

Adaptive Multi-Scale Decomposition Framework for Time Series Forecasting

Yifan Hu^{1,2,*}, Peiyuan Liu^{1,*}, Peng Zhu^{2,†}, Dawei Cheng^{2,†}, Tao Dai^{3,†}

¹Tsinghua Shenzhen International Graduate School

²Tongji University

³Shenzhen University

{huyf0122, peiyuanliu.edu, daitao.edu}@gmail.com

{pengzhu, dcheng}@tongji.edu.cn

Abstract

Transformer-based and MLP-based methods have emerged as leading approaches in time series forecasting (TSF). However, real-world time series often show different patterns at different scales, and future changes are shaped by the interplay of these overlapping scales, requiring high-capacity models. While Transformer-based methods excel in capturing long-range dependencies, they suffer from high computational complexities and tend to overfit. Conversely, MLP-based methods offer computational efficiency and adeptness in modeling temporal dynamics, but they struggle with capturing temporal patterns with complex scales effectively. Based on the observation of multi-scale entanglement effect in time series, we propose a novel MLP-based Adaptive Multi-Scale Decomposition (AMD) framework for TSF. Our framework decomposes time series into distinct temporal patterns at multiple scales, leveraging the Multi-Scale Decomposable Mixing (MDM) block to dissect and aggregate these patterns. Complemented by the Dual Dependency Interaction (DDI) block and the Adaptive Multi-predictor Synthesis (AMS) block, our approach effectively models both temporal and channel dependencies and utilizes autocorrelation to refine multi-scale data integration. Comprehensive experiments demonstrate our AMD framework not only overcomes the limitations of existing methods but also consistently achieves state-of-the-art performance across various datasets.

Code — <https://github.com/TROUBADOUR000/AMD>

Introduction

Time series forecasting (TSF) aims to use historical data to predict future values across various domains, such as finance (Cheng et al. 2022, 2018; Zhu et al. 2024), energy (Xue and Salim 2023), traffic management (Kim et al. 2022), and weather forecasting (Volkovs, Urtans, and Caune 2024). Recently, deep learning has made substantial and reliable advancements in TSF, with the most state-of-the-art performances achieved by Transformer-based methods (Nie et al. 2023; Zhou et al. 2022b; Zhang and Yan 2023) and MLP-based methods (Zeng et al. 2023; Tolstikhin et al. 2021; Li et al. 2023b; Qiu et al. 2024).

*These authors contributed equally.

†Corresponding author: Dawei Cheng and Tao Dai.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

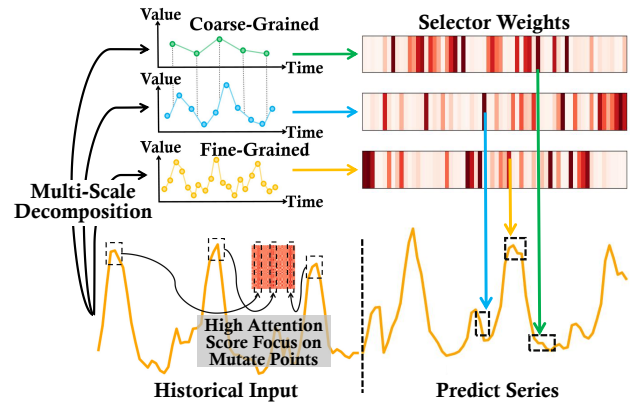


Figure 1: Illustration of multi-scale temporal patterns in time series and the impact of selector weights. Transformer-based methods often overfit by overemphasizing mutation points, weakening temporal relationships. Efficiently modeling and integrating distinct temporal patterns at various scales is crucial for accurate predictions.

It is worth noting that time series exhibit distinctly different temporal patterns at various sampling scales (Wang et al. 2024). Moreover, the weight of these time scales in predicting future variations is not uniform, as future variations are jointly determined by the entangle of multiple scales (see Fig. 1). For example, weather data sampled hourly reflects fine-grained, sudden changes, while monthly sampled data captures coarse-grained climate variations. Similarly, while short-term hourly data might highlight immediate weather shifts, long-term monthly data provides a broader view of climatic trends. Therefore, efficiently modeling the multi-scale changes in time series and adeptly integrating information across different scales, which requires high-capacity models, remains a critical challenge.

Although Transformer-based methods excel at modeling long-range dependencies due to the self-attention mechanisms (Vaswani et al. 2017), they come with computational complexity that scales quadratically with the length of the sequence. Despite their recognized performance, these methods still face challenges in industrial applications, including low training efficiency and high memory consumption. Additionally, self-attention can diminish the tempo-

ral relationships when extracting semantic correlations between pairs in long sequences (Zeng et al. 2023), leading to an overemphasis on mutation points and resulting in overfitting (see Fig. 1). In contrast, MLP-based methods boast significantly lower computational complexity compared to Transformer-based methods. Moreover, MLP-based methods can chronologically model the temporal dynamics in consecutive points, which is crucial for time series analysis (Zeng et al. 2023; Das et al. 2023). However, the simplicity of linear mappings in existing MLP-based methods presents an information bottleneck (Liu et al. 2024b), hindering their ability to capture diverse temporal patterns and limiting their predictive accuracy (Ni et al. 2024).

Motivated by the above observations, we decompose the time series at multiple scales to precisely discern the intertwined temporal patterns within the complex series, rather than merely breaking it down into seasonal and trend components (Wu et al. 2021; Zeng et al. 2023). Subsequently, we model the correlations across different scales in both time and channel dimensions. Unlike average aggregation of TimeMixer (Wang et al. 2024), to account for the varying impacts of different temporal patterns on the future, we employ an autocorrelation approach to model their contributions and adaptively integrate these multi-scale temporal patterns based on their respective influences.

Technically, we propose an MLP-based Adaptive Multi-Scale Decomposition (AMD) Framework to better disentangle and model the diverse temporal patterns within time series. In concrete, the AMD initiates by employing the Multi-Scale Decomposable Mixing (MDM) block, which first decomposes the original time series into multiple temporal patterns through average downsampling and then aggregates these scales to provide aggregate information in a residual way. Subsequently, the Dual Dependency Interaction (DDI) block simultaneously models both temporal and channel dependencies within the aggregated information. Finally, the Adaptive Multi-predictor Synthesis (AMS) block uses the aggregated information to generate specific weights and then employs these weights to adaptively integrate the multiple temporal patterns produced by the DDI. Through comprehensive experimentation, our AMD consistently achieves state-of-the-art performance in both long-term and short-term forecasting tasks, with superior efficiency across multiple datasets.

Our contributions are summarized as follows:

- We decompose time series across multiple scales to precisely identify dominant intertwined temporal patterns within complex sequences and adaptively aggregate predictions of temporal patterns at different scales, addressing their varied impacts on future forecasts. We also demonstrate the feasibility through theoretical analysis.
- We propose a simple but effective MLP-based Adaptive Multi-Scale Decomposition (AMD) framework that initially decomposes time series into diverse temporal patterns, models both temporal and channel dependencies of these patterns, and finally synthesizes the outputs using a weighted aggregation approach to focus on the changes of dominant temporal patterns, thereby enhancing pre-

diction accuracy across scales.

- Comprehensive experiments demonstrate that our AMD consistently delivers state-of-the-art performance in both long-term and short-term forecasting across various datasets, with superior efficiency.

Related Works

Time Series Forecasting

Recently, deep learning methods for TSF have gained prominence, such as CNN (Wang et al. 2023; Dai et al. 2024), RNN (Lin et al. 2023; Jia et al. 2023), GNN (Wu et al. 2020; Yi et al. 2023), Transformer (Wu et al. 2021; Liu et al. 2024a) and MLP (Das et al. 2023; Zhang et al. 2022). Transformer-based models, renowned for superior performance in handling long sequential data, have gained popularity in TSF. PatchTST (Nie et al. 2023) divides the time series into patches to enhance locality. Beyond cross-time dependencies, Crossformer (Zhang and Yan 2023) also mines cross-variable dependencies. However, Transformer-based models always suffer from efficiency problems due to high computational complexity. In contrast, MLP-based models have a smaller memory footprint. FITS (Xu, Zeng, and Xu 2024) proposes a new linear mapping for the transformation of complex inputs, with only 10k parameters. However, due to the inherent simplicity and information bottleneck, MLP-based models struggle to effectively capture diverse temporal patterns (Ni et al. 2024). In this work, we decompose time series across multiple scales and use separate linear models for each, effectively addressing such representational limits.

Mixture of Experts for TSF

The concept of Mixture of Experts (MoE) has a long history (Jacobs et al. 1991; Jordan 1994). Recent advances in large language models (LLM) have reignited interest in MoE (Kim, Lim, and Han 2024; Xue et al. 2024). However, in the TSF domain, the utilization of MoE remains limited. FiLM (Zhou et al. 2022a) assigns input sequences with different time horizons to various experts for forecasting. FEDformer (Zhou et al. 2022b) extracts a set of data-dependent weights by MoE for combining multiple trend components as the final trend. In our work, we harness the adaptability of MoE to craft distinct predictors customized for individual temporal patterns, thereby optimizing model performance and enhancing overall model interpretability.

Series Decomposition in TSF

Lately, with high sampling rates leading to high-frequency data (such as daily, hourly, or minutely data), real-world time series data often contains multiple underlying temporal patterns. To competently harness different temporal patterns at various scales, several series decomposition designs are proposed (Liu et al. 2022; Oreshkin et al. 2020). Seasonal-Trend decomposition (Wu et al. 2021) uses moving averages to separate seasonal and trend components. TimesNet (Wu et al. 2023) uses Fast Fourier Transform to extract multiple dominant frequencies, while SCINet (LIU et al. 2022) employs a hierarchical downsampling tree to iteratively decompose multi-scale information. TimeMixer (Wang et al. 2024)

downsamples time series into different scales of seasonal and trend sequences. However, they ignore high-order interactions across different scales. Refining these designs, we introduces an adaptive multi-scale decomposition method to model the dynamic interactions and accurately discern dominated temporal patterns.

Preliminary: Linear Models with Multi-Scale Information for TSF

We consider the following problem: given a collection of time series samples with historical observations $\mathbf{X} \in \mathbb{R}^{C \times L}$, where C denotes the number of variables and L represents the length of the look-back sequence. The objective is to predict $\mathbf{Y} \in \mathbb{R}^{M \times T}$, where M is the number of target variables to be predicted ($M \leq C$) and T is the length of the future time steps to be predicted. A linear model learns parameters $\mathbf{A} \in \mathbb{R}^{L \times T}$ and $\mathbf{b} \in \mathbb{R}^T$ to predict the values of next T time steps as:

$$\hat{\mathbf{Y}} = \mathbf{X}\mathbf{A} \oplus \mathbf{b} \in \mathbb{R}^{C \times L} \quad (1)$$

where \oplus means column-wise addition. The corresponding M rows in $\hat{\mathbf{Y}}$ can be used to predict \mathbf{Y} .

After that, we introduce the multi-scale information. For time series forecasting, the most influential real-world applications typically exhibit either smooth or periodicity. Without these characteristics, predictability tends to be low, rendering predictive models unreliable. If the time series only exhibits periodicity, linear models can easily model it (Chen et al. 2023). We define the original sequence as $f(x) = f_0(x) = [x_1, x_2, \dots, x_L]$ and assume that $f(x)$ possesses smoothness. After k downsampling operations with a downsampling rate of d , we obtain n sequences $f_i(x), \forall i = 1, 2, \dots, n$, where $f_i(x) = \frac{1}{d} \sum_{j=xd-d+1}^{xd} f_{i-1}(j)$ and $x = 1, 2, \dots, \lfloor \frac{L}{d^i} \rfloor$. It is noteworthy that $f_i(x) \in \mathbb{R}^{C \times \lfloor \frac{L}{d^i} \rfloor}, \forall i = 0, 1, \dots, n$. Then, the sequence $f_i(x), \forall i = 0, 1, \dots, n$ is transformed into $g_i(x)$ through linear mapping and residual calculation. Specifically, $g_n(x) = f_n(x)$, then through top-down recursion for $i = n-1, \dots, 0$, the operation

$$g_i(x) = f_i(x) + g_{i+1}(x)W_i \quad (2)$$

is performed recursively, where $W_i \in \mathbb{R}^{\lfloor \frac{L}{d^{i+1}} \rfloor \times \lfloor \frac{L}{d^i} \rfloor}$. In this case, we derive the **Theorem 1**

Theorem 1. *Let multi-scale mixing representation $g(x)$, where $g(x) \in \mathbb{R}^{1 \times L}$ (for simplicity, we consider univariate sequences) and the original sequence $f(x)$ is Lipschitz smooth with constant \mathcal{K} (i.e. $\left| \frac{f(a)-f(b)}{a-b} \right| \leq \mathcal{K}$), then there exists a linear model such that $|y_t - \hat{y}_t|$ is bounded, $\forall t = 1, \dots, T$.*

This derivation demonstrates that linear models are well-suited to utilize multi-scale information. For nonperiodic patterns provided they exhibit smoothness, which is often observed in practice, the error remains bounded.

Method

As mentioned above, the key challenge for accurate forecasting is to deal with multi-scale temporal patterns that have an

entangled influence on the future evolution of variables. In this paper, we propose AMD to enhance forecasting accuracy by capturing multi-scale entanglement through multi-scale decomposition and an autocorrelation approach (see Fig. 2). Specifically, AMD mainly consists of three components: Multi-Scale Decomposable Mixing (MDM) Block, Dual Dependency Interaction (DDI) Block, and Adaptive Multi-predictor Synthesis (AMS) Block. The details of each essential module are explained in the following subsections.

Multi-Scale Decomposable Mixing Block

Time series exhibit both coarse-grained temporal patterns and fine-grained patterns. Together, these complementary scales of information provide a comprehensive view of the time series. Therefore, we first decompose the time series into individual temporal patterns, and then mix them to enhance the time series data for a more nuanced analysis and interpretation.

Specifically, the raw input information \mathbf{X} already contains fine-grained details, while coarse-grained information is extracted through average pooling. First-level temporal pattern τ_1 is the input of one channel \mathbf{x} . Next, distinct coarse-grained temporal patterns $\tau_i \in \mathbb{R}^{1 \times \lfloor \frac{L}{d^{i-1}} \rfloor} (\forall i = 2, \dots, h)$ are extracted by applying average pooling over the previous layer of temporal patterns, where h denotes the number of downsampling operations and d denotes the downsampling rate. The decomposition of the i^{th} layer of temporal patterns can be represented as:

$$\tau_i = \text{AvgPooling}(\tau_{i-1}) \quad (3)$$

Then, distinct temporal patterns are mixed from the coarse-grained τ_h to the fine-grained τ_1 through a feedforward residual network, while ξ_i represents the mixed data and $\xi_h = \tau_h$. The mixing of the i^{th} layer of temporal patterns can be represented by the following formula:

$$\xi_i = \tau_i + \text{MLP}(\xi_{i+1}) \quad (4)$$

Finally, after completing the mixing of temporal patterns across h scales, we obtain mixed-scale information ξ_1 , with the output of one channel being $\mathbf{u} = \xi_1 \in \mathbb{R}^{1 \times L}$.

Dual Dependency Interaction Block

We observed potential dynamic interactions across different scales in various time series. Intuitively, in a stock price series, daily fluctuations are influenced by monthly economic trends, which in turn are shaped by annual market cycles. As a result, changes at each scale propagate through the system, with the effects at one scale spilling over and altering the dynamics at other scales. However, ScaleFormer (Shabani et al. 2023) and TimeMixer (Wang et al. 2024) ignore these potential high-order interactions among different scales.

To model the interactions across different scales, including both temporal and channel dependencies, we propose the DDI block with time-mixing and channel-mixing. DDI first stacks the aggregated information \mathbf{u} from various channels of the MDM into the matrix $\mathbf{U} \in \mathbb{R}^{C \times L}$ and then performs patch operations to transform \mathbf{U} into $\hat{\mathbf{U}} \in \mathbb{R}^{C \times N \times P}$. For each patch, $\hat{\mathbf{V}}_t^{t+P}$ represents the embedding output of the

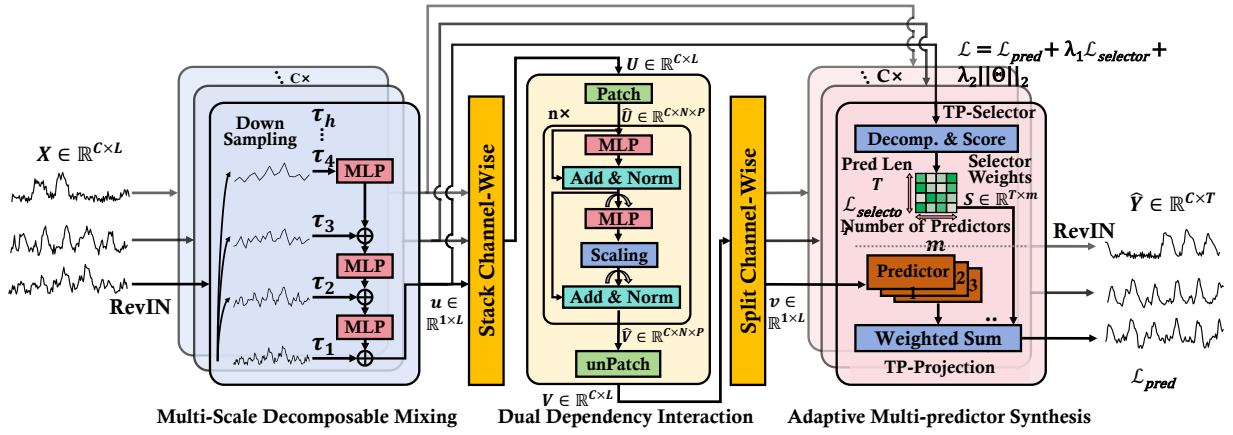


Figure 2: Overall structure of the proposed AMD. MDM decomposes and mixes multi-scale information. DDI models dynamic interactions across different scales. AMS discerns dominated temporal patterns and adaptively makes predictions.

residual network, and $\hat{\mathbf{U}}_t^{t+P}$ represents a patch of aggregated information from MDM. We adopt an MLP shared by time steps to aggregate time-mixing information through the time dimension for each channel to obtain temporal dependencies \mathbf{Z}_t^{t+P} . Next, we perform the transpose operation to fuse channel-mixing information along the time domain through another MLP shared by channels. Finally, we perform the unpatch operation and split the output information into individual channels to obtain $\mathbf{v} \in \mathbb{R}^{1 \times L}$. The residual operation ensures that the model retains its ability to capture temporal dependencies while effectively exploiting cross-channel dependencies. The interaction of the patch $\hat{\mathbf{U}}_t^{t+P}$ can be represented by the following formula:

$$\mathbf{Z}_t^{t+P} = \hat{\mathbf{U}}_t^{t+P} + \text{MLP}(\hat{\mathbf{V}}_t^{t-P}) \quad (5)$$

$$\hat{\mathbf{V}}_t^{t+P} = \mathbf{Z}_t^{t+P} + \beta \cdot \text{MLP}((\mathbf{Z}_t^{t+P})^T)^T \quad (6)$$

where A^T is the transpose of matrix A .

In DDI, dual dependencies are captured under the mixed-scale information \mathbf{U} . Temporal dependencies model the interactions across different periods, while cross-channel dependencies model the relationships between different variables. However, through experiments, we find that cross-channel dependencies are not always effective, especially when the target time series is not correlated with other covariates; instead, they often introduce unwanted interference. Therefore, we introduce the scaling rate β to suppress the noise and balance the emphasis on temporal dependencies and cross-channel dependencies.

Adaptive Multi-predictor Synthesis Block

It is noted that the dominant temporal patterns, which significantly influence future variables, dynamically change over different periods and have stronger predictive power. In contrast to the average aggregation in TimeMixer (Wang et al. 2024), which overlooks these dominant changes, we exploit the adaptive properties of MoE to design specific predictors for each temporal pattern. By dynamically assigning more attention to the dominant scales, we improve both accuracy and generalisability.

The AMS is partitioned into two components: the temporal pattern selector (TP-Selector) and the temporal pattern projection (TP-Projection). The TP-Selector decomposes different temporal patterns and generates the selector weights \mathbf{S} . Unlike the downsampling in MDM, which decomposes individual scales to enhance temporal information, the TP-Selector adaptively untangles highly correlated, intertwined mixed scales through feedforward processing. Meanwhile, the TP-Projection synthesizes the multi-predictions and adaptively aggregates the outputs based on the specific weights.

TP-Selector takes a single channel input \mathbf{u} from MDM, decomposes it through feedforward layers, and then applies a noisy gating design (Shazeer et al. 2017):

$$\mathbf{S} = \text{Softmax}(\text{TopK}(\text{Softmax}(Q(\mathbf{u})), k)) \quad (7)$$

$$Q(\mathbf{u}) = \text{Decomp.}(\mathbf{u}) + \psi \cdot \text{Softplus}(\text{Decomp.}(\mathbf{u}) \cdot \mathbf{W}_{\text{noise}}) \quad (8)$$

where k is the number of dominant temporal patterns, $\psi \in \mathcal{N}(0, 1)$ is standard Gaussian noise, and $\mathbf{W}_{\text{noise}} \in \mathbb{R}^{m \times m}$ is a learnable weight controlling the noisy values.

TP-Projection takes the embedding \mathbf{v} from the output \mathbf{V} of DDI as input. Each predictor consists of multiple feedforward layers. These outputs are then multiplied by the \mathbf{S} and summed to yield the prediction result $\hat{\mathbf{y}}$ for one channel. The final results $\hat{\mathbf{Y}}$ are composed of the outputs $\hat{\mathbf{y}}$ from each channel:

$$\hat{\mathbf{y}} = \sum_{j=0}^m \mathbf{S}_j \cdot \text{Predictor}_j(\mathbf{v}) \quad (9)$$

Compared to sparse MoE (Li et al. 2023a), we adopt dense MoE in TP-Projection with two considerations. Firstly, each temporal pattern contributes to the prediction result. Secondly, this approach can help mitigate issues such as load unbalancing and embedding omissions that are prevalent in sparse MoE architectures. Our TopK method can be formalized with:

$$\text{TopK}(\mathbf{u}, k) = \begin{cases} \alpha \cdot \log(\mathbf{u} + 1), & \text{if } \mathbf{u} < v_k \\ \alpha \cdot \exp(\mathbf{u}) - 1, & \text{if } \mathbf{u} \geq v_k \end{cases} \quad (10)$$

Models	AMD (Ours)		TimeMixer		PatchTST		iTransformer		Crossformer		FEDformer		TimesNet		MICN		DLinear		MTS-Mixers		
Metric	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	
Weather	96	0.145	0.197	0.147	0.197	0.149	0.198	0.174	0.214	0.153	0.217	0.238	0.314	0.172	0.220	0.161	0.226	0.152	0.237	0.156	0.206
	192	0.187	0.238	0.189	0.239	0.194	0.241	0.221	0.254	0.197	0.269	0.275	0.329	0.219	0.261	0.220	0.283	0.220	0.282	0.199	0.248
	336	0.240	0.280	0.241	0.280	0.245	0.282	0.278	0.296	0.252	0.311	0.339	0.377	0.280	0.306	0.275	0.328	0.265	0.319	0.249	0.291
	720	0.315	0.330	0.310	0.330	0.314	0.334	0.358	0.349	0.318	0.363	0.389	0.409	0.365	0.359	0.311	0.356	0.323	0.362	0.336	0.343
ETTh1	96	0.369	0.397	0.361	0.390	0.370	0.399	0.386	0.405	0.386	0.429	0.376	0.415	0.384	0.402	0.396	0.427	0.375	0.399	0.372	0.395
	192	0.401	0.416	0.409	0.414	0.413	0.421	0.441	0.436	0.416	0.444	0.423	0.446	0.457	0.436	0.430	0.453	0.405	0.416	0.416	0.426
	336	0.418	0.427	0.430	0.429	0.422	0.436	0.487	0.458	0.440	0.461	0.444	0.462	0.491	0.469	0.433	0.458	0.439	0.443	0.455	0.449
	720	0.439	0.454	0.445	0.460	0.447	0.466	0.503	0.491	0.519	0.524	0.469	0.492	0.521	0.500	0.474	0.508	0.472	0.490	0.475	0.472
ETTh2	96	0.274	0.337	0.271	0.330	0.274	0.337	0.297	0.349	0.628	0.563	0.332	0.374	0.340	0.374	0.289	0.357	0.289	0.353	0.307	0.354
	192	0.351	0.383	0.317	0.402	0.339	0.379	0.380	0.400	0.703	0.624	0.407	0.446	0.402	0.414	0.409	0.438	0.383	0.418	0.374	0.399
	336	0.375	0.411	0.332	0.396	0.329	0.380	0.428	0.432	0.827	0.675	0.400	0.447	0.452	0.452	0.417	0.452	0.448	0.465	0.398	0.432
	720	0.402	0.438	0.342	0.408	0.379	0.422	0.427	0.445	1.181	0.840	0.412	0.469	0.462	0.468	0.426	0.473	0.605	0.561	0.463	0.465
ETTM1	96	0.284	0.339	0.291	0.340	0.290	0.342	0.334	0.368	0.316	0.373	0.326	0.390	0.338	0.375	0.314	0.360	0.299	0.343	0.314	0.358
	192	0.322	0.362	0.327	0.365	0.332	0.369	0.377	0.391	0.377	0.411	0.365	0.415	0.371	0.387	0.359	0.387	0.335	0.365	0.354	0.386
	336	0.360	0.380	0.360	0.381	0.366	0.392	0.426	0.420	0.431	0.442	0.391	0.425	0.410	0.411	0.398	0.413	0.369	0.386	0.384	0.405
	720	0.421	0.416	0.415	0.417	0.416	0.420	0.491	0.459	0.600	0.547	0.446	0.458	0.478	0.450	0.459	0.464	0.425	0.421	0.427	0.432
ETTM2	96	0.167	0.258	0.164	0.254	0.166	0.256	0.180	0.264	0.421	0.461	0.180	0.271	0.187	0.267	0.178	0.273	0.167	0.260	0.177	0.259
	192	0.221	0.294	0.223	0.295	0.223	0.296	0.250	0.309	0.503	0.519	0.252	0.318	0.249	0.309	0.245	0.316	0.224	0.303	0.241	0.303
	336	0.270	0.327	0.279	0.330	0.274	0.329	0.311	0.348	0.611	0.580	0.324	0.364	0.321	0.351	0.295	0.350	0.281	0.342	0.297	0.338
	720	0.356	0.382	0.359	0.383	0.362	0.385	0.412	0.407	0.996	0.750	0.410	0.420	0.497	0.403	0.389	0.409	0.397	0.421	0.396	0.398
ECL	96	0.129	0.224	0.129	0.224	0.129	0.222	0.148	0.240	0.187	0.283	0.186	0.302	0.168	0.272	0.159	0.267	0.153	0.237	0.141	0.243
	192	0.147	0.238	0.140	0.220	0.147	0.240	0.162	0.253	0.258	0.330	0.197	0.311	0.184	0.289	0.168	0.279	0.152	0.249	0.163	0.261
	336	0.160	0.253	0.161	0.255	0.163	0.259	0.178	0.269	0.323	0.369	0.213	0.328	0.198	0.300	0.196	0.308	0.169	0.267	0.176	0.277
	720	0.193	0.286	0.194	0.287	0.197	0.290	0.225	0.317	0.404	0.423	0.233	0.344	0.220	0.320	0.203	0.312	0.233	0.344	0.212	0.308
Exchange	96	0.083	0.201	0.086	0.204	0.093	0.214	0.086	0.206	0.186	0.346	0.136	0.276	0.107	0.234	0.102	0.235	0.081	0.203	0.083	0.201
	192	0.171	0.293	0.176	0.298	0.192	0.312	0.177	0.299	0.467	0.522	0.256	0.369	0.226	0.344	0.172	0.316	0.157	0.293	0.174	0.296
	336	0.309	0.402	0.345	0.426	0.350	0.432	0.331	0.417	0.783	0.721	0.426	0.464	0.367	0.448	0.272	0.407	0.305	0.414	0.336	0.417
	720	0.750	0.652	0.848	0.692	0.911	0.716	0.847	0.691	1.367	0.943	1.090	0.800	0.964	0.746	0.714	0.658	0.643	0.601	0.900	0.715
Traffic	96	0.366	0.259	0.360	0.249	0.360	0.249	0.395	0.268	0.512	0.290	0.576	0.359	0.593	0.321	0.508	0.301	0.410	0.282	0.462	0.332
	192	0.381	0.265	0.379	0.256	0.375	0.350	0.417	0.276	0.523	0.297	0.610	0.380	0.617	0.336	0.536	0.315	0.423	0.287	0.488	0.354
	336	0.397	0.269	0.385	0.270	0.392	0.264	0.433	0.283	0.530	0.300	0.608	0.375	0.629	0.336	0.525	0.310	0.436	0.296	0.498	0.360
	720	0.429	0.292	0.430	0.281	0.432	0.286	0.467	0.302	0.573	0.313	0.621	0.375	0.640	0.350	0.571	0.323	0.466	0.315	0.529	0.370
Solar	96	0.175	0.228	0.167	0.220	0.224	0.278	0.203	0.237	0.181	0.240	0.201	0.304	0.219	0.314	0.188	0.252	0.289	0.337	0.284	0.325
	192	0.186	0.235	0.187	0.249	0.253	0.298	0.233	0.261	0.196	0.252	0.237	0.337	0.231	0.322	0.215	0.280	0.319	0.397	0.307	0.362
	336	0.200	0.246	0.200	0.258	0.273	0.306	0.248	0.273	0.216	0.243	0.254	0.362	0.246	0.337	0.222	0.267	0.352	0.415	0.333	0.384
	720	0.203	0.249	0.215	0.250	0.272	0.308	0.249	0.275	0.220	0.256	0.280	0.397	0.280	0.363	0.226	0.264	0.356	0.412	0.335	0.383

Table 1: Long-term forecasting task. All the results are averaged from 4 different prediction lengths $T \in \{96, 192, 336, 720\}$. To the best for a fairer comparison for all baselines, the input sequence length L is searched among $\{96, 192, 336, 512, 672, 720\}$.

where v_k is the k -th largest value among \mathbf{u} and α is a constant used to adjust the selector weights. The scaling operation within our TopK requires the input values to be restricted to the interval $[0, 1]$. Consequently, we perform an additional Softmax operation according to Eq. (7).

Loss Function

The loss function of AMS consists of two components. (1) For the predictors, the Mean Squared Error (MSE) loss ($\mathcal{L}_{\text{pred}} = \sum_{i=0}^T \|\mathbf{y}_i - \hat{\mathbf{y}}_i\|_2^2$) is used to measure the variance between predicted values and ground truth. (2) For the gating network loss, we apply the coefficient of variation loss function ($\mathcal{L}_{\text{selector}} = \frac{\text{Var}(S)}{\text{Mean}(S)^2 + \epsilon}$, where ϵ is a small positive constant to prevent numerical instability), which optimizes the gating mechanism by promoting a balanced assignment of experts to inputs, thereby enhancing the overall performance of the MoE. The total loss function is defined as:

$$\mathcal{L} = \mathcal{L}_{\text{pred}} + \lambda_1 \mathcal{L}_{\text{selector}} + \lambda_2 \|\Theta\|_2 \quad (11)$$

where $\|\Theta\|_2$ is the L2-norm, $\lambda_{1,2}$ are hyper-parameters.

Experiments

Main Results

Datasets. We conduct experiments on seven real-world datasets, including Weather, ETTh1, ETTh2, ETTm1,

ETTM2, ECL, Exchange, Traffic and Solar Energy for long-term forecasting and PEMS (PEMS03, PEMS04, PEMS07, PEMS08) for short-term forecasting.

Baselines. We carefully select some representative models to serve as baselines in the field of time series forecasting, including (1) MLP-based models: TimeMixer (Wang et al. 2024), TiDE (Das et al. 2023), MTS-Mixers (Li et al. 2023b), RLinear (Li et al. 2024), and DLinear (Zeng et al. 2023); (2) Transformer-based models: PatchTST (Nie et al. 2023), iTransformer (Liu et al. 2024a), Crossformer (Zhang and Yan 2023), and FEDformer (Zhou et al. 2022b); (3) CNN-based models: TimesNet (Wu et al. 2023), and MICN (Wang et al. 2023).

Experimental Settings. To ensure fair comparisons, for long-term forecasting, we rerun all baselines with different input lengths L and choose the best results to avoid underestimating the baselines. For short-term forecasting, the input length is 96. We select two common metrics in time series forecasting: Mean Absolute Error (MAE) and Mean Squared Error (MSE). All experiments are conducted using PyTorch on an NVIDIA V100 32GB GPU and are repeated five times for consistency.

Results. Comprehensive forecasting results are shown in Tab. 1 and Tab. 2, which present long-term and short-term

Models	AMD		PatchTST		Crossformer		FEDformer		TimesNet		TiDE		DLinear		MTS-Mixers		RLinear	
	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
PEMS03	0.084	<u>0.198</u>	0.099	0.216	0.090	0.203	0.126	0.251	<u>0.085</u>	0.192	0.178	0.305	0.122	0.243	0.117	0.232	0.126	0.236
PEMS04	0.083	<u>0.198</u>	0.105	0.224	0.098	0.218	0.138	0.262	<u>0.087</u>	0.195	0.219	0.340	0.148	0.272	0.129	0.267	0.138	0.252
PEMS07	0.074	0.180	0.095	0.150	0.094	0.200	0.109	0.225	<u>0.082</u>	<u>0.181</u>	0.173	0.304	0.115	0.242	0.134	0.278	0.118	0.235
PEMS08	0.093	0.206	0.168	0.232	0.165	0.214	0.173	0.273	<u>0.112</u>	<u>0.212</u>	0.227	0.343	0.154	0.276	0.186	0.286	0.133	0.247

Table 2: Short-term forecasting task. The input sequence length L is 96 and the prediction length T is 12.

Models	AMD (Ours)		AMD (Sparse)		RandomOrder		AverageWeight		w/o DDI		$\beta = 0.5$		$\beta = 1.0$		w/o $\mathcal{L}_{selector}$		w/o MDM		w/o AMS		
	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	
Weather	96	0.145	0.198	0.150	0.205	0.152	0.209	0.149	0.205	0.154	0.208	0.146	0.198	0.147	0.200	0.160	0.211	0.149	0.204	0.150	0.206
	192	0.187	0.238	0.193	0.245	0.194	0.245	0.194	0.246	0.197	0.248	0.194	0.245	0.195	0.246	0.209	0.258	0.192	0.243	0.199	0.248
	336	0.240	0.280	0.244	0.284	0.245	0.286	0.249	0.288	0.249	0.289	0.249	0.287	0.248	0.286	0.278	0.309	0.246	0.285	0.248	0.287
	720	0.315	0.330	0.315	0.334	0.321	0.338	0.325	0.340	0.323	0.340	0.318	0.335	0.320	0.336	0.364	0.366	0.319	0.336	0.320	0.339
ECL	96	0.129	0.224	0.132	0.228	0.133	0.230	0.135	0.234	0.132	0.231	0.140	0.242	0.147	0.248	0.139	0.238	0.135	0.231	0.137	0.234
	192	0.147	0.238	0.151	0.244	0.155	0.248	0.155	0.251	0.151	0.247	0.169	0.268	0.175	0.273	0.156	0.456	0.153	0.247	0.157	0.251
	336	0.160	0.253	0.165	0.259	0.171	0.264	0.169	0.263	0.167	0.266	0.174	0.274	0.183	0.285	0.177	0.271	0.169	0.265	0.171	0.267
	720	0.193	0.286	0.199	0.293	0.202	0.294	0.202	0.295	0.200	0.294	0.207	0.298	0.215	0.308	0.208	0.299	0.201	0.296	0.203	0.298

Table 3: Component ablation of AMD on Weather and ECL.

forecasting results respectively. The best results are highlighted in **red** and the second-best are underlined. The lower the MSE/MAE, the more accurate the forecast. AMD stands out with the best performance in 50 cases and the second best in 27 cases out of the overall 80 cases. Compared with other baselines, AMD performs well on both high-dimensional and low-dimensional datasets. It is worth noting that PatchTST does not perform well on the PEMS datasets for the patching design leads to the neglect of highly fluctuating temporal patterns. In contrast, AMD leverages information from multi-scale temporal patterns. Furthermore, the performance of Crossformer is unsatisfactory for it introduces unnecessary noise by exploring cross-variable dependencies. AMD skillfully reveals the intricate dependencies existing among time steps across various variables. Additionally, DLinear and MTS-mixers perform poorly on high-dimensional datasets, whereas AMD can handle them.

Model Analysis

Ablation Study. To verify the effectiveness of each component of AMD, we perform detailed ablation of each possible design on Weather and ECL datasets. As shown in Tab. 3, we have the following observations.

For MDM, we conduct experiments by removing MDM (**w/o MDM**). The increased error suggests that relying on a single scale is insufficient for accurate prediction.

For DDI, recent studies (Nie et al. 2023; Zhang and Yan 2023) show that both channel independence and cross-channel dependencies strategies can achieve SOTA accuracy in specific tasks. Typically, datasets with strong cross-channel correlations perform better using cross-channel dependencies strategies. However, for datasets with weak inter-channel correlations, cross-channel dependencies often introduce unwanted noise. To prove this, we introduce cross-channel dependencies by adjusting the scaling rate. Specifically, we set the parameter β in Eq. (7) to 0.5 and 1.0, respectively. In addition, we conduct experiments by removing DDI (**w/o DDI**). From the results, it can be seen




Figure 3: AMD guides the prediction by assigning greater weight to the dominant time pattern.

that the introduction of cross-channel dependencies do not always enhance the prediction accuracy, and this was consistent across other datasets as well. Scaling rate β can balance the emphasis on temporal dependencies and cross-channel dependencies in multivariate time series data and boost the model performance.

For AMS, we demonstrate the performance improvement is not solely due to the enlarged model size, but rather the integration of temporal pattern information, we devised the following experiments: To demonstrate the utilization of sequential information, we replaced the temporal pattern embedding with random orders as **RandomOrder**. To demonstrate the effectiveness of AMS aggregation, we replaced the TP-selector with the average weighting as **AverageWeight**, treating different temporal patterns equally. We also remove AMS (**w/o AMS**) and replace it with a single linear predictor. **RandomOrder**, **AverageWeight** and **w/o AMS** all cause a performance decline. This illustrates that, compared to self-attention whose permutation-invariant nature results in the loss of sequential information, AMS inherently maintains the sequential order and makes better use of temporal relationships by identifying dominant patterns that change over time. Compared to simple averaging like timemixer (Wang et al. 2024), AMS adaptively assigns cor-




Figure 4: Cross-channel dependencies may lead to deviations from the original distribution.

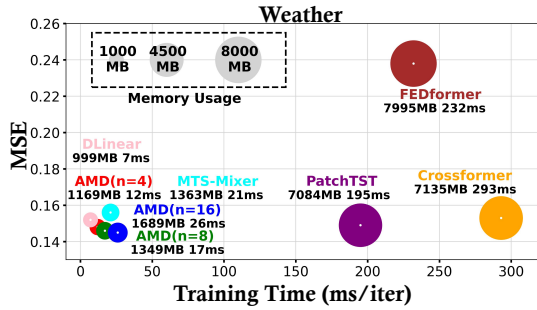


Figure 5: Memory usage (MB), training time (ms/iter), and MSE comparisons on the Weather dataset. The input and predicted lengths are set to 512 and 96, respectively.

responding weights to different temporal patterns, resulting in more accurate predictions.

Model Interpretability. To provide an intuitive understanding of AMS, we plotted selector weights in Fig. 3. Compared to a single linear predictor, AMS offers stronger generalization capabilities and improved interpretability. Before and after the time step T , the temporal variations are respectively dominated by TP16 and TP5. Before T , the predicted data resembles the trend of TP16, both exhibiting a downward fluctuation. However, after T , the predicted data resembles TP5, which suddenly shows a significant increase. AMD recognizes this changing dominant role over time and adaptively assigns them higher weights.

To offer a clear conceptualization of the balance of emphasis on temporal dependencies and cross-channel dependencies mentioned above, we visualize the learned dependencies as shown in Fig. 4. Compared to the temporal dependencies, especially when the target variable is not correlated with other covariates, cross-channel dependencies tend to smooth out the variability in the target variable, causing its distribution to deviate from what would be expected based solely on its own past values.

Dense MoE Strategy and Sparse MoE Strategy. In theory, the information contained in each temporal pattern is useful and discarding any one of them would result in loss of information. Therefore, we adopt the dense MoE strategy, where dominant temporal patterns are given larger weights, while others are given smaller weights instead of being set to 0. We conduct ablation experiments as shown in Tab. 3. Compared to Dense MoE, Sparse MoE shows increased

Models	DLinear		+ MDM & AMS		MTS-Mixers		+ MDM & AMS		
	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	
Weather	96	0.152	0.237	0.146	0.212	0.156	0.206	0.155	0.203
	192	0.220	0.282	0.194	0.261	0.199	0.248	0.202	0.251
	336	0.265	0.319	0.245	0.305	0.249	0.291	0.244	0.286
	720	0.323	0.362	0.313	0.356	0.336	0.343	0.326	0.337
<i>Imp.</i>	-	-	6.46%	5.50%	-	-	1.38%	1.01%	
ECL	96	0.153	0.237	0.150	0.244	0.141	0.243	0.137	0.239
	192	0.152	0.249	0.159	0.256	0.163	0.261	0.160	0.258
	336	0.169	0.267	0.167	0.265	0.176	0.277	0.170	0.271
	720	0.233	0.344	0.221	0.313	0.212	0.308	0.203	0.303
<i>Imp.</i>	-	-	1.41%	1.73%	-	-	3.18%	1.65%	

Table 4: Comparative impact of MDM & AMS on different baselines. *Imp.* represents the average percentage improvement of MDM & AMS compared to the original methods.

prediction errors. This observation highlights the consequence of omitting non-dominant temporal pattern information, which invariably leads to a degradation in performance.

Efficiency Analysis. We thoroughly compare the training time and memory usage of various baselines on the Weather dataset, using the official model configurations and the same batch size. The results, shown in Fig. 5, indicate that the efficiency of AMD outperforms other Transformer-based models and MLP-based models with a relatively small number of parameters. Furthermore, as the number of predictors increases, the prediction error (MSE) gradually decreases, but the training time and memory usage increase significantly. To strike a balance, we set the number of predictors to 8 across all experiments.

Analysis on the balance of the TP-Selector. We conduct experiments on the scaling rate of the load balancing loss, denoted by λ_1 in Eq. (11). As shown in Tab. 3, utilizing $\mathcal{L}_{\text{selector}}$ results in significantly improved performance, exceeding 11.2% in terms of MSE compared to when $\lambda_1 = 0.0$. This underscores the crucial role of implementing load-balancing losses. Furthermore, the selector weights in the Fig. 3 do not tend to favor specific temporal patterns, addressing the load imbalance issue in sparse MoE structures.

Analysis on AMD as a plugin. Finally, we explore whether our proposed MoE-based method can yield an improvement in the performance of other TSF methods. We selected DLinear (Zeng et al. 2023) and MTS-Mixers (Li et al. 2023b) as the baselines. After integrating the MDM and AMS modules, the predictive capabilities of all these models are enhanced as shown in Tab. 4, while maintaining the same computational resource requirements.

Conclusion

In this paper, we propose the Adaptive Multi-Scale Decomposition (AMD) framework for time series forecasting to address the inherent complexity of time series data by decomposing it into multiple temporal patterns at various scales and adaptively aggregating these patterns. Comprehensive experiments demonstrate that AMD consistently achieves state-of-the-art performance in both long-term and short-term forecasting tasks across various datasets, showcasing superior efficiency and effectiveness.

Acknowledgements

This work is supported by the National Key R&D Program of China (Grant no. 2022YFB4501704), the National Science Foundation of China (Grant no. 62472317), the Fundamental Research Funds for the Central Universities, the Shanghai Science and Technology Innovation Action Plan Project (Grant no. 22YS1400600 and 24692118300). This work is also supported in part by the National Natural Science Foundation of China, under Grant (62302309), Shenzhen Science and Technology Program (JCYJ20220818101014030).

References

- Chen, S.-A.; Li, C.-L.; Arik, S. O.; Yoder, N. C.; and Pfister, T. 2023. TSMixer: An All-MLP Architecture for Time Series Forecasting. *Transactions on Machine Learning Research*.
- Cheng, D.; Liu, Y.; Niu, Z.; and Zhang, L. 2018. Modeling Similarities Among Multi-Dimensional Financial Time Series. *IEEE Access*, 6: 43404–43413.
- Cheng, D.; Yang, F.; Xiang, S.; and Liu, J. 2022. Financial time series forecasting with multi-modality graph neural network. *Pattern Recognition*, 121: 108218.
- Dai, T.; Wu, B.; Liu, P.; Li, N.; Bao, J.; Jiang, Y.; and Xia, S.-T. 2024. Periodicity Decoupling Framework for Long-term Series Forecasting. In *International Conference on Learning Representations*.
- Das, A.; Kong, W.; Leach, A.; Mathur, S. K.; Sen, R.; and Yu, R. 2023. Long-term Forecasting with TiDE: Time-series Dense Encoder. *Transactions on Machine Learning Research*.
- Jacobs, R. A.; Jordan, M. I.; Nowlan, S. J.; and Hinton, G. E. 1991. Adaptive Mixtures of Local Experts. *Neural Computation*, 3(1): 79–87.
- Jia, Y.; Lin, Y.; Hao, X.; Lin, Y.; Guo, S.; and Wan, H. 2023. WITRAN: Water-wave Information Transmission and Recurrent Acceleration Network for Long-range Time Series Forecasting. In *Advances in Neural Information Processing Systems*.
- Jordan, M. 1994. Hierarchical mixtures of experts and the EM algorithm. In *IEE Colloquium on Advances in Neural Networks for Control and Systems*, 1/1–1/3.
- Kim, D.; Cho, Y.; Kim, D.; Park, C.; and Choo, J. 2022. Residual Correction in Real-Time Traffic Forecasting. In *Proceedings of the ACM International Conference on Information & Knowledge Management*, 962–971. New York, NY, USA: Association for Computing Machinery. ISBN 9781450392365.
- Kim, Y.; Lim, H.; and Han, D. 2024. Scaling Beyond the GPU Memory Limit for Large Mixture-of-Experts Model Training. In *International Conference on Machine Learning*.
- Li, B.; Shen, Y.; Yang, J.; Wang, Y.; Ren, J.; Che, T.; Zhang, J.; and Liu, Z. 2023a. Sparse Mixture-of-Experts are Domain Generalizable Learners. In *International Conference on Learning Representations*.
- Li, Z.; Qi, S.; Li, Y.; and Xu, Z. 2024. Revisiting Long-term Time Series Forecasting: An Investigation on Affine Mapping.
- Li, Z.; Rao, Z.; Pan, L.; and Xu, Z. 2023b. MTS-Mixers: Multivariate Time Series Forecasting via Factorized Temporal and Channel Mixing. *arXiv preprint arXiv:2302.04501*.
- Lin, S.; Lin, W.; Wu, W.; Zhao, F.; Mo, R.; and Zhang, H. 2023. SegRNN: Segment Recurrent Neural Network for Long-Term Time Series Forecasting. *arXiv preprint arXiv:2308.11200*.
- LIU, M.; Zeng, A.; Chen, M.; Xu, Z.; LAI, Q.; Ma, L.; and Xu, Q. 2022. SCINet: Time Series Modeling and Forecasting with Sample Convolution and Interaction. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 5816–5828. Curran Associates, Inc.
- Liu, S.; Yu, H.; Liao, C.; Li, J.; Lin, W.; Liu, A. X.; and Dustdar, S. 2022. Pyraformer: Low-Complexity Pyramidal Attention for Long-Range Time Series Modeling and Forecasting. In *International Conference on Learning Representations*.
- Liu, Y.; Hu, T.; Zhang, H.; Wu, H.; Wang, S.; Ma, L.; and Long, M. 2024a. iTransformer: Inverted Transformers Are Effective for Time Series Forecasting. In *International Conference on Learning Representations*.
- Liu, Z.; Wang, Y.; Vaidya, S.; Ruehle, F.; Halverson, J.; Soljačić, M.; Hou, T. Y.; and Tegmark, M. 2024b. KAN: Kolmogorov-Arnold Networks. *arXiv preprint arXiv:2404.19756*.
- Ni, R.; Lin, Z.; Wang, S.; and Fanti, G. 2024. Mixture-of-Linear-Experts for Long-term Time Series Forecasting. *arXiv preprint arXiv:2312.06786*.
- Nie, Y.; Nguyen, N. H.; Sinthong, P.; and Kalagnanam, J. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *International Conference on Learning Representations*.
- Oreshkin, B. N.; Carpov, D.; Chapados, N.; and Bengio, Y. 2020. N-BEATS: Neural basis expansion analysis for interpretable time series forecasting. In *International Conference on Learning Representations*.
- Qiu, X.; Wu, X.; Lin, Y.; Guo, C.; Hu, J.; and Yang, B. 2024. DUET: Dual Clustering Enhanced Multivariate Time Series Forecasting. *arXiv preprint arXiv:2412.10859*.
- Shabani, M. A.; Abdi, A. H.; Meng, L.; and Sylvain, T. 2023. Scaleformer: Iterative Multi-scale Refining Transformers for Time Series Forecasting. In *International Conference on Learning Representations*.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *International Conference on Learning Representations*.
- Tolstikhin, I. O.; Houlsby, N.; Kolesnikov, A.; Beyer, L.; and Zhai. 2021. MLP-Mixer: An all-MLP Architecture for Vision. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 24261–24272. Curran Associates, Inc.

- 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.
- Volkovs, K.; Urtans, E.; and Caune, V. 2024. Primed UNet-LSTM for Weather Forecasting. In *Proceedings of the International Conference on Advances in Artificial Intelligence, ICAAI '23*, 13–17. New York, NY, USA: Association for Computing Machinery. ISBN 9798400708985.
- Wang, H.; Peng, J.; Huang, F.; Wang, J.; Chen, J.; and Xiao, Y. 2023. MICN: Multi-scale Local and Global Context Modeling for Long-term Series Forecasting. In *International Conference on Learning Representations*.
- Wang, S.; Wu, H.; Shi, X.; Hu, T.; Luo, H.; Ma, L.; Zhang, J. Y.; and ZHOU, J. 2024. TimeMixer: Decomposable Multiscale Mixing for Time Series Forecasting. In *International Conference on Learning Representations*.
- 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 *International Conference on Learning Representations*.
- Wu, H.; Xu, J.; Wang, J.; and Long, M. 2021. Autoformer: Decomposition Transformers with Auto-Correlation for Long-Term Series Forecasting. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 22419–22430. Curran Associates, Inc.
- Wu, Z.; Pan, S.; Long, G.; Jiang, J.; Chang, X.; and Zhang, C. 2020. Connecting the Dots: Multivariate Time Series Forecasting with Graph Neural Networks. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '20*, 753–763. New York, NY, USA: Association for Computing Machinery. ISBN 9781450379984.
- Xu, Z.; Zeng, A.; and Xu, Q. 2024. FITS: Modeling Time Series with 10^k Parameters. In *International Conference on Learning Representations*.
- Xue, F.; Zheng, Z.; Fu, Y.; Ni, J.; Zheng, Z.; Zhou, W.; and You, Y. 2024. OpenMoE: An Early Effort on Open Mixture-of-Experts Language Models. In *International Conference on Machine Learning*.
- Xue, H.; and Salim, F. D. 2023. Utilizing Language Models for Energy Load Forecasting. In *Proceedings of the ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys '23*, 224–227. New York, NY, USA: Association for Computing Machinery. ISBN 9798400702303.
- Yi, K.; Zhang, Q.; Fan, W.; He, H.; Hu, L.; Wang, P.; An, N.; Cao, L.; and Niu, Z. 2023. FourierGNN: Rethinking Multivariate Time Series Forecasting from a Pure Graph Perspective. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 69638–69660. Curran Associates, Inc.
- Zeng, A.; Chen, M.; Zhang, L.; and Xu, Q. 2023. Are Transformers Effective for Time Series Forecasting? *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(9): 11121–11128.
- Zhang, T.; Zhang, Y.; Cao, W.; Bian, J.; Yi, X.; Zheng, S.; and Li, J. 2022. Less Is More: Fast Multivariate Time Series Forecasting with Light Sampling-oriented MLP Structures. *arXiv preprint arXiv:2207.01186*.
- Zhang, Y.; and Yan, J. 2023. Crossformer: Transformer Utilizing Cross-Dimension Dependency for Multivariate Time Series Forecasting. In *International Conference on Learning Representations*.
- Zhou, T.; MA, Z.; Wang, X.; Wen, Q.; Sun, L.; Yao, T.; Yin, W.; and Jin, R. 2022a. FiLM: Frequency improved Legendre Memory Model for Long-term Time Series Forecasting. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 12677–12690. Curran Associates, Inc.
- Zhou, T.; Ma, Z.; Wen, Q.; Wang, X.; Sun, L.; and Jin, R. 2022b. FEDformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International Conference on Machine Learning*, 27268–27286. PMLR.
- Zhu, P.; Li, Y.; Hu, Y.; Liu, Q.; Cheng, D.; and Liang, Y. 2024. LSR-IGRU: Stock Trend Prediction Based on Long Short-Term Relationships and Improved GRU. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, CIKM '24*, 5135–5142. New York, NY, USA: Association for Computing Machinery. ISBN 9798400704369.