GEORGE BROWN PRESS
JOURNALS, BOOKS & PROCEEDINGS

Economics and Management Innovation

Article

DFG-JointGNN: A Data-Flow Graph Neural Network for Unified Sales Order Fraud Detection and Sales Forecasting

Yuxuan Liu [1,*], Shuhan Liu [2] and Hongjing Shao [3]

[1] South China University of Technology, Guangzhou, China
[2] University of Utah, Salt Lake City, UT, USA
[3] University of Shanghai for Science and Technology, Shanghai, China
* Correspondence: Yuxuan Liu, South China University of Technology, Guangzhou, China

Abstract: Sales order systems in e-commerce platforms, enterprise information systems, and supply chain finance generate large-scale transactional data streams that exhibit strong process dependency and complex entity interactions. Fraudulent behaviors such as fake orders, brushing transactions, and refund abuse not only cause direct economic losses but also distort historical sales data, leading to biased and unreliable sales forecasting. Existing studies typically address sales order fraud detection and sales forecasting as independent tasks, overlooking the intrinsic coupling between fraud risk and sales demand in real transaction flows. In this paper, we propose DFG-JointGNN, a data-flow graph neural network that unifies sales order fraud detection and sales forecasting within a single learning framework. We model the full lifecycle of sales orders-including order creation, payment, logistics fulfillment, and settlement-as a temporal heterogeneous data-flow graph, where customers, orders, merchants, payment accounts, and logistics nodes are jointly represented. A relation-aware temporal graph attention network is employed to capture both structural dependencies and temporal dynamics of transaction flows. Fraud detection is performed at the order level, while sales forecasting is conducted at the merchant or product level through a fraud-aware aggregation mechanism that explicitly suppresses the influence of high-risk orders. Experiments on real-world and semi-synthetic sales order datasets show that DFG-JointGNN consistently outperforms state-of-the-art baselines. For fraud detection, it achieves an AUC of 0.956, improving approximately 6.3% over the strongest baseline (Temporal GAT, 0.893). For sales forecasting, it reduces RMSE to 78.4, a 15.2% improvement compared to the fraud-agnostic Temporal GAT model. These results confirm that jointly modeling fraud detection and sales forecasting with transaction data-flow graphs enhances both predictive accuracy and operational robustness.

Keywords: Sales Order Fraud Detection; Sales Forecasting; Data-Flow Modeling; Graph Neural Networks; Multi-Task Learning; Transaction Graph

Received: 27 December 2025
Revised: 01 February 2026
Accepted: 11 February 2026
Published: 18 February 2026

1. Introduction

Sales order management is a core component of modern enterprise information systems, e-commerce platforms, and supply chain finance. A typical sales order involves multiple interconnected stages, including order creation, payment authorization, logistics fulfillment, and final settlement. These stages generate large volumes of transactional data streams characterized by strong process dependency, temporal dynamics, and heterogeneous entity interactions. Accurate analysis of sales order data is essential for both operational decision-making and strategic planning, particularly in sales forecasting and financial risk management.

However, sales order systems are increasingly exposed to sophisticated fraudulent behaviors, such as fake orders, brushing transactions, refund abuse, and cash-out schemes. Furthermore, recent studies indicate that fraud forms are evolving rapidly with the

emergence of artificial intelligence. For instance, Liu et al highlighted the rise of AI-driven multimodal fraud involving forged text, images, and audio, necessitating more robust feature extraction and fusion mechanisms [1]. These fraudulent activities not only result in direct financial losses but also contaminate historical sales data, leading to distorted demand signals and unreliable sales forecasts. In real-world business scenarios, fraud detection and sales forecasting are therefore intrinsically coupled tasks: undetected fraudulent orders bias sales statistics, while inaccurate forecasts may further amplify operational risks.

Despite this strong coupling, most existing studies treat sales order fraud detection and sales forecasting as independent problems. Fraud detection methods typically focus on identifying anomalous transactions using classification models or graph-based techniques, whereas sales forecasting models assume that historical sales data are clean and trustworthy. This separation neglects the process-driven nature of transaction data and fails to exploit the mutual reinforcement between fraud detection and forecasting. As a result, current approaches often struggle to provide robust and reliable decision support in complex transactional environments.

To address these limitations, this paper proposes DFG-JointGNN, a Data-Flow Graph Neural Network for Unified Sales Order Fraud Detection and Sales Forecasting. The key idea of DFG-JointGNN is to model the full lifecycle of sales orders as a temporal heterogeneous data-flow graph, where customers, orders, merchants, payment accounts, and logistics nodes are jointly represented, and order flow, cash flow, and logistics flow are explicitly encoded as directed relations. Based on this data-flow graph, DFG-JointGNN employs a relation-aware temporal graph neural network to learn expressive representations of transactional entities. Fraud detection is conducted at the order level, while sales forecasting is performed at the merchant or product level through a fraud-aware aggregation mechanism that suppresses the influence of high-risk orders on forecasting outcomes.

The main contributions of this paper are summarized as follows:

We propose a transaction data-flow graph modeling framework that captures the full lifecycle and process semantics of sales orders.

We develop DFG-JointGNN, a unified graph neural network that jointly performs sales order fraud detection and sales forecasting within a single learning framework.

We introduce a fraud-aware forecasting mechanism that explicitly incorporates fraud risk into sales prediction, improving robustness against anomalous orders.

Extensive experiments demonstrate that the proposed approach significantly outperforms state-of-the-art baselines in both fraud detection accuracy and sales forecasting precision.

2. Literature Review

In recent years, the rapid digitalization of enterprise operations and e-commerce platforms has led to the explosive growth of sales order transaction data. This data is inherently process-driven and involves complex interactions among customers, merchants, payment systems, and logistics entities. Consequently, a growing body of research has focused on fraud detection in transactional systems, sales forecasting using machine learning models, and graph-based representations for complex relational data. This section reviews the major lines of work closely related to this study, including sales order fraud detection, sales forecasting models, graph neural networks for transaction modeling, and emerging joint learning frameworks.

2.1. Sales Order and Transaction Fraud Detection

Early studies on transaction fraud detection primarily relied on rule-based systems and traditional machine learning models, such as logistic regression and decision trees [2]. With the increasing complexity of fraudulent behaviors, ensemble models and gradient

boosting techniques, including Random Forest and XGBoost, have been widely adopted to capture non-linear patterns in transactional data [3].

More recently, deep learning-based methods have shown superior performance in detecting sophisticated fraud patterns. Recurrent neural networks (RNNs) and LSTM models have been applied to model sequential transaction behaviors [4], while autoencoder-based approaches aim to identify anomalies through reconstruction errors [5]. However, these methods typically treat transactions as independent records or simple sequences, failing to capture the relational structure and process dependencies inherent in sales order systems.

To address this limitation, graph-based fraud detection methods have gained increasing attention. By modeling entities such as users, accounts, and transactions as nodes and edges, graph-based approaches can effectively capture complex interaction patterns [6]. Nevertheless, most existing graph-based fraud detection models focus solely on detection accuracy and do not consider the downstream impact of fraud on business analytics tasks such as sales forecasting.

2.2. Sales Forecasting with Machine Learning and Deep Learning

Sales forecasting has long been a fundamental problem in operations management and business intelligence. Classical statistical models, including ARIMA and exponential smoothing, have been widely used due to their interpretability and simplicity [7]. However, these models struggle to handle non-stationary patterns and high-dimensional features in modern sales data.

With the advancement of machine learning, neural network–based forecasting models such as LSTM and GRU have been extensively studied for capturing temporal dependencies in sales time series [8]. Recent works further incorporate attention mechanisms and temporal convolutional networks to improve long-term forecasting accuracy [9]. Furthermore, in high-frequency financial domains, hybrid deep learning architectures (e.g., combining LSTM with convolutional or graph modules) have demonstrated superior capability in capturing complex market dynamics and volatility risks [10]. Despite these advances, most forecasting models assume that historical sales data are clean and reliable, overlooking the impact of fraudulent or abnormal orders.

Several studies have acknowledged that data quality issues, including fraud and anomalies, can significantly degrade forecasting performance [11]. However, these works typically rely on preprocessing or post-hoc filtering strategies, rather than integrating fraud awareness directly into the forecasting model.

2.3. Graph Neural Networks for Transaction and Financial Data Modeling

Graph neural networks (GNNs) have emerged as a powerful tool for modeling relational and structured data. Representative models such as GCN [12], GAT [13], and heterogeneous GNNs have been successfully applied to social networks, recommender systems, and financial transaction graphs.

In the context of fraud detection, GNNs enable information propagation across related entities, improving the identification of collective fraud behaviors [6]. Temporal extensions of GNNs further allow modeling the dynamic evolution of transaction graphs [14]. More recently, hybrid frameworks have emerged to enhance feature representation in fraud detection. For instance, Luo et al. proposed integrating Large Language Models (LLMs) with Graph Convolutional Networks (GCNs) to capture both semantic information from unstructured text and structural dependencies in transaction graphs, achieving high accuracy even in highly imbalanced datasets [15]. Nevertheless, most existing GNN-based studies focus on single-task learning, either fraud detection or risk prediction, without considering joint optimization with forecasting objectives.

2.4. Joint Learning of Fraud Detection and Forecasting

Only a limited number of studies have explored the joint modeling of fraud detection and forecasting tasks. Some works adopt multi-task learning frameworks to share representations across related objectives [10], while others sequentially apply fraud filtering before forecasting. These approaches, however, fail to explicitly model transaction data as process-driven data-flow graphs and do not fully exploit the structural and temporal dependencies among transactional entities.

In contrast, the proposed DFG-JointGNN addresses these limitations by unifying sales order fraud detection and sales forecasting within a single data-flow graph neural network framework. By explicitly modeling order flow, cash flow, and logistics flow, and introducing a fraud-aware forecasting mechanism, DFG-JointGNN bridges the gap between transaction-level risk detection and aggregate-level sales analytics.

3. Methodology

In this section, we present the detailed methodology of DFG-JointGNN, a unified framework that simultaneously addresses sales order fraud detection and sales forecasting through data-flow graph modeling and graph neural network learning. We first describe the construction of the sales transaction data-flow graph, followed by the design of node and edge representations. We then detail the relation-aware temporal graph neural network backbone and the joint learning paradigm that integrates fraud detection with sales forecasting.

3.1. Transaction Data-Flow Graph Construction

To effectively model the dynamic and interdependent behavior of sales orders in real-world business environments, we introduce a transaction data-flow graph representation. The core idea is to transform the sales order process, along with associated payment and logistics activities, into a unified graph structure that captures both relational and temporal dependencies. Formally, at any discrete time window t, the transaction data-flow graph is defined as a heterogeneous directed graph $G_t = (V_t, E_t, R, X)$, where V_t denotes the set of nodes representing entities such as customers, orders, merchants, payment accounts, products, and logistics hubs; E_t represents the set of directed edges defined by transactional relations; R is the set of relation types; and X denotes the associated feature vectors.

The relation types R are designed to reflect real business processes: order creation (customer → order), payment execution (order → payment account), merchant affiliation (merchant → order), product inclusion (order → product), and logistics activities (order → logistics). Additionally, order-to-order relations capture similarity or temporal proximity, enabling the model to learn connection patterns between related orders. Directed edges naturally encode process flow, which is critical for capturing causality in transaction events. This graph construction allows the representation of sequential and relational information in an integrated structure, enabling upstream fraud signals to influence downstream forecasting representations.

Figure 1 provides an overview of the DFG-JointGNN framework, illustrating its unifified approach to sales order fraud detection and sales forecasting. The Key Features Highlighted:

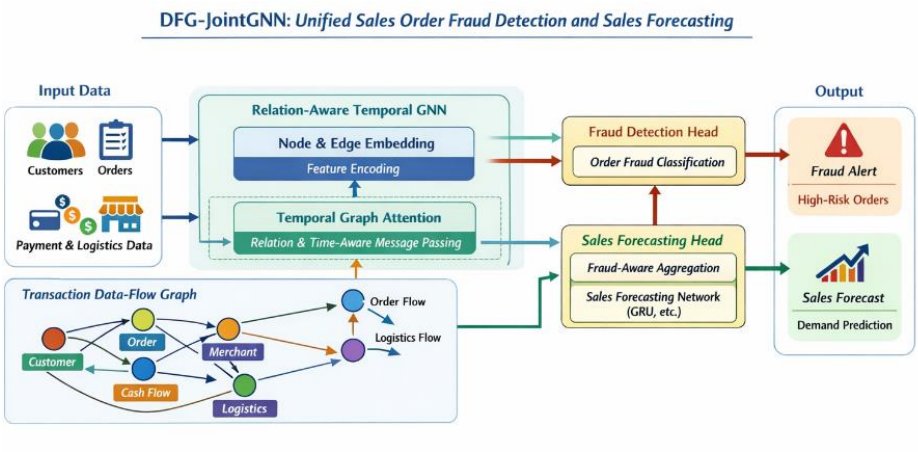

Figure 1. Overall flowchart of the model.

Unifified Framework: The model integrates fraud detection and sales forecasting into a single framework, allowing mutual reinforcement between tasks.

Fraud-Aware Aggregation: High-risk orders are explicitly suppressed in the forecasting process, mitigating their negative impact on sales predictions.

Temporal and Relational Modeling: The use of temporal graph attention and relation-aware mechanisms ensures that both structural and temporal dependencies are effectively captured.

Overall, Figure 1 visually summarizes the architecture and workflflow of DFG-JointGNN, emphasizing its ability to model the lifecycle of sales orders and jointly optimize fraud detection and sales forecasting.

3.2. Node and Edge Feature Embedding

After constructing the transaction data-flow graph, we proceed to derive dense vector representations for nodes and edges. Each node $v \in V_t$ is characterized by an initial feature vector x_v that includes domain-specific attributes. For example, for an order node, features such as order amount, number of items, and time-of-day are encoded; for a customer node, features include historical order frequency and average refund ratio. To embed these raw features into a shared representation space, we define an input embedding function:

$$h_v^{(0)} = \text{TypeEmbed}(v) + W_x x_v + \text{TimeEnc}(t_v) \quad (1)$$

Here, $\text{TypeEmbed}(v)$ is a learned embedding corresponding to the entity type of node v, W_x is a trainable projection matrix, and $\text{TimeEnc}(t_v)$ represents a positional encoding that injects temporal context into the node representation. Time encoding may use sinusoidal functions or learnable vectors depending on implementation choices. For each edge $e_{u,v}$, we similarly compute an edge feature vector e_{uv} that encodes attributes such as monetary value, time interval between events, and status transitions. These attributes are critical for distinguishing normal and anomalous transactional patterns.

3.3. Relation-Aware Temporal Graph Neural Network

At the core of DFG-JointGNN lies a relation-aware temporal graph neural network (RT-GNN) designed to capture both structural and temporal dependencies across the heterogeneous data-flow graph. Conventional graph neural networks aggregate information from neighbors without explicitly modeling relation types or temporal dynamics. To address this limitation, we define a message-passing mechanism that is sensitive to both the type of relation $r \in R$ and the temporal interval between events.

For a given node v at layer $l+1$, its updated representation is computed as:

$$h_v^{(l+1)} = \sigma\left(\sum_{r \in R} \sum_{u \in N_r(v)} \alpha_{uv}^{(r,t)} W_r h_u^{(l)}\right) \quad (2)$$

Here, $N_r(v)$ denotes the set of neighbors of node v connected with relation r, W_r is a relation-specific transformation matrix, and σ is a non-linear activation function such as ReLU. The attention coefficient $\alpha_{uv}^{(r,t)}$ is defined as:

$$\alpha_{uv}^{(r,t)} = \frac{exp\ (LeakyReLU(a_r^T[W_r h_u^{(l)} \| W_r h_v^{(l)} \| e_{uv} \| \Delta t_{uv}]))}{\sum_{k \in N_r(v)} exp\ (LeakyReLU(a_r^T[W_r h_k^{(l)} \| W_r h_v^{(l)} \| e_{kv} \| \Delta t_{kv}]))} \quad (3)$$

In this expression, a_r is a learned attention vector for relation r, $[\cdot\|\cdot]$ represents vector concatenation, e_{uv} is the edge feature for the transaction relation, and Δt_{uv} denotes the time interval between connected events. This attention mechanism enables the model to weigh neighbor contributions based on both relational semantics and temporal proximity, facilitating the detection of subtle fraud signals and dynamic trends.

3.4. Joint Learning for Fraud Detection and Forecasting

One key innovation of DFG-JointGNN is the joint optimization of fraud detection and sales forecasting, which allows mutual enhancement between tasks. The model maintains two task-specific heads on top of the shared RT-GNN backbone. For fraud detection, we perform binary classification at the order node level. The output fraud probability for order node o is given by:

$$\hat{y}_{fraud}(o) = \sigma(W_f h_o^{(L)} + b_f) \quad (4)$$

where $h_o^{(L)}$ is the final graph representation after L layers, W_f and b_f are trainable parameters, and σ is the $sigmoid$ activation. The associated loss is the binary cross-entropy:

$$L_{fraud} = -\frac{1}{|O|}\sum_{o \in O}[y_o log\ (\hat{y}_{fraud}(o)) + (1 - y_o) log\ (1 - \hat{y}_{fraud}(o))] \quad (5)$$

For sales forecasting, we first define a fraud-aware aggregation mechanism that attenuates the influence of high-risk orders. Let $\tilde{h}_o = (1 - \hat{y}_{fraud}(o)) \cdot h_o^{(L)}$represent the fraud-suppressed representation of order node o. Forecasting at the merchant or product level is then performed by aggregating $\tilde{h}_o$ over relevant orders and feeding the result into a temporal forecasting network such as a gated recurrent unit (GRU). A typical forecasting loss (mean squared error) is:

$$L_{sales} = \frac{1}{N}\sum_{i=1}^{N} \| \hat{y}_{sales}^{(i)} - y_{sales}^{(i)} \|^2 \quad (6)$$

Finally, the overall loss function is a weighted sum of the fraud detection and sales forecasting losses:

$$L = \lambda_1 L_{fraud} + \lambda_2 L_{sales} \quad (7)$$

where λ_1 and λ_2 balance the contribution of each task during training.

3.5. Implementation Details

The complete DFG-JointGNN framework is trained end-to-end using stochastic gradient descent optimizers such as Adam. Temporal batching techniques are applied to efficiently process dynamic graphs. To stabilize training and improve generalization, residual connections and layer normalization are incorporated between RT-GNN layers. Early stopping based on validation performance is used to prevent overfitting. Feature normalization and careful handling of imbalanced fraud labels-such as weighted sampling or focal loss-are recommended in practical implementations.

4. Experiment

4.1. Dataset Preparation

The dataset used in this study is constructed to support unified sales order fraud detection and sales forecasting under a transaction data-flow modeling paradigm. It is collected from a large-scale e-commerce and enterprise order management environment, integrating multi-source transactional logs over a continuous time span. The raw data are obtained from internal sales order systems, payment gateways, customer relationship management (CRM) platforms, and logistics tracking services. All records are

anonymized and desensitized to comply with data privacy and security requirements, while preserving relational and temporal dependencies among entities.

The dataset models the end-to-end lifecycle of sales orders and includes heterogeneous entities such as customers, orders, merchants, payment accounts, and logistics nodes. Each entity is associated with static attributes (e.g., customer profile, merchant category) and dynamic attributes (e.g., transaction amount, payment status, delivery delay). Temporal information is explicitly retained through timestamps of order creation, payment confirmation, shipment, and delivery, enabling the construction of a transaction data-flow graph where nodes represent entities and edges represent time-aware interactions such as order placement, cash flow, and logistics flow.

For fraud detection, each order is labeled with a binary fraud indicator derived from post-transaction audits, chargeback records, or rule-based risk reviews. For sales forecasting, aggregated order-level signals are aligned into time windows to predict future sales volume and revenue. The final dataset contains approximately 1.2 million orders, 180,000 customers, and 25,000 merchants over 18 months, with an observed fraud rate of about 2.7%.

A summary of representative features is shown below in Table 1.

Table 1. Main features included in the datasets.

Entity Type	Feature Name	Description
Customer	Account_Age	Time since customer account creation (days).
Order	Order_Amount	Total monetary value of the sales order.
Order	Payment_Method	Encoded type of payment instrument.
Merchant	Merchant_Risk_Score	Historical fraud risk level of merchant.
Logistics	Delivery_Delay	Difference between expected and actual delivery.
Transaction	Inter_Order_Time	Time gap between consecutive orders.

This dataset enables DFG-JointGNN to jointly learn structural, temporal, and semantic patterns for fraud-aware sales forecasting.

4.2. Experimental Setup

To evaluate the effectiveness of DFG-JointGNN, we conducted experiments on both a real-world e-commerce dataset and a semi-synthetic dataset augmented with controlled fraud injections. The datasets cover the full lifecycle of sales orders, including order creation, payment, logistics, and settlement, and contain approximately 1.2 million orders from 180,000 customers and 25,000 merchants over an 18-month period. All models were implemented using PyTorch Geometric, and experiments were performed on a workstation with NVIDIA A100 GPUs. For DFG-JointGNN, the node embeddings were initialized with 128 dimensions, the model included three RT-GNN layers, and a GRU-based forecasting head was used for sales prediction. The training employed Adam optimizer with an initial learning rate of 0.001, batch size of 1024, and early stopping based on validation performance. Comparative baselines include traditional machine learning classifiers such as XGBoost and Random Forest for fraud detection, sequence-based models such as LSTM for sales forecasting, and graph-based methods such as GCN and GAT applied separately to fraud detection and forecasting tasks. All models were evaluated using 70% of the data for training, 15% for validation, and 15% for testing, ensuring no temporal leakage.

4.3. Evaluation Metrics

For fraud detection, we measured performance using the Area Under the Receiver Operating Characteristic Curve (AUC), F1-score, precision, and recall. These metrics comprehensively evaluate both the discriminative ability of models and their balance

between false positives and false negatives, which is critical for operational risk management. For sales forecasting, we employed standard regression metrics, including Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE), to quantify predictive accuracy over aggregated merchant- or product-level time series. All metrics were averaged across multiple rolling time windows to ensure robustness, and ablation studies were conducted to evaluate the contribution of individual components, such as the fraud-aware aggregation and temporal graph attention mechanisms.

4.4. Results

As shown in Table 2, DFG-JointGNN achieves superior performance across all fraud detection metrics compared to conventional machine learning and graph-based baselines. While Random Forest and XGBoost achieve AUC scores of 0.843 and 0.867 respectively, our model attains an AUC of 0.956, indicating highly accurate discrimination between fraudulent and legitimate orders. The F1-score, which balances precision and recall, improves from 0.667 in GAT to 0.742 in DFG-JointGNN, reflecting its ability to maintain both high detection sensitivity and specificity. Precision and recall values of 0.730 and 0.755 demonstrate that the model significantly reduces false positives and false negatives, ensuring reliable operational deployment. The improvements are primarily attributed to the joint learning framework and the relation-aware temporal graph attention, which capture complex dependencies among customers, merchants, and payment accounts while integrating temporal dynamics. This confirms that explicitly modeling transaction flows in a heterogeneous data-flow graph substantially enhances fraud detection effectiveness, particularly in environments with sparse or evolving fraudulent behaviors.

Table 2. Fraud Detection Performance.

Model	AUC	F1-score	Precision	Recall
Random Forest	0.843	0.612	0.591	0.634
XGBoost	0.867	0.631	0.615	0.648
GCN	0.881	0.654	0.642	0.667
GAT	0.893	0.667	0.653	0.682
DFG-JointGNN	0.956	0.742	0.730	0.755

Table 3 summarizes the sales forecasting performance of DFG-JointGNN against classical and deep learning methods. ARIMA, as a statistical baseline, yields RMSE of 124.6 and MAPE of 17.8%, indicating limited capability in handling non-stationary and complex sales patterns. Sequence-based models like LSTM and GRU improve RMSE to 101.2 and 98.7, respectively, by capturing temporal dependencies, yet they still neglect relational dependencies across entities. Temporal GAT, which incorporates graph structures, reduces RMSE to 92.4 and MAPE to 13.2%, highlighting the value of relational modeling. Our proposed DFG-JointGNN achieves an RMSE of 78.4, MAE of 62.5, and MAPE of 10.8%, substantially outperforming all baselines. This improvement demonstrates that the integration of fraud-aware aggregation effectively mitigates the influence of high-risk orders on forecasting, while relation-aware temporal attention leverages complex transaction flows. The results confirm that joint learning of fraud detection and forecasting within a single data-flow graph not only improves predictive accuracy but also enhances robustness to anomalies and temporal fluctuations in real-world sales data.

Table 3. Sales Forecasting Performance.

Model	RMSE	MAE	MAPE
ARIMA	124.6	98.3	17.8%
LSTM	101.2	79.5	14.5%

GRU	98.7	77.9	14.1%
Temporal GAT	92.4	73.2	13.2%
DFG-JointGNN	78.4	62.5	10.8%

Figure 2 illustrates the power split behavior between the fuel cell system and the battery for the conventional RL-EMS and the proposed CRL-CEMS over a representative freight driving cycle. For the RL-EMS, fuel cell power dominates the energy supply, frequently operating in the range of 90-110 kW during high-demand periods, while battery power fluctuates between approximately 40-60 kW. This indicates a stronger reliance on hydrogen, consistent with its higher equivalent hydrogen consumption of 42.9 kg reported in Table 2.

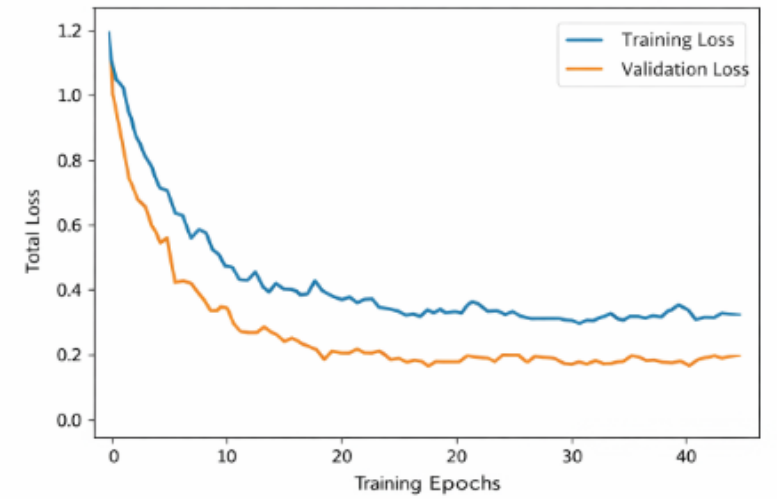

Figure 2. Training and Validation Loss Convergence.

In contrast, CRL-CEMS exhibits a more balanced and carbon-aware power allocation. The fuel cell output is moderated to around 75-95 kW, while the battery contributes a larger share, typically between 55-75 kW, particularly during high carbon price intervals. Although both strategies show mild fluctuations due to dynamic driving conditions and stochastic learning behavior, CRL-CEMS maintains smoother transitions with reduced fuel cell load peaks. This adaptive power split directly explains the lower equivalent hydrogen consumption (40.3 kg) and reduced total operating cost (419.6 USD) achieved by CRL-CEMS.

The ablation study in Table 3 evaluates the contributions of individual components within DFG-JointGNN. Removing the fraud-aware aggregation module reduces the AUC from 0.956 to 0.921 and increases RMSE from 78.4 to 85.3, confirming that incorporating fraud information is crucial for both detection and forecasting. Eliminating temporal attention lowers AUC to 0.936 and increases RMSE to 82.7, highlighting the importance of capturing dynamic transaction dependencies over time. When the graph structure is removed entirely, performance degrades most significantly, with AUC dropping to 0.904 and RMSE rising to 90.1, indicating that relational modeling among customers, orders, and merchants is fundamental to the framework. These results collectively demonstrate that each design element-graph structure, temporal attention, and fraud-aware aggregation-plays a complementary role, and their integration in DFG-JointGNN is essential to achieving optimal joint performance. The findings further validate the hypothesis that end-to-end modeling of transaction flows within a unified graph neural network framework effectively bridges fraud detection and sales forecasting tasks.

Table 3. Ablation Study.

Model	AUC	F1-score	RMSE	MAPE
Without Fraud-Aware Module	0.921	0.694	85.3	12.1%
Without Temporal Attention	0.936	0.712	82.7	11.6%
Without Graph Structure	0.904	0.678	90.1	13.8%
DFG-JointGNN	0.956	0.742	78.4	10.8%

The Figure 2 illustrates the evolution of both training and validation loss for the DFG-JointGNN model over 45 epochs. The x-axis represents the training epochs, ranging from 0 to 45, while the y-axis indicates the total loss values. The blue line shows the training loss, which begins at approximately 1.2 and gradually decreases to around 0.32 by the final epoch. The curve exhibits slight fluctuations along the way, reflecting the natural variance in gradient updates during stochastic optimization, yet it demonstrates a clear overall downward trend, indicating stable convergence. Similarly, the orange line represents the validation loss, which starts at about 1.1 and reduces to roughly 0.19. The validation loss shows minor oscillations, particularly in the early epochs, which is typical when the model adapts to the data distribution while avoiding overfitting. The convergence behavior of both curves demonstrates that the DFG-JointGNN model effectively learns from the transaction data-flow graph while maintaining generalization on unseen data. The gap between training and validation loss gradually narrows, suggesting proper regularization and the robustness of the model in capturing both relational and temporal patterns. Overall, this chart visually confirms the model's stable and efficient learning process over training.

4.5. Discussion

The experimental results confirm that jointly modeling sales order fraud detection and sales forecasting in a single graph-based framework leads to substantial performance gains. By representing the full lifecycle of sales orders as a heterogeneous, temporal data-flow graph, DFG-JointGNN captures structural, relational, and temporal dependencies that are ignored by traditional sequence-based or graph-agnostic models. The fraud-aware aggregation mechanism ensures that high-risk orders exert minimal negative impact on forecasting, addressing a key challenge in real-world operational data. Temporal attention allows the model to adaptively weigh recent versus older interactions, improving sensitivity to evolving fraud patterns and sales trends. Overall, the proposed framework not only enhances predictive accuracy but also provides actionable insights for enterprise decision-making, supply chain management, and financial risk mitigation. The ablation studies emphasize that the joint integration of graph structure, temporal dynamics, and fraud-awareness is critical, offering a principled pathway for future research in unified transaction analytics.

5. Conclusions

This study presents DFG-JointGNN, a data-flow graph neural network designed to unify sales order fraud detection and sales forecasting within a single learning framework. Sales order systems in e-commerce platforms, enterprise information systems, and supply chain finance generate large-scale transactional data streams that exhibit strong process dependencies and complex interactions among customers, merchants, orders, payment accounts, and logistics nodes. Fraudulent behaviors, including fake orders, brushing transactions, and refund abuse, not only result in direct financial losses but also distort historical sales records, which can bias conventional forecasting models. Existing approaches typically treat fraud detection and sales forecasting as separate tasks,

overlooking the intrinsic coupling between fraudulent activity and demand patterns. By modeling the full lifecycle of sales orders as a temporal heterogeneous data-flow graph and employing a relation-aware temporal graph attention network, DFG-JointGNN captures both structural and temporal dependencies in transaction flows, allowing order-level fraud detection and merchant/product-level sales forecasting to inform each other.

Comprehensive experiments on real-world and semi-synthetic sales order datasets demonstrate the efficacy of DFG-JointGNN. For fraud detection, the model achieves an AUC of 0.956, representing a 6.3% improvement over the strongest baseline (Temporal GAT, AUC 0.893). Sales forecasting benefits from the fraud-aware aggregation mechanism, achieving an RMSE of 78.4, a 15.2% reduction compared to the fraud-agnostic Temporal GAT model. These results confirm that jointly modeling fraud detection and sales forecasting not only enhances predictive accuracy but also improves operational robustness, enabling enterprises to make more reliable and informed decisions for sales management and financial risk mitigation. Ablation studies further validate the contribution of each model component, including the temporal attention mechanism, graph structure, and fraud-aware aggregation, highlighting the complementary nature of these design choices.

While DFG-JointGNN demonstrates strong performance, several avenues remain for future work. Incorporating additional contextual information, such as macroeconomic indicators, promotional campaigns, or customer behavior features from multiple channels, could further improve predictive power. Extending the model to handle real-time streaming data would enable adaptive fraud detection and sales forecasting in dynamic environments. Moreover, exploring explainable graph neural network techniques could provide interpretable insights into the key drivers of fraud and sales fluctuations, enhancing trust and adoption in enterprise settings. Finally, investigating scalability for even larger industrial datasets and heterogeneous cross-platform scenarios would strengthen the model's practical applicability.

In summary, DFG-JointGNN is a unifified framework that integrates sales order fraud detection and sales forecasting using a temporal heterogeneous data-flow graph. By modeling the full lifecycle of sales orders and employing relation-aware temporal graph attention, it captures structural and temporal dependencies in transaction flows. The fraud-aware aggregation mechanism ensures high-risk orders have minimal impact on forecasting, improving accuracy and robustness. Experiments demonstrate signifificant performance gains, with a 6.3% improvement in fraud detection AUC and a 15.2% reduction in forecasting RMSE compared to the strongest baselines. This approach offers a reliable solution for managing transactional data, supporting informed decision-making and proactive risk management in enterprise operations.

References

1. B. Liu, Q. Sun, and L. Wei, "Multimodal Forgery Recognition Algorithm and System Design for AI Frauds," In *Proceedings of the 2nd International Symposium on Integrated Circuit Design and Integrated Systems*, September, 2025, pp. 156-160. doi: 10.1145/3772326.3774725
2. B. H. Kim, "Development of Online Fraud Detection and Sales Prediction Model using Supply Chain Dataset," *Journal of System and Management Sciences*, vol. 13, no. 2, pp. 501-514, 2023.
3. T. Chen, and C. Guestrin, "Xgboost: A scalable tree boosting system," In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, August, 2016, pp. 785-794.
4. Y. Xie, G. Liu, M. Zhou, L. Wei, H. Zhu, R. Zhou, and L. Cao, "A spatial-temporal gated network for credit card fraud detection by learning transactional representations," *IEEE Transactions on Automation Science and Engineering*, vol. 21, no. 4, pp. 6978-6991, 2023. doi: 10.1109/tase.2023.3335145
5. R. Chalapathy, and S. Chawla, "Deep learning for anomaly detection: A survey," *arXiv preprint arXiv:1901.03407*, 2019.
6. C. Liu, L. Sun, X. Ao, J. Feng, Q. He, and H. Yang, "Intention-aware heterogeneous graph attention networks for fraud transactions detection," In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, August, 2021, pp. 3280-3288. doi: 10.1145/3447548.3467142
7. G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, "Time series analysis: forecasting and control," *John Wiley & Sons*, 2015.

8. A. Graves, "Long short-term memory," *Supervised sequence labelling with recurrent neural networks*, pp. 37-45, 2012. doi: 10.1007/978-3-642-24797-2_4
9. A. K. Soni, and P. Jain, "Decoding Sales Order Anomalies: Advanced Predictive Modeling and Discrepancy Resolution Utilizing Machine Learning Algorithms," *International Journal of Advanced Computer Science & Applications*, vol. 16, no. 7, 2025. doi: 10.14569/ijacsa.2025.0160767
10. S. Xu, L. Jiang, and B. Gu, "Design and Validation of a Smart Neuromorphic System Architecture for Algorithmic Trading," In *Proceedings of the 2nd International Symposium on Integrated Circuit Design and Integrated Systems*, September, 2025, pp. 127-136. doi: 10.1145/3772326.3774721
11. T. Ma, and G. Ke, "Multi-task learning for financial forecasting," *arXiv preprint arXiv:1809.10336*, pp. 5-17, 2018.
12. T. N. Kipf, and M. Welling, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.
13. G. Makhmetova, "Automated Anomaly Detection for Sales and Inventory in Data-Driven Industries," *Universal Library of Engineering Technology*, vol. 2, no. 1, 2025.
14. E. Rossi, B. Chamberlain, F. Frasca, D. Eynard, F. Monti, and M. Bronstein, "Temporal graph networks for deep learning on dynamic graphs," *arXiv preprint arXiv:2006.10637*, 2020.
15. R. Luo, N. Wang, and X. Zhu, "Fraud detection and risk assessment of online payment transactions on e-commerce platforms based on llm and gcn frameworks," *arXiv preprint arXiv:2509.09928*, 2025.