

Towards Non-Stationary Time Series Forecasting with Temporal Stabilization and Frequency Differencing

Junkai Lu^{1,2}, Peng Chen¹, Chenjuan Guo¹, Yang Shu¹, Meng Wang², Bin Yang^{1*}

¹School of Data Science and Engineering, East China Normal University, Shanghai, China

²School of Computer Science, Xi'an Polytechnic University, Shaanxi, China

{jklu, pchen}@stu.ecnu.edu.cn, {cjguo, yshu, byang}@dase.ecnu.edu.cn, wangmeng@xpu.edu.cn

Abstract

Time series forecasting is critical for decision-making across dynamic domains such as energy, finance, transportation, and cloud computing. However, real-world time series often exhibit non-stationarity, including temporal distribution shifts and spectral variability, which pose significant challenges for long-term time series forecasting. In this paper, we propose *DTAF*, a dual-branch framework that addresses non-stationarity in both the temporal and frequency domains. For the temporal domain, the Temporal Stabilizing Fusion (TFS) module employs a non-stationary mix of experts (MOE) filter to disentangle and suppress temporal non-stationary patterns while preserving long-term dependencies. For the frequency domain, the Frequency Wave Modeling (FWM) module applies frequency differencing to dynamically highlight components with significant spectral shifts. By fusing the complementary outputs of TFS and FWM, DTAF generates robust forecasts that adapt to both temporal and frequency domain non-stationarity. Extensive experiments on real-world benchmarks demonstrate that DTAF outperforms state-of-the-art baselines, yielding significant improvements in forecasting accuracy under non-stationary conditions.

Introduction

With the acceleration of digital transformation, time series have been generated continuously across various domains, including energy (Qiu et al. 2025g; Wu et al. 2025c; Li et al. 2025), finance (Ye et al. 2025; Liu et al. 2025b; Wu et al. 2025b), transportation (Gao et al. 2025; Qiu et al. 2025f; Wu et al. 2025d), and cloud computing (Hu et al. 2024; Qiu et al. 2025d; Wu et al. 2024). Accurate forecasting of future values from historical time series is essential for decision-making and operational planning (Wang et al. 2025b).

In real-world scenarios, time series often exhibit non-stationary patterns in both the temporal and frequency domains (Qiu et al. 2025a; Wang et al. 2025c). As illustrated in Figure 1, the entire sequence is divided into four patches, each characterized by distinct patterns. The distributions of these patches evolve over time in temporal and frequency domains, as illustrated in Figure 1(f-i) and Figure 1(j-m), reflecting the underlying non-stationarity (Qiu

et al. 2025e). These observations highlight the necessity and difficulty of designing methods that can effectively handle non-stationarity in both domains. Specifically, addressing this problem involves overcoming two key challenges:

Extracting and disentangling heterogeneous non-stationary patterns to adapt to long-term dependency modeling is challenging. Time series data inherently consists of stationary and non-stationary components in the temporal domain (Qiu et al. 2025c; Wu et al. 2025a; Chen et al. 2025). Directly modeling long-term temporal dependencies on raw sequences could be severely impaired by non-stationary dynamics, leading to degraded prediction performance. A promising solution involves explicitly modeling the stationary component and isolating non-stationary effects. However, time series non-stationary patterns are highly complex and heterogeneous, making it impossible for a single architecture model to comprehensively model these non-stationary patterns. Therefore, how to design an architecture capable of extracting and disentangling these heterogeneous non-stationary patterns to enable long-term temporal dependency modeling is a challenge.

Dynamic modeling frequency drifts and spectral variations present a challenge. Time series contains rich frequency characteristics that reflect periodicity and trends. These frequency characteristics are inherently non-stationary, changing over time (as shown in Figure 1(j-m)). However, existing frequency analysis methods such as the Fourier Transform (chan Park, gi Jang, and Kang 2021) assume stationarity within the observation window and produce global frequency representations. Therefore, they fail to capture transient frequency shifts and specific local patterns. This limitation is particularly evident in applications like electricity load forecasting, where frequency components corresponding to daily and seasonal cycles change over time due to user behavior and environmental factors (e.g., climate changes). In such cases, relying on static frequency analysis may obscure important variations necessary for accurate prediction. Therefore, efficient frequency modeling that can dynamically track frequency drifts is essential.

To address these challenges, we propose a novel Dual-branch Modeling with Temporal Fusion and Frequency Differencing (*DTAF*), to address the temporal and spectral non-stationarity inherent in time series. As illustrated in Figure 2, DTAF integrates two core components: *Temporal Stabilizing*

*Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

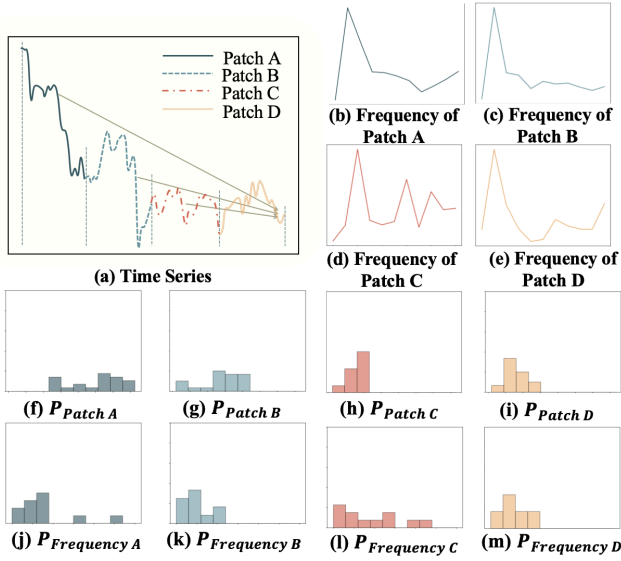


Figure 1: A non-stationary time series split into four patches (A, B, C, D), (b-e) shows the amplitude spectrum of the four patches, (f-i) shows diverse distributions in the temporal domain, and (j-m) shows diverse distributions in the frequency domain, reflecting non-stationarity across both domains.

Fusion (TFS) and Frequency Wave Modeling (FWM).

In the temporal domain, we introduce a Non-stationary MOE Filter in TFS, composed of expert networks specialized in distinct non-stationary patterns. These extracted components are subtracted from the input to yield a stationary residual, enabling more robust modeling. Based on it, we conduct temporal fusion to capture long-term dependencies.

In the frequency domain, FWM addresses frequency non-stationarity via spectral differencing, highlighting components with the most significant changes. This allows FWM to adapt to dynamic spectral patterns. Finally, Dual-branch Attention fuses the outputs of TFS and FWM, enabling robust forecasting that accounts for the two domain shifts. Our contributions are summarized as follows:

- To address the challenge of non-stationarity in time series forecasting, we propose a novel model, DTAF, which leverages dual-branch modeling to capture both temporal and frequency-domain non-stationarity.
- We propose the Temporal Stabilizing Fusion (TFS) module, integrating a non-stationary MOE Filter, which automatically learns diverse temporal non-stationary patterns, then performs global fusion to capture long-term dependency, thus significantly enhancing prediction.
- We introduce the Frequency Wave Modeling (FWM) module, which applies differencing techniques in the frequency domain to enhance the model’s ability to model and adapt to frequency domain shifts.
- Extensive experiments on eleven real-world datasets have demonstrated that DTAF achieves state-of-the-art prediction accuracy.

Related Work

Multivariate Time Series Forecasting

Multivariate time series forecasting aims to predict future values of multiple channels based on their historical data (Qiu et al. 2025b; Cheng et al. 2023; Wang et al. 2023). Traditional methods such as RNN (Cho et al. 2014; Gasthaus et al. 2019) and LSTM (Pillai 2025) capture time dependencies, but as the sequence length increases, the historical patterns will be lost. To address these challenges, deep learning approaches such as DeepAR (Rangapuram et al. 2018) have emerged, utilizing autoregressive probabilistic models to enhance robustness and accuracy in forecasting. Recent advancements have introduced Graph Neural Networks (GNNs) (Cirstea et al. 2021; Chen et al. 2023; Liu et al. 2022a) to model spatial dependencies in time series. For example, in tasks such as energy forecasting (Campagne et al. 2024; Tian et al. 2024), finance prediction (Wang et al. 2025a), and traffic prediction (Prabowo et al. 2023), GNNs are capable of capturing both temporal information and spatial correlations between variables, thereby improving forecasting accuracy. Simultaneously, Transformer-based architectures use attention mechanisms to model long-range dependencies, such as PatchTST (Nie et al. 2023), Pathformer (Chen et al. 2024).

Non-Stationary Modeling for Time Series

Although stationarization is essential for enhancing the predictability of time series data, real-world series invariably exhibit non-stationary patterns. Classical statistical models such as ARIMA (Box, Jenkins, and Reinsel 2015) address this by applying temporal differencing to enforce stationarity. In recent years, some methods addressing these challenges on the **normalization aspect**: Adaptive Norm (Ogasawara et al. 2010), which applies z-score normalization using global training statistics; DAIN (Passalis et al. 2020); RevIN (Kim et al. 2022), which performs two-stage instance normalization on inputs and outputs; and SAN (Liu et al. 2023), which dynamically normalizes within shorter windows to estimate local means and variances. And others from the **model aspect** to capture non-stationary patterns, including Non-stationary Transformer (Liu et al. 2022b), DERITS (Fan et al. 2024), AEFIN (Xiong and Wen 2025), NsDiff (Ye, Xu, and Gui 2025), and TimeBridge (Liu et al. 2025a). AdaRNN (Du et al. 2021), which proposes an adaptive RNN to alleviate the impact of non-stationary factors by characterizing and matching distributions; Triformer (Cirstea et al. 2022), which proposes a light-weight approach to enable variable-specific model parameters, making it possible to capture distinct temporal patterns from different variables.

Although the existing technologies have alleviated non-stationarity to a certain extent, the importance of jointly modeling it from both the temporal and frequency domains is generally ignored. To address this, we propose a new method, DTAF, that models non-stationarity in the temporal and frequency domains, enhancing prediction performance under non-stationary conditions.

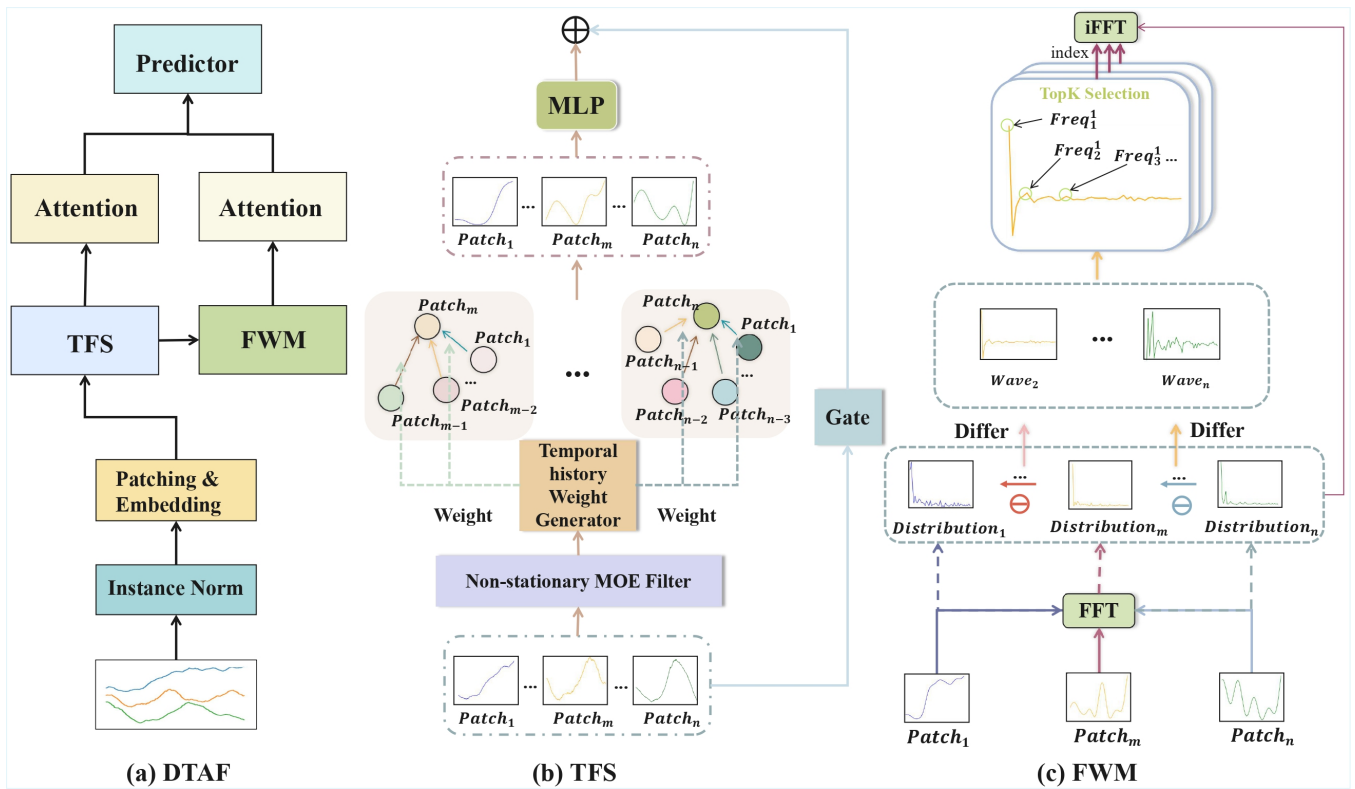


Figure 2: The architecture of DTAF. The two main components of DTAF are TFS and FWM. They are used to aggregate the information in the temporal domain and model in the frequency domain with differencing. Temporal aggregation captures long-term dependencies and local variations, while frequency differencing provides insights into the changing spectral components.

Methodology

Overall Framework

To address the non-stationarity in the temporal and frequency domains, we propose a novel model, Dual-branch Modeling with Temporal Fusion and Frequency Differencing (DTAF). As illustrated in Figure 2, DTAF comprises several core components: Instance Norm (Kim et al. 2022), Patching&Embedding (Nie et al. 2023), Temporal Stabilizing Fusion (TFS), Frequency Wave Modeling (FWM), and a Predictor (Zeng et al. 2023). Instance Norm alleviates distributional shifts between training and inference, enhancing the robustness of the model against variations in data distributions and improving generalization performance. Inspired by the patching strategy in Transformer architectures (Nie et al. 2023), Patching&Embedding segments long sequences into patches and embeds them into high-dimensional representations, preserving local temporal patterns. The TFS and FWM modules are specifically designed to mitigate non-stationarity in the temporal and frequency domains. Finally, the Predictor integrates the features from both branches to produce accurate prediction results.

The core of design is TFS and FWM, which are designed to model non-stationarity in the temporal and frequency domains. TFS utilizes a Non-stationary MoE Filter to dynamically extract and remove non-stationary patterns and then uses the Temporal Fusion module to model long-

term dependency based on the stationary residual. Complementing the temporal modeling, FWM addresses frequency non-stationarity by applying differencing in the frequency domain, highlighting components with the most significant spectral changes. These enable DTAF to adapt to temporal and frequency dynamics simultaneously, leading to more accurate forecasting under non-stationary conditions.

Temporal Stabilizing Fusion

In complex time series, the heterogeneous non-stationary patterns lead to inconsistent distributions among patches, reducing the prediction performance.

As illustrated in Figure 1 (f-i), although these patches may follow similar global trends, their local fluctuations vary drastically. These heterogeneous non-stationary patterns make it difficult to effectively capture long-term temporal dependencies.

To address this, we introduce the Temporal Stabilizing Fusion module (TFS). It first applies a non-stationary MOE Filter to extract and subtract these non-stationary components, and then performs a temporal fusion module to capture long-term dependencies. By extracting and filtering heterogeneous non-stationary patterns, TFS enables robust long-term modeling and significantly enhances forecasting accuracy under dynamical non-stationary conditions.

Non-stationary MOE Filter. To alleviate the impact of

non-stationarity in the temporal modeling, we propose the Non-stationary MOE Filter, which is designed to learn non-stationary patterns and remove them from the original sequence, stabilizing the input and making it more suitable for long-term modeling. Specifically, the Non-stationary MOE Filter consists of a collection of Experts dedicated to extracting non-stationary patterns, along with a routing network that dynamically assigns inputs to the experts.

For notational clarity, we describe our method using a univariate time series $\mathbf{X} \in \mathbb{R}^T$, where T denotes the input sequence length (extension to multivariate cases is straightforward). We first apply instance normalization to standardize the input statistics, mitigating distribution shift. Recognizing that individual time steps often lack meaningful context, we employ a patching strategy that segments the sequence into N overlapping patches $\{\mathbf{X}^1, \mathbf{X}^2, \dots, \mathbf{X}^N\}$, where each patch captures local temporal patterns. These patches are then projected into a high-dimensional latent space through patch embedding, yielding the embeddings $\{\mathbf{X}_{\text{patch}}^1, \mathbf{X}_{\text{patch}}^2, \dots, \mathbf{X}_{\text{patch}}^N\}$ for subsequent processing, where $\mathbf{X}_{\text{patch}}^i \in \mathbb{R}^d$ represents i -th patch embedding, and d denotes the embedding dimension.

For i -th input patch embedding $\mathbf{X}_{\text{patch}}^i$, the routing network first assigns a weight to each expert's contribution based on each input characteristic, and then each expert utilizes special MLP networks to extract potential non-stationary patterns. After aggregating the weighted outputs from all experts, the non-stationary pattern $\mathbf{X}_{\text{patterns}}^i \in \mathbb{R}^d$ for the sequence is obtained, which is subtracted from the input patch $\mathbf{X}_{\text{patch}}^i$, as illustrated by the following equations:

$$\begin{aligned} \mathcal{E} &= \{e_1, e_2, e_3, \dots, e_m\}, \\ \mathbf{router}^i &= \text{Softmax}(\text{Linear}(\mathbf{X}_{\text{patch}}^i)), \\ \mathbf{X}_{\text{patterns}}^i &= \sum_{j=1}^m e_j(\mathbf{X}^i) * \mathbf{router}_j^i, \end{aligned} \quad (1)$$

where \mathcal{E} denotes the expert set and $\text{Linear}(\cdot)$ denotes a special linear layer (Liu et al. 2025c), and each expert denotes multiple linear layers. To ensure that the experts learn the non-stationary patterns, we introduce a KL divergence (Ji et al. 2022) loss to enforce this constraint. Specifically, the loss function is formulated as:

$$\mathbf{X}_{\text{stable}}^i = \mathbf{X}_{\text{patch}}^i - \mathbf{X}_{\text{patterns}}^i, \quad (2)$$

$$\mathcal{L}_{\text{stable}} = \alpha * \sum_{i=1}^N \sum_{j=1}^N \text{KL}(\mathbf{X}_{\text{stable}}^i, \mathbf{X}_{\text{stable}}^j) / N^2, \quad (3)$$

where $\text{KL}(\cdot)$ calculates the KL divergence and α is a hyperparameter. KL divergence can make the distribution of each output patch stationary. Since the outputs are the features extracted by subtracting exports from the original patches, exports can learn non-stationary features.

Temporal Fusion. After extracting and removing non-stationary components from each input patch embedding, we obtain stationary patterns to model long-term temporal dependencies. To further enhance robustness against

potential residual non-stationary effects, we propose a temporal fusion module that dynamically balances historical context and current information.

For each patch $\mathbf{X}_{\text{stable}}^i \in \mathbb{R}^d$, a linear feature extractor ($\text{LinearExtra}(\cdot)$) that incorporates time series decomposition is first employed to capture the patterns of the current patch $\mathbf{X}_{\text{stable}}^i$, formulated as follows:

$$\mathbf{X}_t^i = \text{Avgpool}(\text{padding}(\mathbf{X}_{\text{stable}}^i)), \mathbf{X}_s^i = \mathbf{X}_{\text{stable}}^i - \mathbf{X}_t^i, \quad (4)$$

$$\mathbf{X}_h^i = \mathbf{W}_t^i \cdot \mathbf{X}_t^i + \mathbf{W}_s^i \cdot \mathbf{X}_s^i, \quad (5)$$

where $\mathbf{X}_t^i, \mathbf{X}_s^i \in \mathbb{R}^d$ represent the decomposed trend and seasonal components respectively (Wu et al. 2021), $\mathbf{W}_t^i, \mathbf{W}_s^i \in \mathbb{R}^{d \times d}$ are learnable parameters and $\mathbf{X}_h^i \in \mathbb{R}^d$ is the temporal representation of the i -th patch. Next, a linear layer gets the weights between the current patch and each historical patch, reflecting the degree of information fusion required from each historical patch:

$$\mathbf{Weight}^i = \text{Softmax}(\mathbf{W}_{\text{history}} \cdot \mathbf{X}_h^i + \mathbf{b}_{\text{history}}), \quad (6)$$

where $\mathbf{W}_{\text{history}} \in \mathbb{R}^{N \times d}$, and $\mathbf{b}_{\text{history}} \in \mathbb{R}^N$. To ensure no future information leaks into the model, all weights \mathbf{Weight}_n^i , for $n > i$, are set to zero, enforcing the use of historical patches exclusively. Based on the weights generated by the Temporal history Weight Generator, we then integrate the historical patch information through a weighted aggregation, formulated as:

$$\mathbf{X}_{\text{history}}^i = \text{MLP}\left(\sum_{n=1}^{i-1} \mathbf{Weight}_n^i \cdot \mathbf{X}_{\text{stable}}^n\right), \quad (7)$$

Where $\text{MLP}(\cdot)$ is implemented as one linear transformation layer, facilitating the extraction and reinforcement of key history features. To incorporate the current patch's feature representations, we introduce a gating mechanism designed to dynamically amplify or suppress the contribution of the current patch to the final representations:

$$\mathbf{X}_{\text{current}}^i = \text{Gate}(\mathbf{X}_{\text{patch}}^i) \cdot \mathbf{X}_{\text{patch}}^i, \quad (8)$$

$$\text{Gate}(\mathbf{X}_{\text{patch}}^i) = \mathbf{W}_{\text{gate}} \cdot \text{LinearExtra}(\mathbf{X}_{\text{patch}}^i) + \mathbf{b}_{\text{gate}}, \quad (9)$$

where $\mathbf{W}_{\text{gate}} \in \mathbb{R}^{1 \times d}$, $\mathbf{b}_{\text{gate}} \in \mathbb{R}^1$ are learnable parameters. Finally, the fused representation is obtained by combining the current patch's features with the historical information:

$$\mathbf{H}_t^i = \mathbf{X}_{\text{current}}^i + \mathbf{X}_{\text{history}}^i, \quad (10)$$

where $\mathbf{H}_t^i \in \mathbb{R}^d$ represents global representation to the i -th patch. By integrating historical and current information, DTAF captures essential global patterns while further suppressing residual non-stationary factors, enabling more effective modeling of long-term temporal dependencies.

Overall, DTAF effectively enhances long-term temporal modeling by jointly leveraging the Non-stationary MOE Filter and the Temporal Fusion module. The MOE Filter extracts heterogeneous non-stationary patterns from each patch, producing relatively stationary representations.

Based on the above components, the Temporal Fusion module adaptively integrates historical and current information to further adapt to the non-stationarity. This two-stage process stabilizes temporal dynamics and enhances long-term dependency modeling under non-stationary conditions.

Frequency Wave Modeling

Time series data often reveal trends and seasonal patterns through their frequency characteristics. However, these frequency components are not static, leading to non-stationarity in the frequency domain as shown in Figure 1 (j-m). Existing approaches typically assume that the frequency patterns of time series remain stable over time. However, such assumptions fail to account for the non-stationarity in the frequency domain, where key components may shift over time. This limits the model’s ability to capture evolving seasonal or trend patterns in time series forecasting.

To address this limitation, we propose a novel module named Frequency Wave Modeling, which is specifically designed to model non-stationarity in the frequency domain. By dynamically tracking the temporal changes of frequency components and selectively emphasizing those that exhibit significant variations, the module enables the model to concentrate on the most informative spectral patterns.

Based on the patch representation \mathbf{H}_t^i from the TFS module, we introduce Frequency Wave Modeling to further characterize non-stationary patterns in the frequency domain as shown in Figure 2 (c). First, similar to other approaches (Wu et al. 2023), we apply the Fast Fourier Transform (FFT) to each patch, as expressed as:

$$\mathbf{Freq}^i = \text{rFFT}(\mathbf{H}_t^i), \quad (11)$$

where $\text{rFFT}(\cdot)$ denotes the real-valued Fast Fourier Transform that converts time-domain information into the frequency domain, yielding a set of spectral components. These spectral components capture the distribution of frequency components at different patches. Subsequently, we perform a differencing operation on the spectrum sequence and select the frequency components exhibiting significant variation, which can emphasize non-stationary components:

$$\mathbf{Wave}^i = \mathbf{Freq}^i - \mathbf{Freq}^{i-1} (i \geq 2), \quad (12)$$

$$\mathbf{Picks}^i = \text{Topk}(\mathbf{Wave}^i), \quad (13)$$

$$\mathbf{Freq}_{f_j \notin \mathbf{Picks}^i}^i = 0, \quad (14)$$

where $\mathbf{Wave} \in \mathbb{R}^{N \times d/2+1}$ represents the degree of change of each frequency in each patch, and for the first patch, we set it to \mathbf{Freq}^1 ; The $\text{Topk}(\cdot)$ function selects the k most significantly changing frequency components for every patch, thus obtaining $\mathbf{Picks} \in \mathbb{R}^{N \times k}$, which includes the k most changing frequency components for each patch. For any frequency f_j , which is not selected, the corresponding $\mathbf{Freq}_{f_j}^i$ is set to zero. Finally, we apply the inverse Fast Fourier Transform in $\mathbf{Freq}^i \in \mathbb{R}^{d/2+1}$ to convert the selected frequency features back to the time domain, obtaining the representation $\mathbf{H}_f \in \mathbb{R}^{N \times d}$.

By incorporating differencing operators to track evolving spectral patterns over time, FWM constructs discriminative

spectral representations that capture complex, time-varying frequency dynamics, enhancing forecasting performance.

Dual-branch Attention

Based on temporal and frequency features, we then fuse the two domain features. This fusion allows the model to leverage both temporal and frequency information, leading to more robust and accurate predictions. Traditional fusion approaches, such as direct concatenation or using cross-attention to reconstruct one domain from the other, often fall short in capturing the complementary nuances between temporal and frequency features. We propose a dual-branch attention feature fusion mechanism, referred to as Dual-branch Attention, which can be formulated as follows:

$$\mathbf{Atten}_t = \text{Softmax}\left(\frac{\mathbf{Q}_t \mathbf{K}_t^T}{\sqrt{d}}\right) \mathbf{V}_t, \quad (15)$$

$$\mathbf{Atten}_f = \text{Softmax}\left(\frac{\mathbf{Q}_f \mathbf{K}_f^T}{\sqrt{d}}\right) \mathbf{V}_f, \quad (16)$$

$$\mathbf{H}_{\text{fusion}} = \text{Concat}(\mathbf{Atten}_t, \mathbf{Atten}_f). \quad (17)$$

Here $\mathbf{Q}_t, \mathbf{K}_t, \mathbf{V}_t \in \mathbb{R}^{N \times d}$ are derived from the temporal-domain features $\mathbf{H}_t \in \mathbb{R}^{N \times d}$ through linear transformations, and $\mathbf{Q}_f, \mathbf{K}_f, \mathbf{V}_f \in \mathbb{R}^{N \times d}$ are derived from the frequency-domain features \mathbf{H}_f through similar transformations. After obtaining the attention of the temporal domain ($\mathbf{Atten}_t \in \mathbb{R}^{N \times d}$) and the attention of the frequency domain ($\mathbf{Atten}_f \in \mathbb{R}^{N \times d}$), we concatenate them to form $\mathbf{H}_{\text{fusion}} \in \mathbb{R}^{2N \times d}$. The Dual-branch Attention simultaneously extracts features from both the time and frequency domains, providing richer features for subsequent prediction. Finally, the fused features pass through a prediction module (Predictor), which consists of a fully connected layer:

$$\mathbf{X}_{\text{out}} = \text{Linear}(\text{Flatten}(\mathbf{H}_{\text{fusion}})), \quad (18)$$

where $\mathbf{X}_{\text{out}} \in \mathbb{R}^F$, F is the prediction length of time series. This module adaptively selects the required features from the two domains to produce the final prediction \mathbf{X}_{out} .

Experiment

Dataset

We select 11 widely used time-series datasets for our model, covering a range of domains such as electricity (ETTh1, ETTh2, ETTm1, ETTm2, Electricity), health (ILI, Covid-19), traffic (Traffic), web (Wike2000), environment (Weather), and banking (NN5).

Baselines

We select several state-of-the-art models as baselines and evaluate them in TFB (Qiu et al. 2024), including Amplifier (Fei et al. 2025); Non-stationary Transformer (Stationary) (Liu et al. 2022b); iTransformer (Liu et al. 2024); Pathformer (Chen et al. 2024); PatchTST (Nie et al. 2023); Crossformer (Zhang and Yan 2023); FITS (Xu, Zeng, and Xu 2024); TimeMixer (Wang et al. 2024); TimesNet (Wu et al. 2023), and DLinear (Zeng et al. 2023).

Models	DTAF (ours)		Amplifier (2025)		iTransformer (2024)		Pathformer (2024)		FITS (2024)		TimeMixer (2024)		PatchTST (2023)		Stationary (2022)	
	mse	mae	mse	mae	mse	mae	mse	mae	mse	mae	mse	mae	mse	mae	mse	mae
ILI	1.688	0.801	1.819	0.888	1.857	0.892	1.995	0.909	2.334	1.052	<u>1.819</u>	<u>0.886</u>	1.902	0.879	2.389	1.027
Covid-19	1.351	0.040	5.578	0.112	<u>1.488</u>	<u>0.049</u>	3.326	0.086	5.669	0.123	12.941	0.133	1.697	0.056	2.658	0.078
NN5	0.643	0.538	1.794	1.018	0.660	<u>0.550</u>	0.698	0.582	0.811	0.653	<u>0.655</u>	0.593	0.688	0.589	1.295	0.915
ETTh	0.369	0.398	0.385	0.416	0.404	0.425	0.388	0.410	<u>0.371</u>	<u>0.406</u>	0.388	0.419	0.384	0.415	0.521	0.493
ETThm	0.297	0.338	0.327	0.364	0.314	0.358	<u>0.305</u>	<u>0.341</u>	<u>0.305</u>	0.345	0.306	0.349	0.302	0.347	0.456	0.430
Weather	0.222	0.250	<u>0.222</u>	0.263	0.232	0.269	0.225	<u>0.257</u>	0.243	0.280	0.226	0.264	0.223	0.261	0.293	0.315
Electricity	0.160	0.248	0.174	0.267	<u>0.163</u>	<u>0.258</u>	0.168	0.261	0.169	0.265	0.184	0.284	0.171	0.270	0.194	0.295
Traffic	0.402	0.249	0.423	0.294	0.397	0.281	0.416	<u>0.264</u>	0.429	0.302	0.409	0.279	<u>0.397</u>	0.275	0.621	0.339
1 st Count	7	8	0	0	<u>1</u>	<u>0</u>	0	0	0	0	0	0	0	0	0	0

Table 1: Multivariate forecasting results. Results are averaged from all forecasting horizons. The input lengths of ILI, Covid-19, and NN5 are selected from 36, 104, while the remaining datasets are selected from 96, 336, 512. The best results are highlighted in bold, and the second-best results are underlined. ETTh and ETThm represent the average of ETTh1 and ETTh2 and the average of ETThm1 and ETThm2. Due to space limitations, only a part is displayed.

Settings

For the forecasting horizon, the ILI, Covid-19, NN5, and Wike2000 datasets are evaluated using prediction lengths of 24, 36, 48, 60, while the remaining datasets employ prediction lengths of 96, 192, 336, 720. DTAF is optimized using the AdamW optimizer (Loshchilov and Hutter 2019). We select MAE and MSE as metrics. The loss is combined with three parts: L1 loss as task loss, r-drop loss (Liang et al. 2021) as robust loss to improve robustness of the model, and stable loss, which uses KL to constrain the Non-stationary MOE Filter to learn Non-stationary patterns. It can be formulated as: $\mathcal{L} = \mathcal{L}_{\text{task}} + \alpha \cdot \mathcal{L}_{\text{stable}} + \beta \cdot \mathcal{L}_{\text{robust}}$, where α, β is a hyperparameter. And the evaluation is run in the TFB (Qiu et al. 2024) benchmark. All results of the experiments are the average of running three times.

Main Results

Table 1 presents the performance of our model, DTAF, which consistently surpasses existing baselines across two evaluation metrics. 1) **Significant Improvements over Baselines**: Among the 16 evaluation metrics, DTAF ranks first in 15 cases. Compared with Stationary, which aims to handle non-stationarity through input transformation, DTAF demonstrates outstanding performance on non-stationary datasets such as Covid-19 (reduces MSE by **49.0%** and MAE by **50.0%**) and NN5 (reduces MSE by **50.4%** and MAE by **41.2%**), proving the effectiveness of DTAF in modeling non-stationary data. 2) **Robust Predictive Capabilities**: Compared to other models such as Crossformer, PatchTST, TimeMixer, and other state-of-the-art approaches, DTAF not only achieves notable accuracy im-

provements but also maintains stable performance across diverse time series datasets, demonstrating strong adaptability to varying data distributions and temporal patterns.

Ablation Studies

In this ablation study, we evaluate DTAF on one non-stationary time series dataset and one normal dataset by disabling each of its three core components: Temporal Stabilizing Fusion (TFS), Frequency Wave Modeling (FWM), and the dual-branch attention mechanism. As shown in Table 2, the removal of any single module leads to a clear degradation in both MSE and MAE. Without TFS, the model cannot better adapt to the non-stationary patterns while modeling long-term dependencies; omitting FWM undermines its ability to track changing periodic patterns. These findings confirm that each module contributes indispensably to handling different facets of non-stationarity and that their collaboration is critical to DTAF’s superior forecasting accuracy.

Model Analysis

Visualization of Expert Weights. We select three samples and show the expert weight distribution for each patch in Figure 3. Our observations reveal that the samples possess unique expert weight distributions. Both samples 1 and 3 share similar seasonality and similar trend patterns, which indicates that they may have similar non-stationarity, thereby yielding similar weight distributions. In contrast, the patterns of sample 2 diverge from samples 1 and 3. These findings underscore the adaptability of the Non-stationary MOE Filter and its ability to extract and distinguish different non-stationary patterns.




Figure 3: The weights of different experts for the ETTh1. P1, P2, P3, and P4 denote four patches from the sequence, while E1, E2, E3, and E4 represent four experts to extract non-stationary patterns.

Model	Metrics	DTAF		w/o TFS		w/o FWM		w/ Attention	
		mse	mae	mse	mae	mse	mae	mse	mae
NNS	24	0.716	0.559	0.723	0.567	0.729	0.574	0.725	0.571
	36	0.644	0.537	0.659	0.550	0.659	0.554	0.660	0.554
	48	0.621	0.535	0.623	0.540	0.632	0.549	0.635	0.552
	60	0.593	0.522	0.608	0.535	0.597	0.527	0.608	0.536
ETTh1	96	0.359	0.384	0.367	0.391	0.363	0.391	0.376	0.395
	192	0.400	0.414	0.408	0.421	0.407	0.419	0.418	0.430
	336	0.415	0.424	0.426	0.433	0.420	0.434	0.438	0.445
	720	0.421	0.449	0.457	0.468	0.446	0.460	0.450	0.463

Table 2: Ablation Study. W/O TFS, W/O FWM, represent removing the TFS and FWM, and W/ Attention represents replacing Dual-branch Attention with Cross Attention. Results averaged over 3 seeds; std $\leq \pm 0.0058$.

Non-stationary MOE Filter Analysis. We analyze the Non-stationary MOE Filter and carry out a visual analysis of the input and output of the Non-stationary MOE Filter, which is shown in Figure 4. It can be seen that before the Non-stationary MOE Filter, the distribution among patches had significant differences. However, after the Filter, the distribution similarity among patches is relatively high. This proves that the Non-stationary MOE Filter can indeed remove the Non-stationary part between patches, making the distribution between patches more stable.




Figure 4: Non-stationary MOE Filter Analysis.

TopK selection Analysis. To analyze whether the low-frequency and high-energy components would dominate the selection of TopK, we conduct an experimental analysis. As shown in the left figure of Figure 5, it is a sampled time series. We divided it into four patches and aligned them for frequency-domain differential operations, obtaining three differential spectrograms as shown in the right figure of Figure 5. We found that for the selection of TopK in the differential spectrograms, it is not dominated by low-frequency and high-energy components. It can be seen that in DIFF 1, low-frequency and high-frequency components are selected, while in DIFF 2 and DIFF 3, the intermediate frequency and high-frequency components are chosen. From this, it can be seen that the model’s selection of frequencies is diverse and is not dominated by low-frequency and high-energy components.




Figure 5: TopK selection Analysis.

Conclusion

In this paper, we propose a novel model **DTAF**, which effectively addresses the non-stationarity in both the temporal and frequency domains for time series forecasting. Specifically, our proposed Temporal Stabilizing Fusion (TFS) module is enhanced by a Non-stationary MOE Filter that automatically removes non-stationary patterns in the temporal domain and facilitates the capture of long-term dependencies. Meanwhile, the Frequency Wave Modeling (FWM) module performs frequency domain differencing to handle non-stationarity in the spectral space, allowing the model to better understand periodic trends and sudden shifts. Based on these, DTAF achieves more robust and accurate forecasting across diverse scenarios.

Acknowledgments

This work was partially supported by the National Natural Science Foundation of China (62372179, 62406112), Scientific and Technological Program of Xi'an under Grant 24GXF0016.

References

- Box, G. E. P.; Jenkins, G. M.; and Reinsel, G. C. 2015. *Time Series Analysis: Forecasting and Control*. Hoboken, New Jersey: Wiley, 5th edition.
- Campagne, E.; Amara-Ouali, Y.; Goude, Y.; and Kalogeratos, A. 2024. Leveraging Graph Neural Networks to Forecast Electricity Consumption. In *ECML PKDD Workshop*.
- chan Park, Y.; gi Jang, J.; and Kang, U. 2021. Fast and Accurate Partial Fourier Transform for Time Series Data. In *KDD*.
- Chen, L.; Chen, D.; Shang, Z.; Wu, B.; Zheng, C.; Wen, B.; and Zhang, W. 2023. Multi-Scale Adaptive Graph Neural Network for Multivariate Time Series Forecasting. In *TKDE*.
- Chen, P.; Wang, Y.; Shu, Y.; Cheng, Y.; Zhao, K.; Rao, Z.; Pan, L.; Yang, B.; and Guo, C. 2025. CC-Time: Cross-Model and Cross-Modality Time Series Forecasting. *arXiv preprint arXiv:2508.12235*.
- Chen, P.; Zhang, Y.; Cheng, Y.; Shu, Y.; Wang, Y.; Wen, Q.; Yang, B.; and Guo, C. 2024. Pathformer: Multi-scale Transformers with Adaptive Pathways for Time Series Forecasting. In *ICLR*.
- Cheng, Y.; Chen, P.; Guo, C.; Zhao, K.; Wen, Q.; Yang, B.; and Jensen, C. S. 2023. Weakly guided adaptation for robust time series forecasting. *Proceedings of the VLDB Endowment*, 17(4): 766–779.
- Cho, K.; van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *EMNLP*.
- Cirstea, R.-G.; Guo, C.; Yang, B.; Kieu, T.; Dong, X.; and Pan, S. 2022. Triformer: Triangular, Variable-Specific Attention for Long Sequence Multivariate Time Series Forecasting. In *IJCAI*.
- Cirstea, R.-G.; Kieu, T.; Guo, C.; Yang, B.; and Pan, S. J. 2021. EnhanceNet: Plugin Neural Networks for Enhancing Correlated Time Series Forecasting. In *ICDE*.
- Du, Y.; Wang, J.; Feng, W.; Pan, S.; Qin, T.; Xu, R.; and Wang, C. 2021. AdaRNN: Adaptive Learning and Forecasting for Time Series. In *CIKM*.
- Fan, W.; Yi, K.; Ye, H.; Ning, Z.; Zhang, Q.; and An, N. 2024. Deep Frequency Derivative Learning for Non-stationary Time Series Forecasting. In *IJCAI*.
- Fei, J.; Yi, K.; Fan, W.; Zhang, Q.; and Niu, Z. 2025. Amplifier: Bringing Attention to Neglected Low-Energy Components in Time Series Forecasting. In *AAAI*.
- Gao, H.; Shen, W.; Qiu, X.; Xu, R.; Yang, B.; and Hu, J. 2025. SSD-TS: Exploring the potential of linear state space models for diffusion models in time series imputation. In *SIGKDD*.
- Gasthaus, J.; Benidis, K.; Wang, Y.; Rangapuram, S. S.; Salinas, D.; Flunkert, V.; and Januschowski, T. 2019. Probabilistic Forecasting with Spline Quantile Function RNNs. In *AISTATS*.
- Hu, S.; Zhao, K.; Qiu, X.; Shu, Y.; Hu, J.; Yang, B.; and Guo, C. 2024. MultiRC: Joint Learning for Time Series Anomaly Prediction and Detection with Multi-scale Reconstructive Contrast. *arXiv preprint arXiv:2410.15997*.
- Ji, S.; Zhang, Z.; Ying, S.; Wang, L.; Zhao, X.; and Gao, Y. 2022. Kullback–Leibler Divergence Metric Learning. In *IEEE Transactions on Cybernetics*.
- Kim, T.; Kim, J.; Tae, Y.; Park, C.; Choi, J.-H.; and Choo, J. 2022. Reversible Instance Normalization for Accurate Time-Series Forecasting against Distribution Shift. In *ICLR*.
- Li, Z.; Qiu, X.; Chen, P.; Wang, Y.; Cheng, H.; Shu, Y.; Hu, J.; Guo, C.; Zhou, A.; Wen, Q.; et al. 2025. TSFM-Bench: A Comprehensive and Unified Benchmark of Foundation Models for Time Series Forecasting. In *SIGKDD*.
- Liang, X.; Wu, L.; Li, J.; Wang, Y.; Meng, Q.; Qin, T.; Chen, W.; Zhang, M.; and Liu, T.-Y. 2021. R-Drop: Regularized Dropout for Neural Networks. In *NeurIPS*.
- Liu, P.; Wu, B.; Hu, Y.; Li, N.; Dai, T.; Bao, J.; and tao Xia, S. 2025a. TimeBridge: Non-Stationarity Matters for Long-term Time Series Forecasting. In *ICML*.
- Liu, X.; Qiu, X.; Cheng, H.; Wu, X.; Guo, C.; Yang, B.; and Hu, J. 2025b. ASTGI: Adaptive Spatio-Temporal Graph Interactions for Irregular Multivariate Time Series Forecasting. *arXiv preprint arXiv:2509.23313*.
- Liu, Y.; Hu, T.; Zhang, H.; Wu, H.; Wang, S.; Ma, L.; and Long, M. 2024. iTransformer: Inverted Transformers Are Effective for Time Series Forecasting. *ICLR*.
- Liu, Y.; Liu, Q.; Zhang, J.-W.; Feng, H.; Wang, Z.; Zhou, Z.; and Chen, W. 2022a. Multivariate Time-Series Forecasting with Temporal Polynomial Graph Neural Networks. In *NeurIPS*.
- Liu, Y.; Wu, H.; Wang, J.; and Long, M. 2022b. Non-stationary Transformers: Exploring the Stationarity in Time Series Forecasting. In *NeurIPS*.
- Liu, Z.; Cheng, M.; Li, Z.; Huang, Z.; Liu, Q.; Xie, Y.; and Chen, E. 2023. Adaptive Normalization for Non-stationary Time Series Forecasting: A Temporal Slice Perspective. In *NeurIPS*.
- Liu, Z.; Wang, Y.; Vaidya, S.; Ruehle, F.; Halverson, J.; Soljačić, M.; Hou, T. Y.; and Tegmark, M. 2025c. KAN: Kolmogorov-Arnold Networks. In *ICLR*.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *ICLR*.
- Nie, Y.; H. Nguyen, N.; Sinthong, P.; and Kalagnanam, J. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *ICLR*.
- Ogasawara, E.; Martinez, L. C.; de Oliveira, D.; Zimbrão, G.; Pappa, G. L.; and Mattoso, M. 2010. Adaptive Normalization: A novel data normalization approach for non-stationary time series. In *IJCNN*.

- Passalis, N.; Tefas, A.; Kannianen, J.; Gabbouj, M.; and Iosifidis, A. 2020. Deep Adaptive Input Normalization for Time Series Forecasting. In *TNNLS*.
- Pillai, S. 2025. Graph-Augmented LSTM for Forecasting Sparse Anomalies in Graph-Structured Time Series. In *arXiv*.
- Prabowo, A.; Xue, H.; Shao, W.; Koniusz, P.; and Salim, F. D. 2023. Message Passing Neural Networks for Traffic Forecasting. In *arXiv*.
- Qiu, X.; Cheng, H.; Wu, X.; Hu, J.; and Guo, C. 2025a. A Comprehensive Survey of Deep Learning for Multivariate Time Series Forecasting: A Channel Strategy Perspective. *arXiv preprint arXiv:2502.10721*.
- Qiu, X.; Cheng, H.; Wu, X.; Hu, J.; Guo, C.; and Yang, B. 2025b. A Comprehensive Survey of Deep Learning for Multivariate Time Series Forecasting: A Channel Strategy Perspective. In *arXiv*.
- Qiu, X.; Hu, J.; Zhou, L.; Wu, X.; Du, J.; Zhang, B.; Guo, C.; Zhou, A.; Jensen, C. S.; Sheng, Z.; and Yang, B. 2024. TFB: Towards Comprehensive and Fair Benchmarking of Time Series Forecasting Methods. *VLDB*, 17(9): 2363–2377.
- Qiu, X.; Li, X.; Pang, R.; Pan, Z.; Wu, X.; Yang, L.; Hu, J.; Shu, Y.; Lu, X.; Yang, C.; Guo, C.; Zhou, A.; Jensen, C. S.; and Yang, B. 2025c. EasyTime: Time Series Forecasting Made Easy. In *ICDE*.
- Qiu, X.; Wu, X.; Cheng, H.; Liu, X.; Guo, C.; Hu, J.; and Yang, B. 2025d. DBLoss: Decomposition-based Loss Function for Time Series Forecasting. In *NeurIPS*.
- Qiu, X.; Wu, X.; Lin, Y.; Guo, C.; Hu, J.; and Yang, B. 2025e. DUET: Dual Clustering Enhanced Multivariate Time Series Forecasting. In *SIGKDD*, 1185–1196.
- Qiu, X.; Yang, L.; Cheng, H.; Wu, X.; Wu, R.; Zhang, Z.; Tu, D.; Guo, C.; Yang, B.; Jensen, C. S.; et al. 2025f. Multi-Scale Spatial-Temporal Hypergraph Network with Lead-Lag Structures for Stock Time Series Forecasting. *arXiv preprint arXiv:2509.23668*.
- Qiu, X.; Zhu, Y.; Li, Z.; Cheng, H.; Wu, X.; Guo, C.; Yang, B.; and Hu, J. 2025g. DAG: A dual causal network for time series forecasting with exogenous variables. *arXiv preprint arXiv:2509.14933*.
- Rangapuram, S. S.; Seeger, M.; Gasthaus, J.; Stella, L.; Wang, Y.; and Januschowski, T. 2018. Deep state space models for time series forecasting. In *NeurIPS*.
- Tian, J.; Liang, Y.; Xu, R.; Chen, P.; Guo, C.; Zhou, A.; Pan, L.; Rao, Z.; and Yang, B. 2024. Air Quality Prediction with Physics-Informed Dual Neural ODEs in Open Systems. *CoRR*, abs/2410.19892.
- Wang, Q.; Cheng, M.; Yuan, S.; and Xu, H. 2023. Hierarchical contrastive learning for temporal point processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 10166–10174.
- Wang, Q.; Wu, Y.; Long, Y.; Huang, J.; Ran, F.; Su, B.; and Xu, H. 2025a. A Plug-and-Play Bregman ADMM Module for Inferring Event Branches in Temporal Point Processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 21216–21224.
- Wang, S.; Li, J.; Shi, X.; Ye, Z.; Mo, B.; Lin, W.; Ju, S.; Chu, Z.; and Jin, M. 2024. TimeMixer++: A General Time Series Pattern Machine for Universal Predictive Analysis. In *ICLR*.
- Wang, Y.; Qiu, Y.; Chen, P.; Shu, Y.; Rao, Z.; Pan, L.; Yang, B.; and Guo, C. 2025b. LightGTS: A Lightweight General Time Series Forecasting Model. *arXiv preprint arXiv:2506.06005*.
- Wang, Y.; Qiu, Y.; Chen, P.; Zhao, K.; Shu, Y.; Rao, Z.; Pan, L.; Yang, B.; and Guo, C. 2025c. Towards a General Time Series Forecasting Model with Unified Representation and Adaptive Transfer. In *Forty-second International Conference on Machine Learning*.
- Wu, H.; Hu, T.; Liu, Y.; Zhou, H.; Wang, J.; and Long, M. 2023. TimesNet: Temporal 2D-Variation Modeling for General Time Series Analysis. In *ICLR*.
- Wu, H.; Xu, J.; Wang, J.; and Long, M. 2021. Autoformer: Decomposition Transformers with Auto-Correlation for Long-Term Series Forecasting. In *NeurIPS*.
- Wu, X.; Li, Z.; Cheng, H.; Qiu, X.; Hu, J.; Guo, C.; and Yang, B. 2025a. Unlocking the Power of Mixture-of-Experts for Task-Aware Time Series Analytics. *arXiv preprint arXiv:2509.22279*.
- Wu, X.; Qiu, X.; Cheng, H.; Li, Z.; Hu, J.; Guo, C.; and Yang, B. 2025b. Enhancing time series forecasting through selective representation spaces: A patch perspective. *arXiv preprint arXiv:2510.14510*.
- Wu, X.; Qiu, X.; Gao, H.; Hu, J.; Yang, B.; and Guo, C. 2025c. K²VAE: A Koopman-Kalman Enhanced Variational AutoEncoder for Probabilistic Time Series Forecasting. In *ICML*.
- Wu, X.; Qiu, X.; Li, Z.; Wang, Y.; Hu, J.; Guo, C.; Xiong, H.; and Yang, B. 2025d. CATCH: Channel-Aware multivariate Time Series Anomaly Detection via Frequency Patching. In *ICLR*.
- Wu, X.; Wu, X.; Yang, B.; Zhou, L.; Guo, C.; Qiu, X.; Hu, J.; Sheng, Z.; and Jensen, C. S. 2024. AutoCTS++: zero-shot joint neural architecture and hyperparameter search for correlated time series forecasting. *VLDB J.*, 33(5): 1743–1770.
- Xiong, Y.; and Wen, Y. 2025. Non-Stationary Time Series Forecasting Based on Fourier Analysis and Cross Attention Mechanism. In *IJCNN*.
- Xu, Z.; Zeng, A.; and Xu, Q. 2024. FITS: Modeling Time Series with 10k Parameters. In *ICLR*.
- Ye, H.; Ding, H.; Chen, S.; Jiang, Y.; Changyuan, Z.; and Zhang, X. 2025. Aligning by Misaligning: Boundary-aware Curriculum Learning for Multimodal Alignment. In *NeurIPS*.
- Ye, W.; Xu, Z.; and Gui, N. 2025. Non-stationary Diffusion For Probabilistic Time Series Forecasting. In *ICML*.
- Zeng, A.; Chen, M.; Zhang, L.; and Xu, Q. 2023. Are Transformers Effective for Time Series Forecasting? In *AAAI*.
- Zhang, Y.; and Yan, J. 2023. Crossformer: Transformer Utilizing Cross-Dimension Dependency for Multivariate Time Series Forecasting. In *ICLR*.