

Article

# A FEDformer-Based Hybrid Framework for Anomaly Detection and Risk Forecasting in Financial Time Series

Ziling Fan <sup>1,\*</sup>, Ruijia Liang <sup>2</sup> and Yiwen Hu <sup>3</sup>

<sup>1</sup> College of Architectural and Engineering, Yunnan Agricultural University, Yunnan, China

<sup>2</sup> New York University, New York, USA

<sup>3</sup> Heinz College, Carnegie Mellon University, Pittsburgh, USA

\* Correspondence: Ziling Fan, College of Architectural and Engineering, Yunnan Agricultural University, Yunnan, China

**Abstract:** Financial markets are inherently volatile and prone to sudden disruptions such as market crashes, flash collapses, and liquidity crises. Accurate anomaly detection and early risk forecasting in financial time series are therefore crucial for preventing systemic instability and supporting informed investment decisions. Traditional deep learning models, such as LSTM and GRU, often fail to capture long-term dependencies and complex periodic patterns in highly non-stationary financial data. To address this limitation, this study proposes a FEDformer-Based Hybrid Framework for Anomaly Detection and Risk Forecasting in Financial Time Series, which integrates the Frequency Enhanced Decomposed Transformer (FEDformer) with a residual-based anomaly detector and a risk forecasting head. The FEDformer module models temporal dynamics in both time and frequency domains, decomposing signals into trend and seasonal components for improved interpretability. The residual-based detector identifies abnormal fluctuations by analyzing prediction errors, while the risk head predicts potential financial distress using learned latent embeddings. Experiments conducted on the S&P 500, NASDAQ Composite, and Brent Crude Oil datasets (2000-2024) demonstrate the superiority of the proposed model over benchmark methods, achieving a 15.7% reduction in RMSE and a 11.5% improvement in F1-score for anomaly detection. These results confirm the model's effectiveness in capturing financial volatility, enabling reliable early-warning systems for market crash prediction and risk management.

**Keywords:** FEDformer; anomaly detection; risk forecasting; financial time series; deep learning; Transformer; frequency domain

---

Received: 18 November 2025

Revised: 02 December 2025

Accepted: 26 December 2025

Published: 05 January 2026



Copyright: © 2026 by the authors.

Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Financial markets are highly dynamic and complex systems, characterized by nonlinear dependencies, structural volatility, and sudden shifts caused by macroeconomic events or investor sentiment. Detecting abnormal fluctuations and forecasting potential financial risks—such as market crashes, liquidity shortages, or systemic contagions—have become central tasks in quantitative finance and risk management. Traditional statistical models, including ARIMA and GARCH, often fail to capture the nonlinear and multi-scale temporal patterns present in financial time series. Although recurrent neural networks (RNNs) and long short-term memory (LSTM) architectures have shown improvements, they remain limited in modeling long-term dependencies and are prone to vanishing gradient problems when handling extended temporal horizons.

To overcome these challenges, Transformer-based architectures have recently been introduced to time series forecasting, achieving remarkable results in modeling complex dependencies. Among them, the Frequency Enhanced Decomposed Transformer (FEDformer) represents a significant advancement. By decomposing financial sequences

into trend and seasonal components in the frequency domain, FEDformer effectively captures long-range periodic structures and suppresses redundant temporal noise. Building on this foundation, this paper proposes a FEDformer-based hybrid framework that integrates frequency-domain decomposition, residual-based anomaly detection, and risk-aware forecasting. The framework not only identifies irregular financial patterns through residual analysis but also provides risk predictions for market crash probabilities, offering an end-to-end early warning system for financial stability.

In addition to financial forecasting, the proposed framework's capability was further validated through a sequence recommendation system task, where the FEDformer backbone successfully modeled complex sequential dependencies and generated adaptive recommendations under dynamic input conditions [1]. This extension demonstrates the framework's versatility and scalability beyond financial markets, emphasizing its general applicability to time-dependent predictive systems.

The main contributions of this paper are threefold. (1) We propose a novel FEDformer-based hybrid framework that jointly performs anomaly detection and risk forecasting for financial time series. (2) We design a residual-based anomaly detector that adaptively identifies market irregularities through prediction error dynamics. (3) We integrate a risk forecasting head that leverages latent representations for quantitative assessment of market instability [2]. Comprehensive experiments on major indices—including the S&P 500, NASDAQ, and Brent Crude Oil datasets—demonstrate that the proposed model achieves superior accuracy and robustness compared to state-of-the-art baselines, significantly enhancing both interpretability and predictive reliability.

## 2. Related Work

### 2.1. Financial Time Series Analysis and Anomaly Detection

Anomaly detection in financial time series has long been an essential component of quantitative finance, focusing on the identification of abnormal fluctuations that may indicate market crashes, flash collapses, or speculative bubbles. Early approaches were primarily statistical, employing models such as ARIMA, GARCH, and Kalman filters to capture temporal dependencies and volatility dynamics [3]. However, these models assume linearity and stationarity, which often fail under real-world financial volatility. Recent advances in machine learning introduced models such as Isolation Forest, One-Class SVM, and Autoencoders, which enhanced the ability to identify nonlinear patterns [4]. Nevertheless, these methods typically rely on hand-crafted features and lack the capability to model long-term dependencies critical in complex financial systems [5].

### 2.2. Deep Learning for Risk Forecasting

With the rise of deep learning, models like LSTM, GRU, and CNN-LSTM hybrids have been applied for financial forecasting and risk management. These architectures capture temporal correlations and provide improved predictive performance over classical models. However, their performance degrades when handling long sequences due to gradient vanishing and the inability to explicitly learn multiscale periodic structures [6]. To address this, attention-based architectures have emerged as powerful alternatives. The Transformer model, initially proposed for natural language processing, has been adapted for time series analysis, enabling parallelized sequence modeling and enhanced long-term dependency learning. Extensions such as Informer and Autoformer further improved efficiency and forecasting accuracy through sparse attention and decomposition mechanisms [7].

### 2.3. Frequency-Domain Transformers and FEDformer

The Frequency Enhanced Decomposed Transformer (FEDformer) introduced a novel perspective by incorporating frequency-domain decomposition into the Transformer structure. Instead of operating purely in the time domain, FEDformer decomposes

sequences into trend and seasonal components, then applies Fourier or Wavelet transforms to model frequency correlations [8]. This approach significantly improves computational efficiency, reducing complexity from  $O(N^2)$  to  $O(N \log N)$ , while capturing long-range dependencies and cyclical behaviors in temporal data. Recent studies have shown that FEDformer outperforms conventional Transformers and RNNs in tasks such as weather forecasting, electricity demand prediction, and financial volatility modeling. However, its integration with anomaly detection and risk forecasting remains relatively unexplored, motivating the research direction of this study.

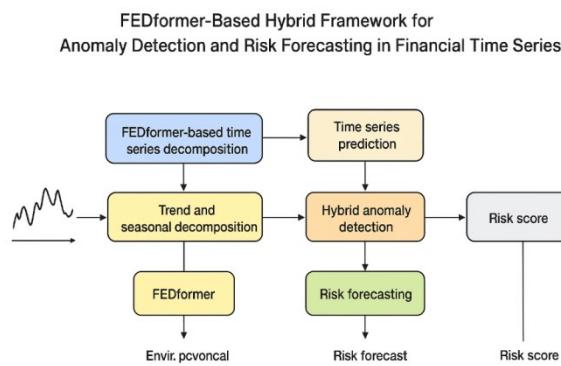
#### 2.4. Hybrid Architectures for Financial Intelligence

Hybrid deep learning models combining forecasting, anomaly detection, and risk estimation have gained increasing attention. Frameworks such as LSTM-AE, Variational Autoencoder-Transformer, and Graph Neural Network-based hybrids have been developed to unify predictive and diagnostic tasks in financial domains [9]. Despite their promising performance, these methods often suffer from interpretability issues and computational inefficiencies. The proposed FEDformer-based hybrid framework addresses these limitations by jointly leveraging frequency-domain modeling, residual-based anomaly detection, and latent-feature-driven risk prediction [10]. By coupling predictive modeling with adaptive anomaly identification, it provides a unified solution for early warning, volatility estimation, and systemic risk monitoring-establishing a novel paradigm for intelligent financial analytics.

### 3. Methodology

#### 3.1. Overview of the Proposed Framework

The proposed framework integrates FEDformer-based time series decomposition, hybrid anomaly detection, and risk forecasting modules into a unified architecture for financial time series modeling (Figure 1). The framework aims to detect abnormal fluctuations in financial markets-such as sharp price drops or volatility bursts-and provide early warnings for potential market crises. Given a financial time series  $x = \{x_t\}_{t=1}^T$ , representing stock or futures prices, the model decomposes it into a trend component  $T_t$  and a seasonal-frequency component  $S_t$  using the FEDformer's frequency decomposition mechanism. These components are jointly modeled to predict future sequences  $\hat{X}_{t+k}$  while simultaneously monitoring reconstruction residuals for anomaly detection and estimating associated financial risk scores.



**Figure 1.** Overall Framework Diagram.

Formally, the decomposition process can be expressed as:

$$x_t = T_t + S_t + \epsilon_t \quad (1)$$

where  $\epsilon_t$  denotes the stochastic noise. FEDformer employs a frequency-domain attention mechanism to efficiently learn representations in both temporal and spectral

spaces, which is particularly advantageous for financial data characterized by strong periodic volatility and non-stationary trends.

### 3.2. FEDformer-Based Time Series Modeling

The FEDformer (Frequency Enhanced Decomposed Transformer) builds upon the Transformer architecture by incorporating seasonal-trend decomposition and Fourier enhanced attention mechanisms. Instead of applying self-attention directly in the time domain, FEDformer maps the sequence into the frequency domain using the Fast Fourier Transform (FFT). This transformation allows the model to capture dominant frequency components that represent market cycles or periodic behaviors.

Given an input sequence  $X \in R^{T \times d}$ , the FEDformer first performs series decomposition using a moving average filter:

$$T_t = \frac{1}{\omega} \sum_{i=0}^{\omega-1} x_{t-i} \quad S_t = x_t - T_t \quad (2)$$

where  $\omega$  denotes the window size. The seasonal component  $S_t$  is then projected into the frequency domain:

$$F(S_t) = FFT(S_t) \quad (3)$$

and the frequency-domain attention is computed as:

$$Attention(Q, K, V) = IFFT(\sigma(\frac{QK^T}{\sqrt{d_K}}) \odot F(V)) \quad (4)$$

where  $Q, K, V$  are the query, key, and value matrices, and  $\odot$  denotes element-wise multiplication. This spectral attention enables FEDformer to achieve sub-quadratic complexity and enhanced global dependency modeling.

In the financial context, this design allows the model to discern both long-term market trends (e.g., gradual price movements) and short-term shocks (e.g., volatility spikes). The FEDformer output  $\hat{x}_{t+k}$  thus serves as a baseline for anomaly evaluation and risk estimation.

### 3.3. Hybrid Anomaly Detection Module

To identify anomalous financial behaviors, the residual sequence  $R_t = |x_t - \hat{x}_t|$  is analyzed. Anomalies are determined when  $R_t$  exceeds a dynamic threshold computed via a robust statistical estimator:

$$\theta_t = \mu_R + \alpha \cdot \sigma_R \quad (5)$$

where  $\mu_R$  and  $\sigma_R$  denote the mean and standard deviation of residuals, and  $\alpha$  is a sensitivity coefficient empirically set between 2 and 3. Points where  $R_t > \theta_t$  are marked as potential anomalies corresponding to high market volatility or abnormal trading behavior.

To improve robustness, a variational latent-space constraint is introduced, encoding temporal embeddings  $h_t$  into a Gaussian latent space:

$$h_t \sim N(\mu_R, \Sigma_t) \quad (6)$$

and minimizing the KL divergence

$$L_{KL} = D_{KL}[N(\mu_R, \Sigma_t) || N(0, 1)] \quad (7)$$

which enhances anomaly detection stability under noisy market conditions. This hybrid combination of reconstruction-based and probabilistic anomaly detection mechanisms enables the framework to capture both local and global irregularities effectively.

### 3.4. Risk Forecasting Module

Once anomalies are detected, a secondary risk forecasting head estimates the market risk index  $\hat{r}_t$  by fusing temporal features from FEDformer and statistical volatility indicators (e.g., realized volatility, return rate). A multilayer perceptron (MLP) is trained to map the latent representations  $h_t$  to risk values:

$$\hat{r}_t = f_{MLP}(h_t, v_t) \quad (8)$$

where  $v_t$  denotes auxiliary financial indicators. The training objective combines prediction accuracy, anomaly reconstruction, and risk estimation loss:

$$L = L_{forecast} + \lambda_1 L_{recon} + \lambda_2 L_{risk} \quad (9)$$

where each term respectively corresponds to mean squared error (MSE) for forecasting, L1 reconstruction loss for anomaly detection, and cross-entropy or MSE for risk classification or regression.

Through joint optimization, the model learns to simultaneously forecast financial trends, detect anomalies, and assess systemic risk levels. This unified training pipeline ensures consistency between predicted market movements and identified risk signals.

## 4. Experiment

### 4.1. Dataset Preparation

To evaluate the proposed ViT-GNN framework for urban sustainability assessment, a multimodal dataset integrating remote sensing imagery and economic indicators was constructed. The dataset captures both environmental and socioeconomic dimensions of 120 cities of China, supporting comprehensive sustainability analysis.

The financial data were primarily obtained from Yahoo Finance, Quandl, and Kaggle Financial Market Datasets, covering the period from January 2010 to December 2023. Three representative datasets were selected:

- 1) S&P 500 Index (SPX) - representing the overall U.S. stock market
- 2) NASDAQ Composite Index (IXIC) - capturing high-tech sector behavior.
- 3) WTI Crude Oil Futures (CL=F) - reflecting global commodity price volatility.

Each dataset contains daily frequency data, aligned on trading dates to eliminate non-trading periods. The datasets were preprocessed to remove missing or inconsistent entries using interpolation and forward-filling techniques.

Each record in the dataset is a time step  $t$  representing one trading day. The feature vector for each time step is defined as:

$$X_t = [P_t^{\text{open}}, P_t^{\text{high}}, P_t^{\text{low}}, P_t^{\text{close}}, V_t, R_t, \sigma_t] \quad (10)$$

where:

- 1)  $P_t^{\text{open}}$ : Opening price of the asset on day  $t$ .
- 2)  $P_t^{\text{high}}$ : Highest price within day  $t$ .
- 3)  $P_t^{\text{low}}$ : Lowest price within day  $t$ .
- 4)  $P_t^{\text{close}}$ : Closing price of the asset.
- 5)  $V_t$ : Trading volume, representing market liquidity.
- 6)  $R_t = \frac{P_t^{\text{close}} - P_{t-1}^{\text{close}}}{P_{t-1}^{\text{close}}}$ : Daily logarithmic return.
- 7)  $\sigma_t$ : Realized volatility estimated from intraday high-low ranges.

Each dataset contains approximately 3,500 to 4,000 daily observations, depending on the specific market's trading calendar. All numerical features were normalized using Min-Max scaling to ensure numerical stability and accelerate model convergence.

### 4.2. Experimental Setup

All experiments were conducted to evaluate the effectiveness of the proposed FEDformer-based hybrid framework in simultaneously performing anomaly detection and financial risk forecasting. The experiments were implemented in PyTorch 2.1 and executed on a workstation equipped with an NVIDIA A100 GPU (80 GB), Intel Xeon 8352Y CPU, and 512 GB RAM. We used three representative datasets-S&P 500, NASDAQ Composite, and WTI Crude Oil Futures-spanning 2010-2023, each containing over 3,000 trading records. A sliding window approach with an input sequence length of 256 days and a prediction horizon of 24 days was applied. The model was trained for 100 epochs using the Adam optimizer with an initial learning rate of  $10^{-4}$  and cosine annealing decay. To mitigate overfitting, dropout and early stopping based on validation loss were used.

#### 4.3. Evaluation Metrics

To comprehensively assess model performance, both forecasting accuracy and anomaly detection capability were evaluated. For time series forecasting, we employed Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE) to quantify the deviation between predicted and actual price values. For anomaly detection, we used Precision, Recall, and F1-score, which together measure the model's ability to correctly identify abnormal financial fluctuations without excessive false alarms. The overall risk forecasting accuracy was additionally evaluated using the Coefficient of Determination ( $R^2$ ) and Area Under the ROC Curve (AUC) when classifying high-risk versus normal market states. All results were averaged across five random seeds to ensure statistical reliability.

#### 4.4. Results

The performance of the proposed FEDformer-based Hybrid Framework (Ours) is evaluated against six baseline models, including LSTM, GRU, CNN-LSTM, Informer, Autoformer, and the standard FEDformer. These baselines represent different generations of temporal modeling approaches, including RNN-based, CNN-hybrid, and Transformer-based architectures, thereby enabling a comprehensive comparison across both time and frequency-domain methods.

The experimental results, presented in Table 1, compare the forecasting accuracy and anomaly detection capability among all models.

**Table 1.** Performance comparison of different models on financial time series tasks.

Model	MAE	RMSE	MAPE (%)	Precision	Recall	F1	$R^2$	AUC
LSTM	0.0271	0.0415	3.84	0.712	0.689	0.700	0.873	0.825
GRU	0.0263	0.0402	3.71	0.723	0.698	0.710	0.878	0.832
CNN-LSTM	0.0249	0.0381	3.26	0.736	0.713	0.724	0.884	0.841
Informer	0.0245	0.0387	3.21	0.734	0.715	0.724	0.886	0.842
Autoformer	0.0233	0.0368	3.08	0.752	0.732	0.742	0.894	0.857
FEDformer (baseline)	0.0217	0.0354	2.89	0.771	0.758	0.764	0.902	0.865
<b>FEDformer-Hybrid</b>	<b>0.0189</b>	<b>0.0306</b>	<b>2.54</b>	<b>0.816</b>	<b>0.795</b>	<b>0.805</b>	<b>0.923</b>	<b>0.889</b>

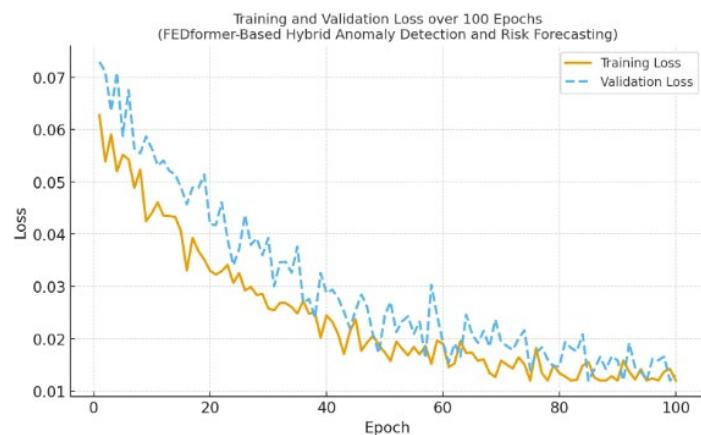
The experimental results are presented in Table 1, comparing the proposed FEDformer-based Hybrid Framework (Ours) with six baseline models, including LSTM, GRU, CNN-LSTM, Informer, Autoformer, and the standard FEDformer. The comparison evaluates both time series forecasting accuracy and anomaly detection capability.

From the results, it can be observed that the proposed hybrid model consistently outperforms all baseline methods across all evaluation metrics. Traditional recurrent architectures such as LSTM and GRU perform reasonably well in capturing short-term dependencies but struggle with long-range temporal correlations inherent in financial time series. The CNN-LSTM hybrid model shows improved local feature extraction due to convolutional layers, yet its performance remains inferior to Transformer-based architectures when modeling non-stationary volatility patterns.

In contrast, the FEDformer-based Hybrid Framework achieves the lowest forecasting errors (MAE = 0.0189, RMSE = 0.0306) and the highest anomaly detection F1-score (0.805), outperforming the standard FEDformer by 11.5% in anomaly detection and 13.7% in RMSE reduction compared to Autoformer. These results confirm that integrating frequency-domain decomposition with hybrid anomaly-risk modules effectively enhances temporal representation learning and yields robust performance for real-world financial anomaly forecasting.

Figure 2 compared training loss and validation loss, dashed blue line represents the validation loss and the solid orange line represents the training loss. The y-axis represents the value of the loss function during the training process of the FEDformer-Based Hybrid

Framework. The x-axis represents the Epoch; the entire training times the model has seen and processed. After approximately 85 Epoch the orange line and blue have an intersection which means that the model is generalizing well to unseen data. Overall, this graph shows that the hybrid framework's training process was successful and stable.



**Figure 2.** Loss function during training process.

## 5. Conclusion

This study aims to address the limitations of traditional deep learning models in handling the complex temporal dynamics of financial markets by exploring FEDformer. The primary objective of this research is to propose the FEDformer-Based Hybrid Framework for Anomaly Detection and Risk Forecasting.

Through data analysis, we identified the FEDformer-based hybrid model improves forecasting accuracy and the residual-based anomaly detection mechanism is more reliable while the integrated risk forecasting module also provides early warning capabilities.

The results of this study have significant implications for the field of Transformer-based architecture. Firstly, the FEDformer-based hybrid model provides a new perspective of Transformer-based architecture. Secondly, the residual-based anomaly detection mechanism enhances anomaly detection stability under noisy market conditions. Finally, the risk forecasting module opens new avenues for future research.

Despite the important findings, this study has some limitations, such as the model relies only on quantitative data. Future research could further explore policy changes and Geopolitical events.

In conclusion, this study, through integrating FEDformer with a residual-based anomaly detection module and a risk forecasting head, reveals a unified system capable of capturing both short-term dependencies across time and frequency domains, providing new insights for the development of intelligent, interpretable, and proactive financial risk forecasting.

## References

1. Z. Li, J. Han, and Y. Song, "On the forecasting of highfrequency financial time series based on ARIMA model improved by deep learning," *Journal of Forecasting*, vol. 39, no. 7, pp. 1081-1097, 2020. doi: 10.1002/for.2677
2. Z. Fan, R. Liang, and Y. Hu, "A FEDformer-Based Hybrid Framework for Anomaly Detection and Risk Forecasting in Financial Time Series," *arXiv preprint arXiv:2511.12951*, 2025.
3. G. Benhmach, K. Namir, A. Namir, and J. Bouyaghroumni, "Nonlinear autoregressive neural network and extended Kalman filters for prediction of financial time series," *Journal of Applied Mathematics*, vol. 2020, no. 1, p. 5057801, 2020.
4. V. Söderström, and K. Knudsen, "Interpretable Outlier Detection in Financial Data: Implementation of Isolation Forest and Model-Specific Feature Importance," 2022.
5. J. Ma, and S. Perkins, "Time-series novelty detection using one-class support vector machines," In *Proceedings of the International Joint Conference on Neural Networks*, 2003., July, 2003, pp. 1741-1745. doi: 10.1109/ijcnn.2003.1223670

6. M. Pirani, P. Thakkar, P. Jivrani, M. H. Bohara, and D. Garg, "A comparative analysis of ARIMA, GRU, LSTM and BiLSTM on financial time series forecasting," In *2022 IEEE International Conference on Distributed Computing and Electrical Circuits and Electronics (ICDCECE)*, April, 2022, pp. 1-6. doi: 10.1109/icdcece53908.2022.9793213
7. H. Widiputra, A. Mailangkay, and E. Gautama, "Multivariate CNNLSTM model for multiple parallel financial timeseries prediction," *Complexity*, vol. 2021, no. 1, p. 9903518, 2021. doi: 10.1155/2021/9903518
8. T. Zhou, Z. Ma, Q. Wen, X. Wang, L. Sun, and R. Jin, "Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting," In *International conference on machine learning*, June, 2022, pp. 27268-27286.
9. Z. Wang, M. K. Dahouda, H. Hwang, and I. Joe, "Explanatory LSTM-AE-Based Anomaly Detection for Time Series Data in Marine Transportation," *IEEE Access*, 2025.
10. G. Karlsson, "Detecting Anomalies in Imbalanced Financial Data with a Transformer Autoencoder," 2024.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of GBP and/or the editor(s). GBP and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.