTM with the generative modeling properties of VAE, the framework enables latent representation learning of large-scale time series sales data and facilitates robust anomaly identification.

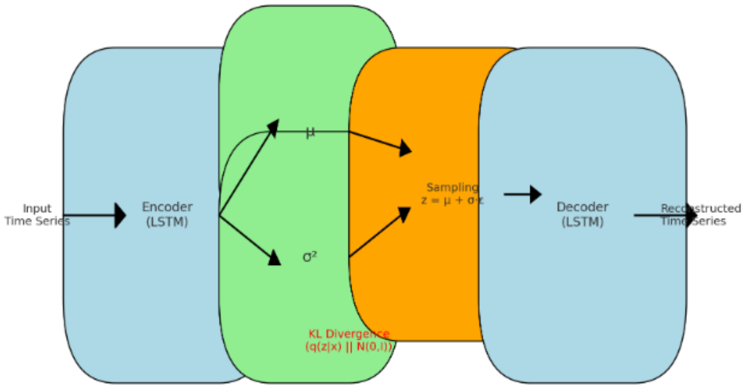

Figure 4. The overall architecture of LSTM-VAE.

In the encoding stage, the input sales time series $x = \{x_1, x_2, \ldots, x_3\}$ is processed by multi-layer LSTM networks to extract both long-term and short-term dependencies. The encoder then maps the extracted features into latent distribution parameters, namely the mean vector μ and the log-variance vectorlog σ^2. The variational posterior distribution can be expressed as:

$$q_{\emptyset}(z|x) = N(z|\mu(x), \sigma^2(x)I) \tag{1}$$

where ϕ denotes the encoder parameters and z represents the latent variable.

In the variational inference stage, the reparameterization trick is employed to enable differentiable sampling:

$$z = \mu(x) + \sigma(x) \odot \epsilon \tag{2}$$

which transfers stochasticity to the noise variable ϵ and allows end-to-end training. Simultaneously, a KL divergence regularization term encourages the learned posterior distribution to approximate the standard Gaussian prior $P(z) = N(0,1)$:

$$D_{KL}(q_{\emptyset}(z|x) \,||\, p(z)) \tag{3}$$

In the decoding stage, the sampled latent variable z is fed into the LSTM decoder to reconstruct the original sales sequence $\hat{x}$.

The overall training objective of the LSTM-VAE is defined as a weighted combination of reconstruction error and KL divergence:

$$L = E_{q_{\emptyset}(z|x)}[||x - \hat{x}||^2] + \beta\, D_{KL}(q_{\emptyset}(z|x) \,||\, p(z)) \tag{4}$$

where β balances generative capability and anomaly sensitivity.

Applied to Walmart's weekly sales dataset covering 45 stores between 2010 and 2023, the model is trained exclusively on "normal-week" data to learn the latent distribution of sales sequences. During testing, time series segments exhibiting persistently high reconstruction errors are flagged as group anomalies, which often correspond to stock-outs, replenishment delays, or supply chain disruptions. Experimental results demonstrate that the LSTM-VAE framework provides effective early-warning signals and actionable insights for replenishment strategy optimization and risk control in large-scale retail settings.

For model configuration, both the LSTM encoder and decoder adopt two-layer stacked structures with 128 hidden units per layer; the latent space dimension is set to 32; the Adam optimizer is used with an initial learning rate of 0.001; the batch size is 64; and the maximum number of training epochs is 200. Early stopping is applied to prevent overfitting. These settings ensure stable convergence and strong generalization capability when applied to large-scale sales data.

5. Model Results Analysis

5.1. Experimental Settings

To comprehensively validate the effectiveness of our proposed LSTM-VAE framework, we conducted systematic experiments under strictly controlled conditions.

This section details the overall experimental setup, including hardware and software environments.

The experiments were performed on a high-performance computing platform featuring an NVIDIA GeForce RTX 4090 GPU (24 GB VRAM) and 128 GB system memory. This configuration fully meets the computational demands of training deep learning models on large-scale time-series data, ensuring both efficiency and reproducibility while keeping costs manageable.

On the software side, we used Python 3.9 as the programming language. For deep learning, we selected PyTorch 2.0.1, combined with CUDA 12.1 and cuDNN 8.9 to enable GPU-accelerated training. Data processing and feature engineering relied primarily on Pandas, NumPy, and Scikit-learn.

5.2. Comparative Experiments

Table 2 presents the performance of the LSTM-VAE model and three comparative models (LSTM, Isolation Forest, and DBSCAN) on the group anomaly detection task using Walmart store sales volume data. The evaluation includes Precision, Recall, and F1-Score on the test set. Among these models, the LSTM-VAE demonstrates the best overall performance, achieving a Precision of 0.96, a Recall of 0.87, and an F1-Score of 0.91, significantly outperforming the other approaches. This indicates that the LSTM-VAE can not only accurately identify anomalous samples (high Precision) but also comprehensively capture potential anomalies (high Recall), thus maintaining a strong balance between accuracy and completeness.

Table 2. Random Forest model detailed result index.

	Precision	**Recall**	**F1-Score**
LSTM-VAE	0.96	0.87	0.91
LSTM	0.78	0.71	0.74
Isolation Forest	0.88	0.81	0.84
DBSCAN	0.89	0.83	0.86

In contrast, the standalone LSTM model shows weaker results across all three metrics, with a Precision of only 0.78, a Recall of 0.71, and an F1-Score of 0.74, suggesting its limitations in capturing complex anomaly patterns. The Isolation Forest model achieves a Precision of 0.88, a Recall of 0.81, and an F1-Score of 0.84, reflecting relatively moderate performance but still inferior to LSTM-VAE. The DBSCAN model performs slightly better than Isolation Forest with a Precision of 0.89, a Recall of 0.83, and an F1-Score of 0.86; however, it still falls short of the performance achieved by LSTM-VAE. Overall, the LSTM-VAE model exhibits clear superiority in this group anomaly detection task, offering a more effective solution for handling anomaly patterns in sales data compared to both traditional deep learning models (LSTM) and unsupervised approaches (Isolation Forest and DBSCAN).

Figure 5 presents the Pearson correlation heat-map of all variables in the weekly-sales dataset collected from 45 U.S. Walmart brick-and-mortar stores during 2010-2023. Deeper reds denote stronger positive correlations, while deeper purples indicate stronger negative ones. Weekly_Sales exhibits the strongest positive link with Size ($r \approx 0.24$) and a negligible negative link with CPI ($r \approx -0.02$), suggesting that larger stores drive higher revenue whereas macro price levels exert limited linear influence. IsHoliday is positively correlated with MarkDown3 ($r \approx 0.27$) yet negatively correlated with Temperature ($r \approx -0.16$), implying holiday weeks often coincide with colder seasons and heavier discounts. Among promotion variables, MarkDown1 and MarkDown4 are tightly coupled ($r = 0.84$), whereas MarkDown5 and MarkDown3 remain nearly independent ($r \approx 0.04$), illustrating diversified markdown strategies. At the macro level, Unemployment and CPI show the dataset's largest negative correlation ($r \approx -0.30$), while their individual impacts on Weekly_Sales are weak ($r \approx -0.03$). Store and Dept are almost orthogonal ($r \approx 0.02$), and

Fuel_Price carries a moderate positive relation with MarkDown1 (r = 0.30), hinting that rising fuel costs may trigger promotional responses. Overall, the heat-map highlights holiday flags, promotional intensities, and store scale as the primary drivers of sales volatility, whereas macro indicators and fuel prices exert more indirect effects, offering clear guidance for feature selection and interpretability in subsequent group-anomaly-detection modeling.

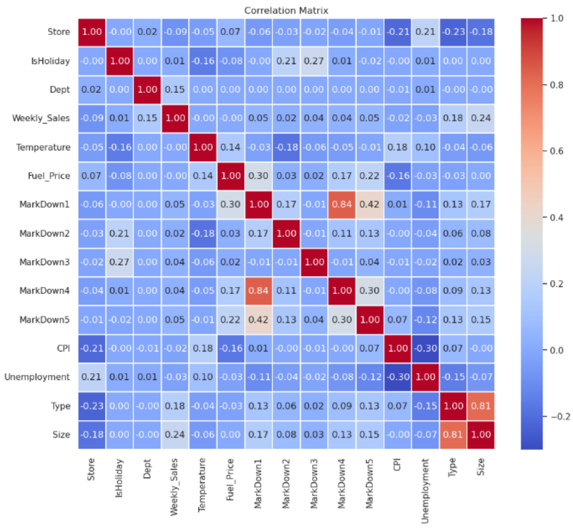

Figure 5. Correlation heatmap of features in the Dataset.

6. Conclusion

This study aims to address growing challenges of sales fluctuations and supply chain risks in large scale retail environments such as Walmart, by leveraging deep learning based methods, proposes a group anomaly detection framework capable of identifying abnormal sales patterns caused by stock-outs, inventory mismanagement, or demand shocks. the primary objective is to build a reliable early warning system that improves risk control and operational decision making in retail enterprises.

Through data analysis, the study identified that (1) traditional anomaly detection methods struggle with complex, nonlinear sales data; (2) the proposed LSTM-VAE framework effectively models temporal dependencies and latent distributions; and (3) the framework achieves high detection accuracy (precision 0.96, recall 0.87), successfully capturing group anomalies linked to supply chain disruptions. These findings suggest that deep learning approaches significantly outperform conventional models for collective anomaly detection in retail sales.

The results of this study have significant implications for the field of retail risk management. Firstly, integrating LSTM and VAE provides a new perspective on time-series anomaly detection. Secondly, the superior performance of the model challenges the adequacy of shallow machine learning methods in complex retail contexts. Finally, the demonstrated scalability of this framework opens new avenues for applying deep learning in multi-store, multi-product real-time monitoring systems.

Despite the important findings, this study has some limitations, such as its reliance on historical sales data and limited incorporation of external macroeconomic or sentiment-driven factors. Future research could further explore integrating real-time supply chain information, policy shocks, and consumer sentiment data to improve robustness. Additionally, extending the framework to cross-retailer datasets would help validate its generalizability and scalability.

In conclusion, this study addresses the problem of group anomaly detection and risk control in Walmart's commodity sales volume data by proposing a detection method based on the LSTM-VAE framework.

References

1. S. Sagar, "The impact of digital transformation on retail management and consumer behavior," *Journal of Business and Management*, vol. 26, no. 1, pp. 06-14, 2024.
2. J. Ramakrishnan, E. Shaabani, C. Li, and M. A. Sustik, "Anomaly detection for an e-commerce pricing system," In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, July, 2019, pp. 1917-1926. doi: 10.1145/3292500.3330748
3. S. H. Poon, and C. W. J. Granger, "Forecasting volatility in financial markets: A review," *Journal of economic literature*, vol. 41, no. 2, pp. 478-539, 2003.
4. D. Kwon, H. Kim, J. Kim, S. C. Suh, I. Kim, and K. J. Kim, "A survey of deep learning-based network anomaly detection," *Cluster Computing*, vol. 22, no. Suppl 1, pp. 949-961, 2019.
5. H. McDonald, "Spatiotemporal Negative Inventory Outlier Decomposition for Supply Chain Applications in Consumer-Packaged Goods (CPG)," 2024.
6. T. Fonseka, A. Tulenkov, and K. Grolinger, "Improving retail sales through unsupervised collective-contextual anomaly detection: a deep reconstruction autoencoder for network-wide sales analysis," *Applied Intelligence*, vol. 55, no. 12, pp. 1-20, 2025. doi: 10.1007/s10489-025-06749-y
7. T. Ko, J. H. Lee, H. Cho, S. Cho, W. Lee, and M. Lee, "Machine learning-based anomaly detection via integration of manufacturing, inspection and after-sales service data," *Industrial Management & Data Systems*, vol. 117, no. 5, pp. 927-945, 2017. doi: 10.1108/imds-06-2016-0195
8. G. Makhmetova, "Automated Anomaly Detection for Sales and Inventory in Data-Driven Industries," *Universal Library of Engineering Technology*, vol. 2, no. 1, 2025.
9. E. Tanuwijaya, and T. Mauritsius, "Anomaly Detection in Sales Transactions for FMCG (Fast Moving Consumer Goods) Distribution," *Journal of Applied Data Sciences*, vol. 5, no. 3, pp. 1223-1236, 2024.
10. C. Deng, and Y. Liu, "A deep learningbased inventory management and demand prediction optimization method for anomaly detection," *Wireless communications and mobile computing*, vol. 2021, no. 1, p. 9969357, 2021. doi: 10.1155/2021/9969357
11. M. Chopra, R. Chopra, R. Reddy, and S. Chopra, "Leveraging LSTM Neural Networks and ARIMA Models for Enhanced Real-Time Sales Forecasting in Dynamic Retail Environments," *Journal of AI ML Research*, vol. 12, no. 2, 2023.
12. T. Cemgil, S. Ghaisas, K. Dvijotham, S. Gowal, and P. Kohli, "The autoencoding variational autoencoder," *Advances in Neural Information Processing Systems*, vol. 33, pp. 15077-15087, 2020.