Industry Science and Engineering Vol. 2 No. 10, 2025

Multivariate Time Series Anomaly Detection Based on Time-Domain and Frequency-Domain Fusion

Jiming Xu, Hanqi Liu, Jun Xu
College of Science, North China University of Science and Technology, Tangshan, China

Abstract: In complex multivariate time-series anomaly detection tasks, traditional methods often rely on time-domain modeling while neglecting frequency-domain information, leading to limited performance. To address this issue, this study proposes an innovative time-frequency anomaly detection framework, FTAD, which combines time-domain and frequency-domain features through adaptive graph attention and frequency-domain attention mechanisms, significantly improving anomaly detection performance for multivariate time-series data. the model first decomposes the time-series data using Exponential Moving Average (EMA), then processes the time-domain and frequency-domain features separately using Adaptive Graph Attention (AGAT) and Frequency-domain Attention (FreRA), respectively. Finally, the mutual inverse entropy weighting (REFusion) mechanism is used to dynamically and adaptively fuse time-domain and frequency-domain features. Experimental results show that FTAD outperforms existing methods on multiple datasets, demonstrating its effectiveness in complex time-series data.

Keywords: Industrial Internet of Things; Multivariate Time Series Data; Time-Domain Feature; Frequency-Domain Feature; Adaptively Fusion

1. Introduction

In these applications, time series anomaly detection, as a key technology, can effectively identify potential abnormal conditions of the system by mining abnormal behaviors that deviate from normal patterns in the data, thereby providing strong support for ensuring system stability and operational security[1] . For example, in industrial control systems, anomaly detection on sensor and network data can realize timely early warning of equipment faults and potential attacks [2] , and in multi-channel physiological signal analysis, real-time electrocardiogram waveform anomaly identification based on deep learning models helps to improve the accuracy and efficiency of diagnosis [3] .

Although anomaly detection technology has been widely applied in fields such as industrial monitoring and financial risk control, accurate identification of anomalies remains a challenging task when dealing with time series data with complex and diverse structures [4] . However, schemes based on a single domain all have inherent limitations. Time-domain methods, although good at characterizing short-term dependencies and nonlinear dynamics, often ignore spectral structures and are insufficient in representing strong periodicity, harmonic patterns, and cross-scale oscillations; frequency-domain methods, although able to highlight periodicity and repeated fluctuations, find it difficult to capture the local temporal context of sudden events or non-stationary anomalies, and performing Fourier transform on the whole sequence easily introduces noisy frequencies and causes spectral leakage and aliasing, thereby masking key patterns and interfering with effective information extraction [5] . For example, FEDformer [6] introduces Fourier enhanced blocks and wavelet enhanced blocks into the Transformer framework, mapping time series signals to the frequency domain, capturing key structures, and then returning to the time domain for prediction. F-SE-LSTM [7] combines FFT with the attention channel mechanism SENet to extract frequency-domain features, and uses LSTM to model cross-period frequency dependencies. Therefore, relying only on a single domain makes it difficult to comprehensively cover multiple types of anomalies.

We propose FTAD, an unsupervised time–frequency fusion anomaly detection framework that uses EMA-based decomposition, adaptive graph attention, improved frequency-domain attention, and entropy-based feature fusion, and

combines prediction and reconstruction errors to robustly detect complex anomalies, achieving superior performance on multiple public datasets.

2. FTAD Model

2.1 Problem Definition

Given a multivariate time series X, as shown in Eq. (1),

$$X = \{x_1, x_2, ..., x_T\}, x_t \in \mathbb{R}^d \quad (1)$$

where the two parameters respectively denote the dimension of sensors or indicators and the sequence length. the main task is, in an unsupervised setting, to determine whether there exist potential anomalies within a sliding time window. During the inference stage, the weighted sum of the two absolute errors, prediction error and reconstruction error, is used. A threshold is set, and the observations are evaluated according to the anomaly score; that is, when the anomaly score obtained for a time window exceeds the threshold, the window is marked as anomalous. Here, a binary indicator is used to represent whether the data point at time stamp ttt is anomalous (1 for anomalous and 0 for normal).

2.2 Overall Framework

This paper proposes a time–frequency adaptive fusion framework FTAD based on trend–seasonal decomposition and an uncertainty-driven mechanism. the overall architecture of FTAD is shown in Fig. 1 and mainly consists of five stages: time series decomposition, time- and frequency-domain analysis, feature fusion, prediction and reconstruction, and anomaly detection.

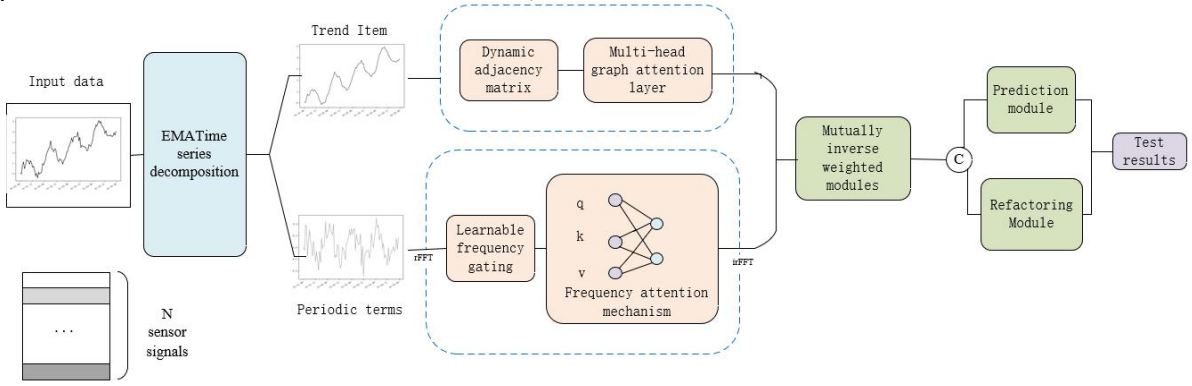

Figure 1. Overall Architecture of the Model

2.3 Adaptive Graph Attention Module

This paper proposes an adaptive graph attention mechanism (AGAT), which, during end-to-end training, simultaneously recovers the topology and learns attention, and uses a task-adaptive soft adjacency matrix to replace hard-coded graph structures, thereby more robustly characterizing implicit and time-varying channel dependencies.

Let the input be the windowed time-series block $H \in \mathbb{R}^{B\times C\times T}$ (batch size B, number of variables C, feature length T).

First, a trainable node embedding is learned for each variable $\mathbf{E} \in \eth^{C\times d}$, and perform row-wise vector normalization to obtain a similarity matrix, and then construct a dynamic adjacency matrix using a combination of tanh and ReLU functions:

$$A_{ij} = ReLU\left(\tanh\left(\mathbf{e}^i \cdot \mathbf{e}_j\right)\right) \quad (2)$$

where A_{ij} is a nonnegative, dense, and differentiable soft adjacency matrix, which is updated via backpropagation with respect to the downstream task loss and thus adapts to both the data domain and the objective function, enabling it to approximate the true time-varying topology without any external graph prior. Parallel multi-head graph attention is then applied on this soft graph. the input features are linearly projected, and edge scores are computed following the idea of GAT:

$$e_{ij} = LeakyReLU\left(a \left[\mathbf{h}_i' \| \mathbf{h}_j'\right]\right), a \in \eth^{2T_{out}} \quad (3)$$

and employ adjacency-gated masked normalization to suppress non-neighborhood interference, and obtain after multi-head aggregation

$$h_i' = \|_{k=1}^{H} \sigma\left(\sum_j \alpha_{ij}^{(k)} W^{(k)} h_j\right) \quad (4)$$

After concatenating the outputs of all heads along the feature dimension, a linear projection and a second-level single-head attention are applied to obtain the final output $\mathbf{H}^{\text{out}} \in \eth^{B\times C\times T}$.

This module restores the topological structure by dynamically learning a soft adjacency matrix and uses this structure as a prior for attention computation. In this way, AGAT can adaptively adjust the graph structure in highly variable time series data instead of relying on a pre-defined graph, which enables it to handle dependency relationships in different scenarios more flexibly, especially when dealing with data with complex time-varying topologies, and to cooperate with the subsequent frequency-domain analysis and feature fusion modules (such as REFusion) to jointly improve the sensitivity and robustness of anomaly detection. the model framework is shown in Fig. 2.

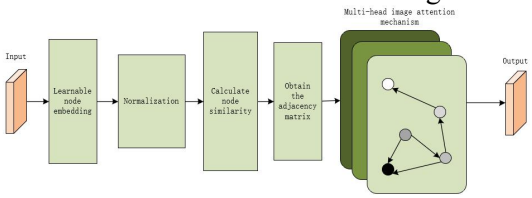

Figure 2. Framework of the Adaptive Graph Attention Module

2.4 Reciprocal Entropy-Weighted Fusion Module

In anomaly detection tasks for multivariate time series data, traditional time–frequency fusion methods usually rely on fixed weights or prior-based fusion schemes, which have certain limitations in dealing with the complexity and dynamics of the data. To address this problem, this study proposes an innovative Reciprocal Entropy-weighted Fusion module (REFusion), which adaptively adjusts the contributions of time-domain and frequency-domain features, breaking through the limitations of traditional fixed weighting and truly realizing dynamic time–frequency fusion. In traditional methods, the fusion of time-domain and frequency-domain features usually depends on fixed weights, ignoring the differences in the importance of each feature under different data patterns. In contrast, the REFusion module dynamically evaluates the uncertainty of each branch based on information entropy, thereby automatically adjusting the contributions of time-domain and frequency-domain features, and adopts a reciprocal weighting strategy—large entropy with small weight and small entropy with large weight—to achieve a parameter-free, differentiable, and stable adaptive fusion.

Let the output of one branch be $z \in \mathbb{R}^{B\times C\times T}$. For each sample, it is mapped to a differentiable probability distribution, and then a temperature-scaled softmax is applied to it:

$$P = \mathrm{softmax}\left(\frac{\mathrm{vec}(|Y|)}{T}\right), T > 0 \tag{5}$$

Accordingly, the normalized sample-level Shannon entropy is computed and further normalized to [0, 1] by log N:

$$H(z) = -\sum_{i=1}^{N} p_i \log p_i, \tag{6}$$

$$\hat{H}(z) = \frac{H(z)}{\log N} \in [0,1] \tag{7}$$

The larger $\hat{H}$ is, the more dispersed the energy and the more chaotic the structure; the smaller $\hat{H}$ is, the more concentrated the energy and the more definite the pattern. For the two branches, we obtain $\hat{H}_t = \hat{H}(t_{\text{out}})$ and $\hat{H}_f = \hat{H}(f_{\text{out}})$, respectively, and define the reciprocal weights.

$$\alpha_t = \frac{\hat{H}_f}{\hat{H}_t + \hat{H}_f + \varepsilon} \tag{8}$$

$$\alpha_f = \frac{\hat{H}_t}{\hat{H}_t + \hat{H}_f + \varepsilon}, \alpha_t + \alpha_f \approx 1 \tag{9}$$

and then reweight the outputs of the time-domain and frequency-domain branches:

$$\tilde{t} = \alpha_t t_{\text{out}}, \tilde{f} = \alpha_f f_{\text{out}} \tag{10}$$

After concatenating the inputs, the fused term $\text{fused} = \text{Concat}(x, \tilde{t}, \tilde{f})$ is obtained and then fed into the GRU for temporal fusion.

The FTAD model proposed in this paper is different from existing traditional methods; by integrating adaptive graph attention, the frequency-domain attention mechanism, and the reciprocal entropy-weighted fusion module, it realizes an organic combination of time-domain and frequency-domain features. In particular, the REFusion module adaptively adjusts the weights by computing the uncertainty of the features, thereby enhancing the effectiveness of time–frequency feature fusion, fully reflecting the collaborative optimization among the modules, avoiding the shortcomings of static weighting in traditional methods, and improving the performance of anomaly detection.

2.5 Anomaly Threshold Detection Module

In the anomaly detection stage, the model first obtains the predicted values of the future sequence and the reconstructed values of the input sequence through the prediction branch

and the reconstruction branch, respectively. Using these two outputs, the prediction error and reconstruction error at each time point are computed, and the two are combined in a weighted manner with a weighting parameter γ to obtain the final anomaly score S_t:

$$S_t = \sqrt{(y_t - \hat{y}_t)^2} + \gamma \cdot \sqrt{(x_t - \hat{x}_t)^2} \quad (11)$$

Among them, y_t and x_t denote the true future values and the input values, respectively, and $\hat{y}_t$ and $\hat{x}_t$ are the predicted values and reconstructed values of the model. This design enables the anomaly score to simultaneously reflect the deviation of the model's prediction of future trends and its reconstruction ability for normal patterns, thereby enhancing the discriminability of anomaly identification. Subsequently, the module automatically determines the detection boundary θ using the percentile threshold of the anomaly scores on the training set. In the testing stage, when the anomaly score at a certain time exceeds $S_t \geq \theta$, that time point is marked as an anomalous point.

3. Experimental Results and Analysis

3.1 Datasets

In terms of datasets, three multivariate real-world datasets were adopted to evaluate the performance of the model, and their detailed information is summarized in Table 1. Among them, the SMD (Server Machine Dataset) [8] is a five-week dataset collected from the computing clusters of a large Internet company, capturing 38 resource utilization indicators from 28 server machines, with an anomaly rate of about 4.16%. the MSL (Mars Science Laboratory dataset) [9] was collected by NASA and presents the status of sensor and actuator data of the Mars rover; the SMAP (Soil Moisture Active Passive Dataset) [9] also comes from NASA and reflects the soil samples and telemetry data used by the Mars rover, with 25 dimensions and obvious point anomalies compared with the other datasets.

Table 1. Description of Datasets

Dataset	Feature dimension	Training set	Test set	Anomaly ratio/%
SMAP	25	135183	427617	13.13
MSL	55	58317	73729	10.72
SMD	38	708405	708420	4.16

3.2 Performance Comparison

Table 2. Experimental Results of the Proposed Model and Baseline Models (unit:%)

Datasets	SMAP			MSL			SMD		
Model/Metric	Pre	R	F1	Pre	R	F1	Pre	R	F1
LSTM-VAE[10]	85.51	63.66	72.98	52.57	94.46	67.80	88.73	51.11	64.86
OmniAnomaly[11]	74.16	97.76	84.34	88.67	91.17	89.89	83.34	94.49	88.57
USAD[12]	90.96	85.29	88.03	93.08	89.17	91.08	99.89	80.26	89.00
GDN[13]	89.32	88.72	89.02	91.35	86.12	88.66	71.70	99.74	83.42
TranAD[14]	93.12	71.33	80.78	90.72	94.73	92.68	74.30	81.65	77.80
MTAD-GAT[15]	89.06	91.23	90.13	87.54	94.40	90.84	94.09	85.24	89.45
TFMAE[16]	94.36	98.03	96.16	92.60	92.69	92.64	91.54	87.46	89.46
Ours	96.82	97.26	97.04	93.98	96.73	95.34	89.45	99.98	94.43

By examining Table 2, we find that prediction- and reconstruction-based methods (e. g., LSTM-VAE, OmniAnomaly, USAD) can learn latent representations but tend to overfit anomalous patterns, weakening detection. Graph-based methods (e. g., GDN, MTAD-GAT) emphasize cross-variable correlations yet rely on static, often unweighted graphs, limiting their ability to model time-varying dependencies and variable importance. Contrastive learning methods (e. g., TranAD, TFMAE) depend on generic augmentations and instance discrimination, which may not align well with anomaly detection semantics. Our design alleviates overfitting in reconstruction-based methods, compensates for graph methods' limitations in modeling dynamic weighted structures, and explicitly considers uneven variable importance, thereby reducing false positives and false negatives while maintaining robustness and improving overall detection performance and stability.

3.3 Ablation Study

To verify the effectiveness of each module, we conducted ablation experiments on the SMAP, MSL, and SMD datasets by removing different components and comparing the F1-scores of the

resulting variants. Specifically, we removed the EMA decomposition module (w/o EMADecomposition), the adaptive graph attention module (w/o AdaptiveGAT), the frequency-domain attention module (w/o FourierCrossAttention), and the reciprocal entropy-weighted fusion module (w/o REFusion). the corresponding results are reported in Table 3.

Table 3. Comparison of F1-Scores of Different Model Variants (unit:%)

Models	SMAP	MSL	SMD
w/o EMADecomposition	93.43	91.06	92.49
w/o AdaptiveGAT	93.16	89.94	92.16
w/o FourierCrossAttention	92.72	90.02	91.87
w/o REFusion	93.97	92.26	92.50
FTAD	97.04	95.34	94.43

From Table 3, FTAD consistently outperforms its ablated variants on all three datasets, confirming the contribution of each module. Removing EMADecomposition, AdaptiveGAT, FourierCrossAttention, or REFusion leads to F1 drops of about 2–5 percentage points, showing that EMA enhances temporal feature separability, the adaptive graph structure is crucial for modeling complex sensor relations, frequency-domain modeling is indispensable for detection, and REFusion effectively allocates weights based on branch uncertainty to improve time–frequency complementarity and anomaly separability. Additional ablations using only frequency-domain attention or its learnable-weight variant further verify the effectiveness of the combined fusion design.

4. Conclusion

In this paper, aiming at the problems of limited capture of time-domain features and insufficient utilization of frequency-domain information in multivariate time series anomaly detection, a dual-branch model FTAD that fuses the time and frequency domains is proposed. the model segments the sequence using a sliding window, introduces exponential moving average (EMA) to achieve time series decomposition, and models the trend and periodic components through an adaptive graph attention network and a frequency-domain cross-attention mechanism, respectively, thereby significantly enhancing the representation ability for complex dynamic patterns and the sensitivity to anomalies.

Experimental results on public datasets (SMAP, MSL, SMD) show that FTAD is clearly superior to existing methods in terms of Precision, Recall, and F1-score, verifying its potential for industrial applications. In addition, ablation experiments are further conducted to analyze the contribution of each component, confirming the key roles of EMA decomposition, adaptive graph modeling, and the improved frequency-domain attention mechanism in improving model performance.

In the future, multi-scale frequency-domain feature modeling and multi-task learning mechanisms can be further explored to improve the generalization ability and robustness of the model to anomalous patterns. At the same time, more diverse graph neural network architectures can be investigated to fully capture the potential spatial and temporal dependencies in multivariate time series.

References

[1] Wen Q, Yang L, Zhou T, et al. Robust time series analysis and applications: An industrial perspective [C] // Proc of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. New York: ACM, 2022:4836-4837.

[2] Kheddar H, Himeur Y, Awad A I. Deep transfer learning for intrusion detection in industrial control networks: A comprehensive review [J]. Journal of Network and Computer Applications, 2023, 220:103760.

[3] Wang F, Jiang Y, Zhang R, et al. A Survey of Deep Anomaly Detection in Multivariate Time Series: Taxonomy, Applications, and Directions [J]. Sensors (Basel, Switzerland), 2025, 25(1):190.

[4] Si H, Li J, Pei C, et al. Timeseriesbench: An industrial-grade benchmark for time series anomaly detection models [C]//2024 IEEE 35th International Symposium on Software Reliability Engineering (ISSRE). IEEE, 2024:61-72.

[5] Wang Z, Pei C, Ma M, et al. Revisiting VAE for unsupervised time series anomaly detection: A frequency perspective [C] // Proceedings.

[6] Zhou T, Ma Z, Wen Q, et al. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting [C] // International Conference on Machine Learning. PMLR, 2022:27268-27286.

[7] Lu Y X, Jin X B, Chen J, et al. F-SE-LSTM: A Time Series Anomaly Detection Method with Frequency Domain Information [J/OL]. arXiv:2412.02474, 2024[2025-08-11]. https://arxiv. org/abs/2412.02474

[8] Su Y, Zhao Y, Niu C, et al. Robust anomaly detection for multivariate time series through stochastic recurrent neural network [C] // Proc of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. New York: ACM, 2019:2828-2837.

[9] Hundman K, Constantinou V, Laporte C, et al. Detecting spacecraft anomalies using LSTMs and nonparametric dynamic thresholding [C] // Proc of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. New York: ACM, 2018:387-395.

[10]. Park D, Hoshi Y, Kemp C C. A multimodal anomaly detector for robot-assisted feeding using an lstm-based variational autoencoder [J]. IEEE Robotics and Automation Letters, 2018, 3(3):1544-1551.

[11]. SU Y, ZHAO Y J, NIU CH H, et al. Robust anomalydetectionformultivariatetimeseriest hrough stochasticrecurrentneuralnetwork [C]. 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2019:2828-2837.

[12]. Audibert J, Michiardi P, Guyard F, et al. USAD: Unsupervised anomaly detection on multivariate time series [C] // Proc of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. New York: ACM, 2020:3395-3404.

[13]. Tuli S, Casale G, Jennings N R. TranAD: Deep transformer networks for anomaly detection in multivariate time series data [J/OL]. arXiv:2201.07284, 2022[2025-08-11]. https://arxiv. org/abs/2201.07284.

[14]. Deng A, Hooi B. Graph neural network-based anomaly detection in multivariate time series [C] // Proceedings of the AAAI Conference on Artificial Intelligence. Palo Alto, CA: AAAI Press, 2021, 35(5):4027-4035.

[15]. Fang Y, Xie J, Zhao Y, et al. Temporal-frequency masked autoencoders for time series anomaly detection [C]//2024 IEEE 40th International Conference on Data Engineering (ICDE). IEEE, 2024:1228-1241.

[16]. ZHAO H, WANG Y J, DUAN J Y, et al. Multivariatetime-seriesanomaly detection via graph attention network [C]. 2020 IEEE International ConferenceonData Mining(ICDM), 2020:841-850.