

T3Time: Tri-Modal Time Series Forecasting via Adaptive Multi-Head Alignment and Residual Fusion

Abdul Monaf Chowdhury¹, Rabeya Akter¹, Safaeid Hossain Arib¹

¹Robotics & Mechatronics Engineering, University of Dhaka
{monafabdul15, rabeyaakter231023, safaeid48}@gmail.com

Abstract

Multivariate time series forecasting (MTSF) seeks to model temporal dynamics among variables to predict future trends. Transformer-based models and large language models (LLMs) have shown promise due to their ability to capture long-range dependencies and patterns. However, current methods often rely on rigid inductive biases, ignore inter-variable interactions, or apply static fusion strategies that limit adaptability across forecast horizons. These limitations create bottlenecks in capturing nuanced, horizon-specific relationships in time-series data. To solve this problem, we propose **T3Time**, a novel trimodal framework consisting of time, spectral, and prompt branches, where the dedicated frequency encoding branch captures the periodic structures along with a gating mechanism that learns prioritization between temporal and spectral features based on the prediction horizon. We also proposed a mechanism which adaptively aggregates multiple cross-modal alignment heads by dynamically weighting the importance of each head based on the features. Extensive experiments on benchmark datasets demonstrate that our model consistently outperforms state-of-the-art baselines, achieving an average reduction of 3.28% in MSE and 2.29% in MAE. Furthermore, it shows strong generalization in few-shot learning settings: with 5% training data, we see a reduction in MSE and MAE by 4.13% and 1.91%, respectively; and with 10% data, by 3.62% and 1.98% on average. Code will be released upon publication.

Introduction

Multivariate time-series forecasting (MTSF) lies at the heart of modern decision-making, powering everything from energy-load balancing (Liu et al. 2023a) and urban traffic management (Liu et al. 2024c) to high-frequency trading (Xu et al. 2021) and weather forecasting (Schneider and Dickinson 1974). While the objective appears straightforward—predicting future values based on past observations—the underlying challenge is profoundly complex. Effective models must simultaneously capture short-term temporal fluctuations, long-range dependencies, and intricate inter-variable dynamics, all while maintaining computational efficiency and robustness in data-sparse regimes.

Recent deep learning advances have led to the development of numerous models for time series forecasting (Miao

et al. 2024; Challu et al. 2023). However, early approaches were constrained by limited parameters and poor representations (Liu et al. 2022a), hindering generalization. Following the development of transformer-based architectures (Vaswani et al. 2017), frameworks like Informer (Zhou et al. 2021), Autoformer (Wu et al. 2021), and Fedformer (Zhou et al. 2022) addressed the quadratic complexity of standard self-attention and introduced seasonal decomposition.

Subsequent efforts have advanced in two complementary directions for better representation learning in the subspace of time-series. The first focuses on token restructuring, exemplified by PatchTST (Nie et al. 2023), which segments time series into sub-series patches and processes each channel independently. The second explores alternative representation spaces; for instance, Freeformer (Yue et al. 2025) operates in the frequency domain, demonstrating that spectral tokens can encode global periodic patterns more compactly than conventional time-domain attention. However, single-modality encoder architectures remain limited in their capacity to fully capture the intricate and multi-scale structure of temporal dynamics.

To make contextual time-series representations even robust, pre-trained LLM based (Radford et al. 2018) prompting techniques have been utilized in some frameworks. Prompts are either used to encapsulate the time-series information (Jia et al. 2024; Huang et al. 2024) or to provide further context to the time series (Jin et al. 2023; Xue and Salim 2023; Xue, Voutharoja, and Salim 2022). Although time and prompt-based dual modal embeddings have accomplished promising performance, they have been plagued with embedding overlapping issues (Chang et al. 2024; Jia et al. 2024), which weakened representation integrity as seen here by Fig. 1(a).

Recently, TimeCMA (Liu et al. 2025) has presented cross-modal alignment to integrate temporal-language dual modalities to represent embeddings as disentangled yet robust. Regardless, they miss out on aligning spectral modality to further represent global periodicity. By and large, existing MTSF frameworks exhibit three fundamental limitations. First, existing models often adopt a modality-isolated architecture, emphasizing a single modality while disregarding the other, which leads to fragmented representations. Second, if they utilize multi-modalities, they suffer from limited alignment capacity that constrain the model’s ability to




Figure 1: Comparison of bimodal vs. tri-modal framework for time series forecasting. (a) Bimodal models use static fusion of time and prompt features, lacking frequency awareness and horizon adaptivity. (b) T3Time introduces tri-modal encoding with horizon-aware gating and adaptive multi-head cross-modal alignment for robust, horizon-sensitive representations.

capture rich, fine-grained interactions between multi modalities. Third, these models exhibit horizon rigidity by applying static processing across forecast lengths, ultimately hindering their ability to adapt modality emphasis based on the temporal scope of the prediction.

To address these challenges, we propose **T3Time** - illustrated in Fig. 1(b), a tri-modal framework for MTSF that integrates temporal, spectral, and prompt-based semantic representations. Each modality captures complementary structures in the data. To enable effective multimodal integration, T3Time employs an adaptive multi-head cross-modal alignment module that dynamically weighs modality-specific embeddings based on their relevance to the forecasting task. Additionally, a horizon-aware gating mechanism modulates the influence of temporal and spectral features by prediction horizon, while a channel-wise residual fusion preserves variable-specific priors and enhances representational granularity. These components collectively enable T3Time to learn horizon-aware, variable-sensitive representations for robust generalization and improved forecasting. Our contributions are summarized as follows:

- We propose **T3Time**, a novel tri-modal forecasting framework that unifies temporal, spectral, and prompt-based semantic representations via an adaptive multi-head cross-modal alignment mechanism. This design enables dynamic fusion of heterogeneous modalities for more expressive and context-aware representation.
- We introduce horizon aware gating module and channel-wise residual fusion mechanism to elevate temporal adaptability and fine-grained feature representation.
- T3Time consistently outperforms state-of-the-art baselines and demonstrates strong generalization across benchmark datasets.

Related Work

Forecasting in Time Domain. Time series forecasting has witnessed significant advances driven by deep learning (Lai et al. 2018; Franceschi, Dieuleveut, and Jaggi 2019; Jin et al. 2022), with Transformer-based architectures (Wen et al. 2022) emerging as powerful tools due to their capacity to model long-range dependencies. Early transformer-based methods (Zhou et al. 2021; Wu et al. 2021; Liu et al.

2022b) introduce various attention mechanisms to reduce the quadratic complexity of standard self-attention and improve the efficiency of forecasting. However, most of them rely on fixed inductive biases (e.g., decomposition, sparse priors) and can not adapt to varying forecasting horizons. iTransformer (Liu et al. 2023b) models each time series as an independent token, enabling flexible cross-variable attention. Regardless, its reliance on a single-stream encoder and static fusion limits its ability to capture complex, modality-specific dynamics in multivariate forecasting.

Forecasting in Frequency Domain. Frequency-domain methods (Cao et al. 2020; Woo et al. 2022; Sun and Boning 2022; Yi et al. 2023; Chen et al. 2023) offer another perspective by modeling temporal periodicity and long-term structure. Autoformer (Wu et al. 2021) integrates series decomposition with autocorrelation in the frequency domain, while FEDformer (Zhou et al. 2022) further enhances this approach employing a mixture-of-expert structure and Fourier-based attention. Although effective, these methods typically employ fixed or global spectral representations, lacking mechanisms for adapting frequency-domain importance based on forecast length or contextual features.

Cross-modal Alignment. Recent work has explored adapting large language models (LLMs) (Lu et al. 2021) for time series forecasting, either by replacing standard tokenizers with learned embeddings for time series inputs (Zhou et al. 2023; Liu et al. 2024a) or by formatting time series as textual prompts (Xue and Salim 2023; Jin et al. 2023; Pan et al. 2024). Although these methods leverage the representational power of LLMs, they often face modality mismatches. Moreover, fusion between language and numerical representations is typically static, limiting adaptability across tasks. Several studies have proposed the transfer of knowledge from pre-trained models using self-supervised learning (Zhang et al. 2022; Deldari et al. 2022; Zhang et al. 2024), multimodal reprogramming (Chen 2024), or instruction-based fine-tuning (Yin et al. 2024). TimeCMA (Liu et al. 2025) represents a recent effort in this direction, introducing a cross-modal alignment to integrate time series and prompt embeddings for multivariate forecasting. However, its use of a single head alignment mechanism can be limiting, as it constrains the model’s ability to capture diverse and fine-grained interactions between semantic and temporal signals. To this end, we introduce trimodal encodings to better represent time series representations and ensure robust performance.

Methodology

Our three-stage framework, as illustrated in Figure 2 contains Tri-Modal Encoding, Adaptive Multi-Head Cross-Modal Alignment (CMA), and channel-wise residual connection. **Tri-Modal Encoding** comprises three branches: the frequency encoding branch, the time series encoding branch and the LLM encoding branch, designed to extract three different representations from the input time series. We also designed a representation-rich horizon-gating to fuse the frequency and time encoder branches dynamically based on the forecast horizon. **Adaptive Multi-head CMA**, inspired by (Vaswani et al. 2017), aligns the fused temporal-

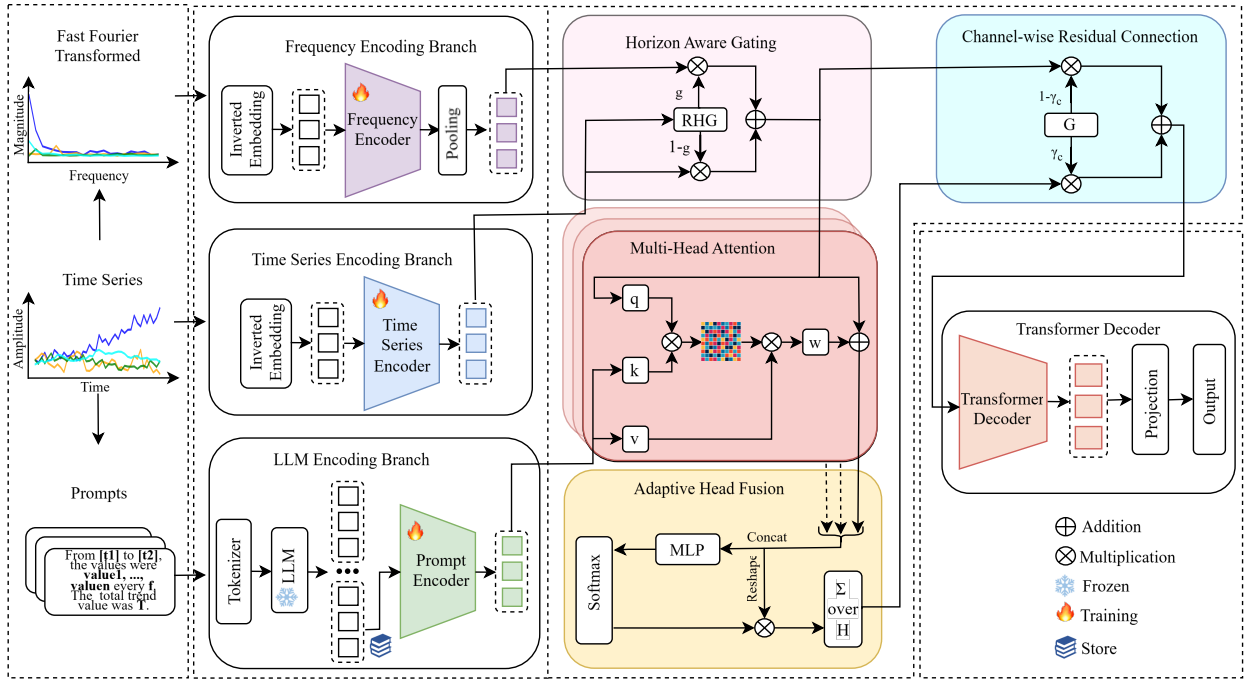


Figure 2: Overview of our framework. The model comprises tri-modal encoding (time, frequency, prompt), a horizon-aware gating module for dynamic temporal-spectral fusion, adaptive multi-head cross-modal alignment with per-head importance weighting, and a channel-wise residual connection for fine-grained representation mixing prior to decoding.

spectral features with the prompt embeddings using multiple instances of a Cross-Modal Attention module. Each attention head independently computes cross-attention between the fused temporal-spectral embeddings and prompt embeddings. Adaptive head fusion mechanism combines all the attention heads based on their importance and generates an expressive aligned representation. Finally, **Channel-wise Residual Connection** creates a fine-grained residual fusion between original temporal-spectral embeddings and aligned cross-modal representations, which are ultimately passed through the decoder for forecasting.

Tri-Modal Encoding

Frequency Encoding Branch: To capture periodic and frequency-aware patterns from the time series, we designed a dedicated Frequency Encoding Branch, which transforms raw temporal inputs into the frequency domain using the real-valued fast Fourier transform and processes the resulting spectral features through a Transformer-based encoder.

Given the normalized input $\mathbf{X}_t \in \mathbb{R}^{B \times N \times L}$, where B is the batch size, N is the number of variables (or nodes), and L is the sequence length, we apply the real-valued Fourier transform along the temporal dimension to obtain complex-valued spectra:

$$\hat{\mathbf{X}}_t = \mathcal{F}_r(\mathbf{X}_t) \in \mathbb{C}^{B \times N \times L_f}, \quad L_f = \left\lfloor \frac{L}{2} \right\rfloor + 1 \quad (1)$$

Only the magnitude spectrum, \mathbf{F} , is retained as input to the frequency encoder. Each frequency bin is treated as a token. The tensor \mathbf{F} is reshaped to $(BN) \times L_f$ and projected

to the embedding dimension C using a learnable projection matrix $\mathbf{W}_f \in \mathbb{R}^{C \times 1}$:

$$\mathbf{Z}_f = \phi(\mathbf{F} \cdot \mathbf{W}_f^T) \in \mathbb{R}^{BN \times L_f \times C} \quad (2)$$

where $\phi(\cdot)$ denotes an element-wise ReLU nonlinearity.

To model dependencies across frequency components, the projected frequency tokens are passed through a single-layer Transformer encoder. Let, \mathcal{T} denote the self-attention block with pre-normalization. The encoded representation is given by:

$$\tilde{\mathbf{Z}}_f = \mathcal{T}(\mathbf{Z}_f) \in \mathbb{R}^{BN \times L_f \times C} \quad (3)$$

To aggregate the encoded spectral features, we compute a learnable attention-weighted pooling over the frequency bins to summarize the frequency information for each feature. Specifically, we first project the encoded frequency tokens $\tilde{\mathbf{Z}}_f \in \mathbb{R}^{BN \times L_f \times C}$ through a two-layer perceptron consisting of weights $\mathbf{W}_1 \in \mathbb{R}^{C \times d}$ and $\mathbf{W}_2 \in \mathbb{R}^{d \times 1}$, with an intermediate nonlinearity $\phi(\cdot)$. The resulting scalar logits are normalized via a softmax function across the frequency dimension L_f , yielding attention weights $\alpha \in \mathbb{R}^{BN \times L_f \times 1}$:

$$\alpha = \frac{\exp\left(\left[\phi\left(\tilde{\mathbf{Z}}_f \mathbf{W}_1\right)\right] \mathbf{W}_2\right)}{\sum_{j=1}^{L_f} \exp\left(\left[\phi\left(\tilde{\mathbf{Z}}_f \mathbf{W}_1\right)\right] \mathbf{W}_2\right)_j} \in \mathbb{R}^{BN \times L_f \times 1} \quad (4)$$

The final pooled embedding is computed as a weighted

sum:

$$\mathbf{F}_{\text{pooled}} = \sum_{l=1}^{L_f} \alpha_l \cdot \tilde{\mathbf{Z}}_{f, :, l} \in \mathbb{R}^{B \times N \times C} \quad (5)$$

Eventually, the pooled frequency features are reshaped back to match the node dimension, $\tilde{\mathbf{F}} \in \mathbb{R}^{B \times N \times C}$. This frequency-aware representation $\tilde{\mathbf{F}} \in \mathbb{R}^{B \times N \times C}$ captures periodic patterns and serves as one of the three modality-specific encodings for subsequent fusion.

Time Series Encoding Branch: To model temporal dependencies and evolving patterns in raw time-domain signals, we construct a dedicated **Time Series Encoding Branch**. This branch transforms normalized time series into contextualized representations using a shared projection followed by a Transformer-based encoder.

Given the normalized input $\mathbf{X}_t \in \mathbb{R}^{B \times N \times L}$, we project the temporal dimension into a latent embedding space of dimension C using a shared learnable projection matrix $\mathbf{W}_t \in \mathbb{R}^{L \times C}$:

$$\mathbf{Z}_t = \mathbf{X}_t \mathbf{W}_t \in \mathbb{R}^{B \times N \times C} \quad (6)$$

This projection treats each node’s input sequence as a single vector, linearly embedding the temporal axis into C features per node. To model interactions and dependencies across temporal patterns within each node, we apply a Transformer encoder with pre-normalization. Let \mathcal{T}_t denote the time-domain Transformer encoder:

$$\tilde{\mathbf{Z}}_t = \mathcal{T}_t(\mathbf{Z}_t) \in \mathbb{R}^{B \times N \times C} \quad (7)$$

$\tilde{\mathbf{Z}}_t$ provides positionally-aware, temporally contextualized embeddings for each node and serves as time modality-specific representation in our tri-modal framework.

LLM Encoding Branch: To inject external priors and semantic structure into the forecasting model, we introduce a dedicated **LLM Encoding Branch**, which leverages a pre-trained frozen GPT-2 (Radford et al. 2019) to encode prompt-based descriptions of input time series segments.

Given the normalized input series $\mathbf{X}_t \in \mathbb{R}^{B \times N \times L}$ and their associated temporal markers $\mathbf{M}_t \in \mathbb{R}^{B \times L \times D}$ (e.g., date, hour), we generate natural language prompts describing the sequence statistics. Each prompt is constructed for every feature. Prompt generation strategy is further discussed in Appendix B-1.

Each prompt is tokenized using a GPT-2 tokenizer and fed into a pre-trained GPT-2 model. Let $\mathcal{L}(\cdot)$ denote the language model. For a given tokenized prompt $\mathbf{P}_{i,j}$ associated with the j -th node in the i -th sample, we obtain:

$$\mathbf{E}_{i,j} = \mathcal{L}(\mathbf{P}_{i,j}) \in \mathbb{R}^{T \times d_{\text{LLM}}} \quad (8)$$

where T is the number of tokens in the prompt and d_{LLM} is the embedding dimension of the language model (e.g., 768 for GPT-2). Since different prompts may produce variable-length token sequences, we uniformly pad the output with copies of the final token embedding to ensure consistent shape across all samples and features.

After batching, we extract the final token embedding per prompt as a compact summary:

$$\mathbf{Z}_{\text{LLM}}[i, j] = \mathbf{E}_{i,j}[-1] \in \mathbb{R}^{d_{\text{LLM}}} \quad (9)$$

The final LLM-derived representation is a tensor $\mathbf{Z}_{\text{LLM}} \in \mathbb{R}^{B \times N \times d_{\text{LLM}}}$, which is projected and passed through a Transformer encoder before being used in downstream cross-modal fusion.

We introduce a **Horizon-Aware Gating Module** to adaptively balance contributions from different modalities based on the forecast horizon. The core intuition is that short-term forecasts may benefit more from time-localized representations, whereas long-range forecasts can better leverage global periodic patterns captured in the frequency domain.

Time encoding, $\tilde{\mathbf{Z}}_t \in \mathbb{R}^{B \times N \times C}$, is first pooled over the feature dimension to obtain a global summary per sample. We normalize the prediction length by a constant factor and concatenate it with the pooled time encoding in order to provide the forecast length as a continuous conditioning signal, $\mathbf{g}_{\text{in}} \in \mathbb{R}^{B \times (C+1)}$. The concatenated vector, \mathbf{g}_{in} , is processed through a lightweight two-layer MLP followed by a sigmoid nonlinearity to produce channel-wise gating weights:

$$\mathbf{g} = \sigma(\mathbf{W}_4 \cdot \phi(\mathbf{W}_3 \cdot \mathbf{g}_{\text{in}}^{\top}))^{\top} \in \mathbb{R}^{B \times C} \quad (10)$$

where $\mathbf{W}_3 \in \mathbb{R}^{(C+1) \times d}$, $\mathbf{W}_4 \in \mathbb{R}^{d \times C}$, $\sigma(\cdot)$ denotes the element-wise sigmoid nonlinearity.

The output of this module is a horizon-aware convex combination of the frequency and time representations:

$$\mathbf{Z}_{\mathbf{g}} = \mathbf{g} \odot \tilde{\mathbf{F}} + (1 - \mathbf{g}) \odot \tilde{\mathbf{Z}}_t \in \mathbb{R}^{B \times C \times N} \quad (11)$$

where \odot denotes element-wise multiplication broadcast over the feature dimension. This gating mechanism enables the model to adaptively shift focus between temporally localized and spectrally global features, as a function of both the input content and the desired forecast horizon.

Adaptive Dynamic Head Cross-Modal Alignment

To integrate heterogeneous contextual representations from the time-spectral and semantic domains, we follow the cross-modal alignment strategy introduced in (Liu et al. 2025), wherein the time series encoder output is aligned with the prompt encoder output through cross-attention. However, instead of relying on a single head cross-modal alignment as in the original approach, we extend this paradigm by introducing *adaptive dynamic head fusion*, where multiple CMA heads are independently learned and their outputs are dynamically fused based on data-dependent gating scores.

Each CMA head h maps the fused representation from the time and frequency branches, $\mathbf{Z}_{\mathbf{g}} \in \mathbb{R}^{B \times C \times N}$, and prompt embeddings, $\mathbf{Z}_{\text{LLM}} \in \mathbb{R}^{B \times E \times \tilde{N}}$, into a head-specific aligned output $\mathbf{H}^{(h)} \in \mathbb{R}^{B \times C \times N}$ via an independent cross-attention mechanism with queries from $\mathbf{Z}_{\mathbf{g}}$ and keys/values from \mathbf{Z}_{LLM} . Rather than aggregating these heads through static averaging or fixed linear projections, head-wise outputs are aggregated via a *feature-aware, head-adaptive* fusion to dynamically weight the importance of each head per feature. Necessarily, the outputs are first concatenated along the channel dimension and transposed to shape $\mathbb{R}^{B \times N \times HC}$:

$$\mathbf{U} = [\mathbf{H}^{(1)}; \dots; \mathbf{H}^{(H)}]^{\top} \in \mathbb{R}^{B \times N \times HC} \quad (12)$$

Each node’s fused embedding vector $\mathbf{U}_{b,n} \in \mathbb{R}^{HC}$ (for batch index b and node index n) is passed through a two-layer gating network to compute importance scores $\pi_{b,n} \in \mathbb{R}^H$ over the CMA heads:

$$\mathbf{e}_{b,n} = \mathbf{W}_6 \phi(\text{LN}(\mathbf{W}_5 \mathbf{U}_{b,n}^\top)) \in \mathbb{R}^H, \quad (13)$$

$$\pi_{b,n}^{(h)} = \frac{\exp(e_{b,n}^{(h)})}{\sum_{j=1}^H \exp(e_{b,n}^{(j)})}, \quad \sum_{h=1}^H \pi_{b,n}^{(h)} = 1 \quad (14)$$

where $\mathbf{W}_5 \in \mathbb{R}^{128 \times HC}$ and $\mathbf{W}_6 \in \mathbb{R}^{H \times 128}$ are learnable matrices, $\phi(\cdot)$ is a pointwise ReLU nonlinearity, and $\text{LN}(\cdot)$ denotes layer normalization.

Let $\mathbf{H}_{b,:n}^{(h)} \in \mathbb{R}^C$ be the h -th head output for sample b and node n . The gated fusion is obtained via a convex combination over the head dimension, weighted by the attention scores $\pi_{b,n}^{(h)}$:

$$\mathbf{\Lambda}_{b,:n} = \sum_{h=1}^H \pi_{b,n}^{(h)} \cdot \mathbf{H}_{b,:n}^{(h)} \in \mathbb{R}^C \quad (15)$$

yielding the final cross-modally aligned representation $\mathbf{\Lambda} \in \mathbb{R}^{B \times C \times N}$ across all B samples and N features.

Channel-wise Residual Connection

Before decoding, we apply a channel-wise residual fusion to reconcile horizon-aware spectral-temporal features with cross-modal alignment outputs, allowing each latent feature to balance its dependence on intrinsic patterns and external priors.

Here, $\mathbf{Z}_g \in \mathbb{R}^{B \times C \times N}$ denotes the horizon-gated fusion of time-domain and frequency-domain encodings, and $\mathbf{\Lambda} \in \mathbb{R}^{B \times C \times N}$ denotes the output of the adaptive multi-head cross-modal alignment. We learn a set of trainable channel-wise residual coefficients $\gamma \in \mathbb{R}^C$ that modulate the importance of the two streams during the fusion process.

The fused representation $\Theta \in \mathbb{R}^{B \times C \times N}$ is computed as a convex combination along the channel axis:

$$\Theta_{b,c,n} = \gamma_c \odot \mathbf{\Lambda}_{b,c,n} + (1 - \gamma_c) \odot \mathbf{Z}_{g,b,c,n} \quad (16)$$

where \odot denotes element-wise multiplication, b , c , and n index the batch, channel, and node dimensions respectively, and $\gamma_c \in [0, 1]$ is a learnable scalar specific to channel c .

This $\Theta \in \mathbb{R}^{B \times C \times N}$ formulation allows each latent dimension to adaptively balance cross-modal information against temporal and spectral evidence, enabling fine-grained control over representational mixing. The fused representation Θ is subsequently passed into the decoder for final forecasting.

Decoder: Fused representation $\Theta \in \mathbb{R}^{B \times C \times N}$, which encapsulates spectro-temporal priors and semantically grounded alignment, is transposed to match the input format of the transformer decoder. The decoder module, composed of D stacked multi-head cross-attention blocks with pre-normalization, integrates global dependencies across nodes and channels.

Let $\mathcal{D}(\cdot)$ denote the Transformer decoder. The contextualized output $\mathbf{Z}_d \in \mathbb{R}^{B \times N \times C}$ is given by:

$$\mathbf{Z}_d = \mathcal{D}(\Theta^\top, \Theta^\top), \quad \Theta^\top \in \mathbb{R}^{B \times N \times C} \quad (17)$$

Following decoding, the representation \mathbf{Z}_d is linearly projected along the channel axis to produce the final forecast sequence of length L_p :

$$\hat{\mathbf{Y}} = \mathbf{Z}_d \mathbf{W}_p^\top + \mathbf{b}_p, \quad \mathbf{W}_p \in \mathbb{R}^{L_p \times C}, \quad \hat{\mathbf{Y}} \in \mathbb{R}^{B \times N \times L_p} \quad (18)$$

Forecast tensor $\hat{\mathbf{Y}}$ is transposed to match the original input format $\mathbf{Y} \in \mathbb{R}^{B \times L_p \times N}$, representing the final output.

Experiments

T3Time demonstrates consistent and superior performance across a wide range of benchmarks, particularly excelling in long-horizon forecasting and low-data regimes. To ensure fair comparison, we adhere strictly to the experimental protocol established by Liu et al. (2025) across all baseline models unless explicitly stated otherwise. Our evaluation includes a comprehensive set of strong baselines: *prompt-based LLMs* such as TimeCMA (Liu et al. 2025), Time-LLM (Jin et al. 2023), and UniTime (Liu et al. 2024b); *Transformer-based architectures* including iTransformer (Liu et al. 2023b), OFA (GPT4TS) (Zhou et al. 2023), PatchTST (Nie et al. 2023), and Fedformer (Zhou et al. 2022); the *linear* model DLinear (Zeng et al. 2023); and the *CNN-based* framework TimesNet (Wu et al. 2022). Details regarding device, datasets, and evaluation metrics are provided in Appendix A.

Long-Term Forecasting

Setups. We evaluate T3Time on eight widely-used multivariate time series benchmarks: ETTh1, ETTh2, ETTm1, ETTm2, ECL, Weather, ILI, and Exchange. Following the standardized protocol from Liu et al. (2025), we set the input sequence length to 96 and vary the forecasting horizon across $\{96, 192, 336, 720\}$, except for ILI where we follow the $\{24, 36, 48, 60\}$ prediction steps. Mean squared error (MSE) and mean average error (MAE) are used as evaluation metrics for all our experiments.

Results. Table 1 summarizes the forecasting performance across all datasets. On average, T3Time consistently achieves state-of-the-art (SOTA) results in 14 out of 16 baselines. The model yields the lowest error across most horizons, outperforming recent competitive approaches such as TimeLLM, (Jin et al. 2023) and iTransformer, (Liu et al. 2023b) by 11.28% in MSE and 6.20% in MAE, and 8.86% in MSE and 6.10% in MAE, respectively. When compared to the strongest prompt-based model, TimeCMA, (Liu et al. 2025), T3Time reduces the MSE by up to 4.36% on average, while demonstrating more stable MAE performance across long horizons. Overall, T3Time achieves an average MSE reduction of 3.28% and an MAE reduction of 2.29% compared to the SOTA baselines. A comprehensive overview of our full results is provided in Appendix C.

Dataset	Ours		TimeCMA		TimeLLM		UniTime		TimesNet		DLinear		iTransformer		PatchTST		OFA	
	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTm1	0.372	<u>0.393</u>	<u>0.380</u>	0.392	0.410	0.409	0.385	0.399	0.400	0.406	0.403	0.407	0.407	0.410	0.392	0.402	0.396	0.401
ETTm2	<u>0.279</u>	0.322	0.275	0.323	0.296	0.340	0.293	0.334	0.291	0.333	0.350	0.401	0.288	0.332	0.285	<u>0.328</u>	0.294	0.339
ETTh1	0.418	0.430	<u>0.423</u>	<u>0.431</u>	0.448	0.443	0.442	0.448	0.458	0.450	0.456	0.452	0.454	0.447	0.463	0.449	0.457	0.450
ETTh2	0.348	0.390	<u>0.372</u>	<u>0.397</u>	0.381	0.404	0.378	0.403	0.414	0.427	0.559	0.515	0.383	0.407	0.395	0.414	0.389	0.414
ECL	0.170	0.266	0.174	0.269	0.195	0.288	0.216	0.306	0.192	0.295	0.212	0.300	0.178	0.270	0.207	0.289	0.217	0.308
Weather	0.244	0.275	<u>0.250</u>	<u>0.276</u>	0.275	0.291	0.253	<u>0.276</u>	0.259	0.287	0.265	0.317	0.258	0.278	0.257	0.280	0.279	0.297
ILI	1.705	0.835	<u>1.922</u>	<u>0.921</u>	2.432	1.012	2.108	0.929	2.139	0.931	2.616	1.090	2.444	1.203	2.388	1.011	2.623	1.060
Exchange	0.353	0.401	0.395	0.429	0.372	0.416	0.364	0.404	0.416	0.443	<u>0.354</u>	0.414	0.360	<u>0.403</u>	0.390	0.429	0.519	0.500
1st Count	7	7	<u>1</u>	<u>2</u>	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Table 1: Multivariate forecasting results. All results are averaged from four different forecasting horizons: $\mathbf{H} \in \{96, 192, 336, 720\}$ for the input sequence length 96. **Bold**: the best, underline: the second best. Full results are in Appendix C.

Dataset	Ours		TimeCMA		TimeLLM		GPT4TS		TimesNet		DLinear		PatchTST		Fedformer	
	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTm1	0.376	0.398	<u>0.387</u>	<u>0.410</u>	0.404	0.427	0.464	0.441	0.677	0.537	0.411	0.429	0.501	0.466	0.722	0.605
ETTm2	0.266	<u>0.327</u>	0.312	0.358	<u>0.277</u>	0.323	0.293	0.335	0.320	0.353	0.316	0.368	0.296	0.343	0.463	0.488
ETTh1	0.449	0.454	<u>0.480</u>	<u>0.479</u>	0.556	0.522	0.590	0.525	0.869	0.628	0.691	0.600	0.633	0.542	0.639	0.561
ETTh2	0.357	0.388	0.398	0.433	<u>0.370</u>	<u>0.394</u>	0.397	0.421	0.479	0.465	0.605	0.538	0.415	0.431	0.466	0.475
Weather	0.226	0.268	<u>0.229</u>	<u>0.272</u>	0.234	0.273	0.238	0.275	0.279	0.301	0.241	0.283	0.242	0.279	0.284	0.324
1st Count	5	4	0	0	0	<u>1</u>	0	0	0	0	0	0	0	0	0	0

Table 2: Few-shot learning on 10% training data. All results are averaged from four different forecasting horizons: $\mathbf{H} \in \{96, 192, 336, 720\}$ for the input sequence length 512. **Bold**: the best, underline: the second best. Full results are in Appendix D.

Dataset	Ours		TimeCMA		TimeLLM		GPT4TS		TimesNet		DLinear		PatchTST		Fedformer	
	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTm1	0.384	0.405	<u>0.396</u>	<u>0.416</u>	0.425	0.434	0.472	0.450	0.717	0.561	0.400	0.417	0.526	0.476	0.730	0.592
ETTm2	0.267	0.330	0.329	0.367	<u>0.274</u>	<u>0.323</u>	0.308	0.346	0.344	0.372	0.399	0.426	0.314	0.352	0.381	0.404
ETTh1	0.442	0.451	<u>0.472</u>	<u>0.470</u>	0.627	0.543	0.681	0.560	0.925	0.647	0.750	0.611	0.694	0.569	0.658	0.562
ETTh2	0.357	0.403	0.395	0.430	<u>0.382</u>	<u>0.418</u>	0.400	0.433	0.439	0.448	0.694	0.577	0.827	0.615	0.463	0.454
Weather	0.226	0.269	<u>0.231</u>	<u>0.273</u>	0.260	0.309	0.263	0.301	0.298	0.318	0.263	0.308	0.269	0.303	0.309	0.353
1st Count	5	5	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Table 3: Few-shot learning on 5% training data. All results are averaged from four different forecasting horizons: $\mathbf{H} \in \{96, 192, 336, 720\}$ for the input sequence length 512. **Bold**: the best, underline: the second best. Full results are in Appendix D.

Few-Shot Forecasting

Setups. Over the years the foundation models have shown exceptional performance in generalization tasks by few-shot learning or zero-shot learning settings (Brown et al. 2020; Achiam et al. 2023). To evaluate the generalization performance of T3Time under few-shot conditions, we follow the setups from Jin et al. (2023) for a fair comparison. Specifically, we adopt a scenario where the training data consists of 10% and 5% of the available time steps, with input sequence length set to 512. We evaluate T3Time on the same benchmarks used in the long-term forecasting experiments.

Results. Tables 2 and 3 summarize the results for 10% and 5% few-shot learning, respectively. In both settings, T3Time outperforms almost all SOTA baselines. Specifically, for the recent SOTA models such as TimeCMA,

TimeLLM, and GPT4TS (OFA) in the 10% few-shot task, average MSE is reduced by 7.13%, 7.42%, and 13.44%, respectively. T3Time reduces average MSE and MAE by 3.62% and 1.98% for 10% few-shot forecasting tasks in all SOTA scores. Similar trend is visible in the 5% few-shot forecasting task, as T3Time improves upon SOTA scores by 4.13% and 1.91% w.r.t. MSE and MAE. When compared with models like TimesNet, DLinear, and PatchTST, T3Time consistently outperforms these approaches, surpassing 30% in MSE reduction. Additionally, relative to recent SOTA methods like TimeCMA, TimeLLM, and GPT4TS, T3Time demonstrates an average MSE improvement exceeding 10%. A comprehensive overview of our full results for few shot forecasting is provided in Appendix D.

Design	ETTm1		ETTm2		ETTh1		ETTh2		Weather		ILI		Exchange Rate		1st Count
	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	
Our Model	0.372	0.393	0.279	0.322	0.418	0.430	0.348	0.390	0.244	0.275	1.705	0.835	0.353	0.401	14
w/o Frequency Module	0.381	0.394	0.279	0.324	0.433	0.433	0.364	0.398	0.250	0.280	1.786	0.884	0.374	0.410	1
w/o Multihead CMA	0.374	0.395	0.283	0.325	0.421	0.432	0.355	0.392	0.249	0.277	1.813	0.865	0.378	0.413	0
w/o Residual Connection	0.404	0.413	0.288	0.329	0.433	0.443	0.384	0.411	0.249	0.280	2.176	0.969	0.396	0.427	0
w/o Gating Mechanism	0.373	0.396	0.280	0.322	0.425	0.432	0.363	0.398	0.249	0.277	1.724	0.837	0.373	0.412	1

Table 4: Results of design choices related to the model. All results are averaged from four different forecasting horizons: $H \in \{96, 192, 336, 720\}$ for the input sequence length 96. **Bold**: the best. Full results are in Appendix E.

Design Variant

Setups. To evaluate the robustness and contribution of various design choices in T3Time, we experiment with several ablations, including removing frequency features, multi-head mechanisms, residual connections, and gating. Specifically, we evaluate the following design variants: (1) *Without Frequency*, where frequency-domain features are excluded; (2) *Without Multi-Head*, where the adaptive multi-head cross-modal alignment is disabled; (3) *Without Residual*, which eliminates the channel-wise residual fusion mechanism; and (4) *Without Gating*, where the horizon-aware gating mechanism is omitted. We maintain the same setup as the long-term forecasting.

Results. Table 4 presents the performance comparison across different design variants. The full model achieves average SOTA performance in 14 out of 14 cases for both MSE and MAE. Notably, removing frequency-domain features results in a 3.22% MSE and 1.85% MAE drop compared to the full model. This emphasizes the crucial role of frequency-domain information in capturing both temporal and spectral patterns. However, the residual connection’s contribution is the most significant as excluding it leads to the largest performance drop, with an average MSE increase of 8.36% and MAE of 5.25%. Excluding the gating mechanism and multi-head CMA results in a more modest decline, with an average MSE decrease of approximately 2%. Full results for the design variants are provided in Appendix E.

t-SNE Visualization

To better understand the effectiveness of the learned embeddings, we employ T-SNE visualization. Figure 3 presents the T-SNE visualization of the embeddings across four key modalities: time series embeddings, frequency embeddings, prompt embeddings, and forecasted embeddings. Panel 3(a) illustrates the time series embeddings, which demonstrate clear clustering of the datasets. Separation between the datasets indicates that the model is effectively capturing their temporal characteristics. Panel 3(b) visualizes the frequency embeddings, where we observe a distinct clustering pattern. This suggests that the model captures meaningful frequency-domain information, which is crucial for distinguishing underlying periodicities in the time series. Panel 3(c) shows the prompt embeddings, which exhibit a relatively more dispersed structure compared to the time series and frequency embeddings, reflecting the diverse nature of the prompts. Finally, panel 3(d) displays the forecasted em-

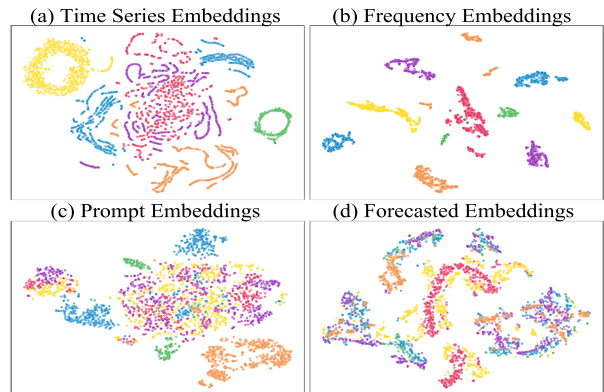


Figure 3: The combined T-SNE visualization of time series, frequency, prompt, and forecasted embeddings across six datasets. A detailed breakdown of each modality and dataset is provided in Appendix F.

beddings, where the clustering is similar to that of the time series embeddings, suggesting that the model is learning representations that align with both historical and forecasted data. These visualizations illustrate T3Time’s ability to integrate temporal, spectral, and prompt information in learning robust multimodal representation.

Conclusion

In this work, we present **T3Time**, a novel framework for multivariate time series forecasting which integrates three distinct modalities—time series, frequency, and prompt embeddings. Along with the tri-modal encoding framework, T3Time adapts the fusion of temporal-spectral features using a dynamic horizon-aware gating mechanism. The introduction of adaptive multi-head cross-modal alignment enables more flexible and contextual interaction between the time-spectral and prompt representations, while the channel-wise residual connection ensures efficient fusion of these representations later in the framework. T3Time consistently outperforms SOTA models, achieving notable reductions in both MSE and MAE across numerous baselines. Additionally, T3Time demonstrates strong generalization in few-shot learning scenarios, with significant performance gains. Future works can explore large-scale pertaining and better representation methods to enrich modalities for stronger time series forecasting.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Cao, D.; Wang, Y.; Duan, J.; Zhang, C.; Zhu, X.; Huang, C.; Tong, Y.; Xu, B.; Bai, J.; Tong, J.; et al. 2020. Spectral temporal graph neural network for multivariate time-series forecasting. *Advances in neural information processing systems*, 33: 17766–17778.
- Challu, C.; Olivares, K. G.; Oreshkin, B. N.; Ramirez, F. G.; Canseco, M. M.; and Dubrawski, A. 2023. Nhits: Neural hierarchical interpolation for time series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 6989–6997.
- Chang, C.; Chan, C.-T.; Wang, W.-Y.; Peng, W.-C.; and Chen, T.-F. 2024. Timedrl: Disentangled representation learning for multivariate time-series. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, 625–638. IEEE.
- Chen, M.; Xu, Z.; Zeng, A.; and Xu, Q. 2023. FrAug: Frequency domain augmentation for time series forecasting. *arXiv preprint arXiv:2302.09292*.
- Chen, P.-Y. 2024. Model reprogramming: Resource-efficient cross-domain machine learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 22584–22591.
- Deldari, S.; Xue, H.; Saeed, A.; He, J.; Smith, D. V.; and Salim, F. D. 2022. Beyond just vision: A review on self-supervised representation learning on multimodal and temporal data. *arXiv preprint arXiv:2206.02353*.
- Franceschi, J.-Y.; Dieuleveut, A.; and Jaggi, M. 2019. Un-supervised scalable representation learning for multivariate time series. *Advances in neural information processing systems*, 32.
- Huang, Q.; Zhou, Z.; Yang, K.; Lin, G.; Yi, Z.; and Wang, Y. 2024. Leret: Language-empowered retentive network for time series forecasting. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*.
- Jia, F.; Wang, K.; Zheng, Y.; Cao, D.; and Liu, Y. 2024. Gpt4mts: Prompt-based large language model for multimodal time-series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 23343–23351.
- Jin, G.; Liu, C.; Xi, Z.; Sha, H.; Liu, Y.; and Huang, J. 2022. Adaptive dual-view wavenet for urban spatial-temporal event prediction. *Information Sciences*, 588: 315–330.
- Jin, M.; Wang, S.; Ma, L.; Chu, Z.; Zhang, J. Y.; Shi, X.; Chen, P.-Y.; Liang, Y.; Li, Y.-F.; Pan, S.; et al. 2023. Time-llm: Time series forecasting by reprogramming large language models. *arXiv preprint arXiv:2310.01728*.
- Lai, G.; Chang, W.-C.; Yang, Y.; and Liu, H. 2018. Modeling long- and short-term temporal patterns with deep neural networks. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, 95–104.
- Liu, C.; Xiao, Z.; Wang, D.; Wang, L.; Jiang, H.; Chen, H.; and Yu, J. 2022a. Exploiting Spatiotemporal Correlations of Arrive-Stay-Leave Behaviors for Private Car Flow Prediction. *IEEE Transactions on Network Science and Engineering*, 9(2): 834–847.
- Liu, C.; Xu, Q.; Miao, H.; Yang, S.; Zhang, L.; Long, C.; Li, Z.; and Zhao, R. 2025. Timecma: Towards llm-empowered multivariate time series forecasting via cross-modality alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 18780–18788.
- Liu, C.; Yang, S.; Xu, Q.; Li, Z.; Long, C.; Li, Z.; and Zhao, R. 2024a. Spatial-temporal large language model for traffic prediction. In *2024 25th IEEE International Conference on Mobile Data Management (MDM)*, 31–40. IEEE.
- Liu, H.; Ma, Z.; Yang, L.; Zhou, T.; Xia, R.; Wang, Y.; Wen, Q.; and Sun, L. 2023a. Sadi: A self-adaptive decomposed interpretable framework for electric load forecasting under extreme events. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Liu, S.; Yu, H.; Liao, C.; Li, J.; Lin, W.; Liu, A. X.; and Dustdar, S. 2022b. Pyraformer: Low-Complexity Pyramidal Attention for Long-Range Time Series Modeling and Forecasting. In *International Conference on Learning Representations*.
- Liu, X.; Hu, J.; Li, Y.; Diao, S.; Liang, Y.; Hooi, B.; and Zimmermann, R. 2024b. Unitime: A language-empowered unified model for cross-domain time series forecasting. In *Proceedings of the ACM Web Conference 2024*, 4095–4106.
- Liu, Y.; Hu, T.; Zhang, H.; Wu, H.; Wang, S.; Ma, L.; and Long, M. 2023b. iTransformer: Inverted Transformers Are Effective for Time Series Forecasting. *arXiv preprint arXiv:2310.06625*.
- Liu, Z.; Miao, H.; Zhao, Y.; Liu, C.; Zheng, K.; and Li, H. 2024c. LightTR: A lightweight framework for federated trajectory recovery. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, 4422–4434. IEEE.
- Lu, K.; Grover, A.; Abbeel, P.; and Mordatch, I. 2021. Pre-trained Transformers as Universal Computation Engines. *arXiv preprint arXiv:2103.05247*.
- Miao, H.; Zhao, Y.; Guo, C.; Yang, B.; Zheng, K.; Huang, F.; Xie, J.; and Jensen, C. S. 2024. A unified replay-based continuous learning framework for spatio-temporal prediction on streaming data. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, 1050–1062. IEEE.
- Nie, Y.; H. Nguyen, N.; Sinthong, P.; and Kalagnanam, J. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *International Conference on Learning Representations*.

- Pan, Z.; Jiang, Y.; Garg, S.; Schneider, A.; Nevmyvaka, Y.; and Song, D. 2024. s^2 IP-LLM: Semantic space informed prompt learning with LLM for time series forecasting. In *Forty-first International Conference on Machine Learning*.
- Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Schneider, S. H.; and Dickinson, R. E. 1974. Climate modeling. *Reviews of Geophysics*, 12(3): 447–493.
- Sun, F.-K.; and Boning, D. S. 2022. Fredo: frequency domain-based long-term time series forecasting. *arXiv preprint arXiv:2205.12301*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wen, Q.; Zhou, T.; Zhang, C.; Chen, W.; Ma, Z.; Yan, J.; and Sun, L. 2022. Transformers in time series: A survey. *arXiv preprint arXiv:2202.07125*.
- Woo, G.; Liu, C.; Sahoo, D.; Kumar, A.; and Hoi, S. 2022. Cost: Contrastive learning of disentangled seasonal-trend representations for time series forecasting. *arXiv preprint arXiv:2202.01575*.
- Wu, H.; Hu, T.; Liu, Y.; Zhou, H.; Wang, J.; and Long, M. 2022. Timesnet: Temporal 2d-variation modeling for general time series analysis. *arXiv preprint arXiv:2210.02186*.
- Wu, H.; Xu, J.; Wang, J.; and Long, M. 2021. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems*, 34: 22419–22430.
- Xu, J.; Wu, H.; Wang, J.; and Long, M. 2021. Anomaly transformer: Time series anomaly detection with association discrepancy. *arXiv preprint arXiv:2110.02642*.
- Xue, H.; and Salim, F. D. 2023. Promptcast: A new prompt-based learning paradigm for time series forecasting. *IEEE Transactions on Knowledge and Data Engineering*, 36(11): 6851–6864.
- Xue, H.; Voutharoja, B. P.; and Salim, F. D. 2022. Leveraging language foundation models for human mobility forecasting. In *Proceedings of the 30th international conference on advances in geographic information systems*, 1–9.
- Yi, K.; Zhang, Q.; Fan, W.; Wang, S.; Wang, P.; He, H.; An, N.; Lian, D.; Cao, L.; and Niu, Z. 2023. Frequency-domain MLPs are more effective learners in time series forecasting. *Advances in Neural Information Processing Systems*, 36: 76656–76679.
- Yin, S.; Fu, C.; Zhao, S.; Li, K.; Sun, X.; Xu, T.; and Chen, E. 2024. A survey on multimodal large language models. *National Science Review*, 11(12): nwae403.
- Yue, W.; Liu, Y.; Ying, X.; Xing, B.; Guo, R.; and Shi, J. 2025. Freeformer: Frequency enhanced transformer for multivariate time series forecasting. *arXiv preprint arXiv:2501.13989*.
- Zeng, A.; Chen, M.; Zhang, L.; and Xu, Q. 2023. Are transformers effective for time series forecasting? In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 11121–11128.
- Zhang, K.; Wen, Q.; Zhang, C.; Cai, R.; Jin, M.; Liu, Y.; Zhang, J. Y.; Liang, Y.; Pang, G.; Song, D.; et al. 2024. Self-supervised learning for time series analysis: Taxonomy, progress, and prospects. *IEEE transactions on pattern analysis and machine intelligence*, 46(10): 6775–6794.
- Zhang, X.; Zhao, Z.; Tsiligkaridis, T.; and Zitnik, M. 2022. Self-supervised contrastive pre-training for time series via time-frequency consistency. *Advances in neural information processing systems*, 35: 3988–4003.
- Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; Xiong, H.; and Zhang, W. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 11106–11115.
- Zhou, T.; Ma, Z.; Wen, Q.; Wang, X.; Sun, L.; and Jin, R. 2022. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International conference on machine learning*, 27268–27286. PMLR.
- Zhou, T.; Niu, P.; Sun, L.; Jin, R.; et al. 2023. One fits all: Power general time series analysis by pretrained Im. *Advances in neural information processing systems*, 36: 43322–43355.