

# Times2D: Multi-Period Decomposition and Derivative Mapping for General Time Series Forecasting

Reza Nematirad, Anil Pahwa, and Balasubramaniam Natarajan

Electrical and Computer Engineering Department, Kansas State University, Manhattan, KS, 66506, USA  
nematirad@ksu.edu

## Abstract

Time series forecasting is an important application in various domains such as energy management, traffic planning, financial markets, meteorology, and medicine. However, real-time series data often present intricate temporal variability and sharp fluctuations, which pose significant challenges for time series forecasting. Previous models that rely on 1D time series representations usually struggle with complex temporal variations. To address the limitations of 1D time series, this study introduces the Times2D method that transforms the 1D time series into 2D space. Times2D consists of three main parts: first, a Periodic Decomposition Block (PDB) that captures temporal variations within a period and between the same periods by converting the time series into a 2D tensor in the frequency domain. Second, the First and Second Derivative Heatmaps (FSDH) capture sharp changes and turning points, respectively. Finally, an Aggregation Forecasting Block (AFB) integrates the output tensors from PDB and FSDH for accurate forecasting. This 2D transformation enables the utilization of 2D convolutional operations to effectively capture long and short characteristics of the time series. Comprehensive experimental results across large-scale data in the literature demonstrate that the proposed Times2D model achieves state-of-the-art performance in both short-term and long-term forecasting.

**Code** — <https://github.com/Tims2D/Times2D>

## Introduction

Time series forecasting plays a critical role in several real-world applications, including electricity load demand forecasting, lifetime estimation of industrial machinery, traffic planning, weather prediction, and the stock market. Due to their critical relevance and wide application, there has been considerable interest in time series forecasting in recent years. The range of data resolutions in time series has historically been broad, from annual records such as GDP rates to monthly indicators like unemployment rates, and daily measurements such as stock market indices. However, advancements in data collection technologies have significantly increased the frequency of data recording. For example, observations range from hourly metrics like electricity use, to second-by-second updates in internet traffic

and even millisecond-level data in rotary machines. Low-resolution time series data were typically smoother and exhibited fewer fluctuations, making prediction relatively straightforward. However, shifting toward high-resolution time steps introduces intricate temporal patterns and sharp fluctuations, making the time series forecasting task significantly challenging. Traditional methods such as Autoregressive Integrated Moving Average (ARIMA) and Exponential Smoothing were effective in handling well-defined seasonal patterns and trends. However, they struggle with intricate temporal and nonlinear patterns in time series data (Fu, Wu, and Boulet 2022). Recently, due to advances in deep learning models, the landscape of time series forecasting has achieved promising progress. These models, by leveraging non-linear modeling, high-dimensional pattern recognition, and adaptive learning capabilities, have been able to capture complex temporal variations in real-world time series (He et al. 2022). Recurrent Neural Network (RNN) models are designed to process sequences by handling time points successively with internal state updates that capture temporal information. However, RNNs often struggle with long-term dependencies due to vanishing gradients (Massart and Abrol 2022). Long Short-Term Memory (LSTM) models were proposed to address the vanishing gradient problem. However, LSTMs have difficulty capturing local trends (Huang et al. 2024). Temporal Convolutional Networks (TCNs) are utilized to extract the variation information by applying convolutional neural networks. TCNs are efficient at capturing relationships between adjacent time points but may struggle to model long-term dependencies (Wu et al. 2022). Recently, transformer-based models have shown superior performance over TCN models due to their attention mechanism. However, the attention mechanism can struggle with capturing dependencies among scattered time points (Wu et al. 2021) and can be computationally extensive, especially for long-term forecasting (Dai et al. 2024).

It should be noted that intricate temporal variations in real-world data often manifest as sharp fluctuations, abrupt rises and falls, and overlapping patterns. Additionally, most real-world datasets exhibit multi-periodicity, which further complicates the modeling process. These periods may overlap, creating additional challenges when attempting to model them using a 1D perspective. Motivated by the limitations of the existing deep learning models and the complex

characteristics observed in real-world time series data, this study introduces a novel framework for general time series forecasting. The proposed approach specifically addresses the challenges of multi-periodicity, sharp fluctuations, and turning points in time series data. Technically, Times2D consists of three major components. First, the Periodic Decomposition Block (PDB) identifies the top  $k$  dominant periods and their corresponding frequencies using the Fast Fourier Transform (FFT). For each dominant period-frequency pair, the 1D time series is reshaped into a 2D tensor, where the rows represent the period and the columns correspond to the frequency. As a result, the 1D input time series is transformed into  $k$  distinct 2D tensors. This transformation facilitates the model ability to capture both short-term and long-term dependencies. Second, the First and Second Derivative Heatmaps (FSDH) operate separately from the PDB. This component calculates the first and second derivatives of the original 1D time series, constructing 2D tensors that highlight sharp changes and turning points in the data. Finally, the Aggregation Forecasting Block (AFB) integrates the outputs generated by the PDB and FSDH to produce accurate forecasts, effectively capturing the intricate temporal variations inherent in the time series. Figure 1 provides an overview of the proposed Times2D framework. The contributions of this study are summarized as follows:

- This study introduces the Times2D framework, which transforms 1D time series data into 2D space to capture complex temporal patterns, including multi-periodicity, sharp fluctuations, and turning points for general time series forecasting.
- We propose the PDB, which decomposes the time series into segments based on dominant periods identified through FFT. This block effectively captures both short-term and long-term dependencies.
- The study introduces the FSDH block, which calculates the first and second derivatives to emphasize sharp changes and turning points in the time series.
- Comprehensive experiments on various large-scale datasets demonstrate that the proposed Times2D model achieves state-of-the-art performance in both short-term and long-term forecasting.

### Related Work

Classical methods such as ARIMA and Exponential Smoothing aim to extract predefined temporal variations, such as trends and seasonality. While these methods are effective for capturing well-defined seasonal patterns and trends, they often struggle with more complex, non-linear patterns that are present in real-world time series data (Zhang, Meng, and Gu 2024). In recent years, deep learning has made significant contributions to time series forecasting. Based on the literature review in this field, deep learning models to capture temporal variations in time series data can be categorized into four groups.

### Recurrent Neural Networks

RNNs and their variants, LSTMs and Gated Recurrent Units (GRUs), have been widely adopted for modeling sequen-

tial data. RNNs process sequences by maintaining a hidden state that captures temporal dependencies (Tan et al. 2023). However, they often suffer from vanishing gradients, which makes it difficult for RNNs to capture long-term relationships (Liu et al. 2022). LSTM networks overcome this issue by introducing memory cells and gate mechanisms to control the flow of information. This modification enables LSTMs to capture long-term dependencies. An example of an LSTM-based model is proposed in (Li, Xu, and Anastasiu 2023) to predict 3-day ahead hourly water level. However, LSTMs still face challenges in capturing local trends and rapid fluctuations.

### Convolutional Neural Networks

TCN models in time series analysis take advantage of causal convolutions, dilated convolutions, and residual connections. Causal convolutions ensure that outputs at time step  $t$  depend only on inputs from time step  $t$  and earlier. This aspect of TCNs prevents future data leakage, which is essential for accurate time series forecasting. Dilated convolutions expand the receptive field exponentially without increasing computational complexity. This enhances the ability of the model to capture long-term dependencies. Residual connections mitigate the vanishing gradient problem. They enable the training of deeper networks by providing direct paths for gradient flow (Wang et al. 2023). A TCN is developed in (Ahmadi et al. 2014) to predict the weather, and experimental results show that the TCN outperforms the LSTM and other classical machine learning methods in predicting weather. A Multi-scale Iso-metric Convolution Network (MICN) is proposed in (Wang et al. 2023) to efficiently capture the local and global patterns in time series data for long-term forecasting. Temporal 2D-Variation Modeling to capture the intricate temporal variations in time series data analysis is proposed in (Wu et al. 2022). Despite the advantages of TCNs, the locality of one-dimensional convolution kernels and receptive fields restricts their ability to capture relationships between distant time points (Luo and Wang 2024).

### Transformers

Transformers have revolutionized time series forecasting models with their powerful attention mechanisms. Unlike RNNs and TCNs, transformers use a self-attention mechanism to weigh the importance of time steps. Accordingly, they can effectively capture long-term dependencies and relationships between distant time points (Zeng et al. 2023). Pyramidal Attention Module (PAM) is introduced in (Liu et al. 2022) to model temporal dependencies with low complexity in long-term forecasting tasks. In Autoformer model (Wu et al. 2021), a novel decomposition architecture with an Auto-Correlation mechanism is designed to effectively capture the seasonal and trend patterns of time series data. Besides, to further enhance the seasonal and trend decomposition, a sparse attention within the frequency domain is designed as the Frequency Enhanced Decomposed Transformer (FEDformer) (Zhou et al. 2022), which enhances the performance of Transformers for long-term forecasting. Despite their strengths, it is challenging for the attention mechanism to reliably discern dependencies directly from

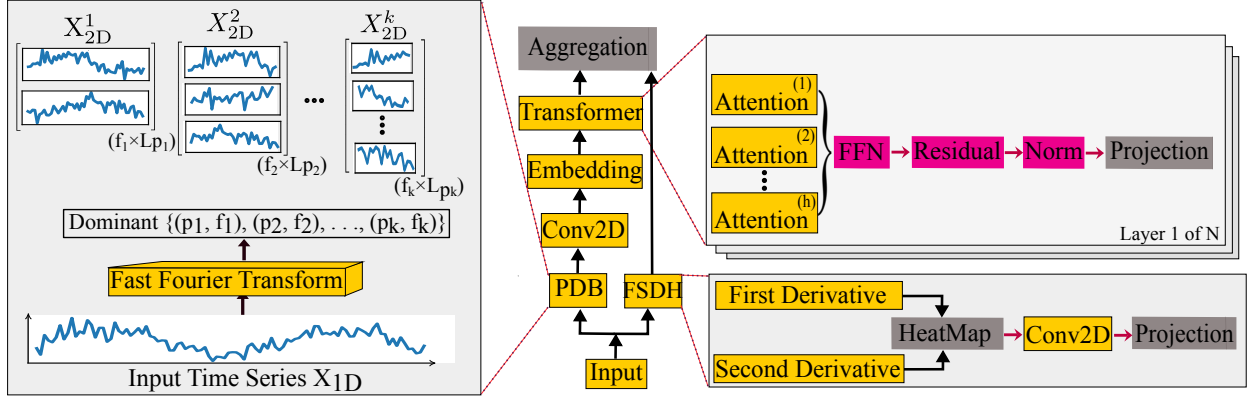


Figure 1: Overview of the proposed Times2D framework for general time series forecasting.

scattered time points because temporal dependencies can be deeply hidden by intricate temporal patterns (Wu et al. 2022).

### Multi-layer Perceptron

Multi-layer Perceptrons (MLPs) are fundamental neural network architectures consisting of fully connected layers. A deep neural architecture called N-BEATS is proposed by (Oreshkin et al. 2019) for univariate time series forecasting. The architecture uses backward and forwards residual links and a very deep stack of fully-connected layers. A lightweight deep learning architecture called LightTS is introduced by (Zhang et al. 2022) for multivariate time series forecasting. LightTS utilizes simple MLP-based structures and incorporates two down-sampling strategies, interval sampling, and continuous sampling, to preserve the majority of the time series information. Besides, to address volatility and computational complexity in long-horizon forecasting, Neural Hierarchical Interpolation for Time Series Forecasting (NHITS) is introduced in (Challu et al. 2023).

Although valuable studies have been done in time series forecasting tasks, we reveal intricate temporal information within time series data by employing multi-periodicity and derivative representations to explore temporal patterns in 2D space. In this transformation from 1D to 2D, 2D kernels are able to capture various complex temporal variations such as periodicity, rising, falling, sharp fluctuations, and turning points for the first time.

### Times2D

Given a multivariate time series dataset  $X = [x_1, x_2, \dots, x_T] \in R^{T \times N}$ , where  $T$  represents the length of the time series and  $N$  denotes the number of variables or dimensions. The dataset must be appropriately formatted by defining the input and output sequences. First, a segment of the time series, referred to as the sequence length  $S$ , is selected as the input. A subsequent segment of the time series, referred to as the prediction length  $P$ , is selected as the output. By following this approach, the original time series of total length  $T$  is transformed into multiple rows, where each row contains  $S$  input time points

and  $P$  corresponding output. Additionally, the data is organized into batches for parallel processing during training. As a result, the samples are organized into batches of size  $B$ , resulting in input tensors with dimensions  $[B, S, N]$  and output tensors with dimensions  $[B, P, N]$ .

### Periodic Decomposition Block

The PDB identifies dominant periods within the time series data in frequency domain. Given the input tensor  $X_{1D}$  with dimensions  $[B, S, N]$ , the FFT can be calculated as:

$$X_f = \sum_{t=0}^{S-1} X(t)e^{-2\pi itf/S} \quad (1)$$

where  $X_f$  represents the transformed tensor in the frequency domain and  $f$  is the frequency. The magnitude  $A_f$  can be calculated as follows:

$$A_f = |X_f| = \sqrt{Re(X_f)^2 + Im(X_f)^2} \quad (2)$$

Since  $X_f$  is symmetric, we maintain only the first half of the frequency components, corresponding to the non-negative frequencies. Accordingly, the magnitude  $A_f$  becomes a tensor of dimensions  $B \times \frac{S}{2} \times N$ . To identify the dominant periods for each sequence, the magnitudes are averaged over the  $B$  and  $N$  dimensions. The resulting magnitudes are then sorted in descending order, and the dominant frequencies are selected based on their associated magnitudes. As a result, the dominant periods  $p_1, p_2, \dots, p_k$ , associated with the dominant frequencies  $f_1, f_2, \dots, f_k$ , can be calculated as:

$$P_k = \frac{S}{f_k} \quad (3)$$

Thus, the original 1D input tensor  $X \in R^{B \times S \times N}$  can be decomposed into  $k$  distinct 2D tensors  $X_{2D}^i \in R^{B \times p_i \times f_i \times N}$ , each corresponding to one of the identified dominant periods. Then the tensors  $X_{2D}^i$  are passed through 2D convolutional layers to extract intricate relationships within and between periods. This convolution operates across both the periodic dimension  $p_i$  and the frequency dimension  $f_i$  as follows:

$$\hat{X}_{2D}^i = Conv2D(X_{2D}^i) \quad (4)$$

Where  $\hat{X}_{2D}^i$  represents the output features after applying the convolutional kernel. Subsequently, the 2D tensors  $\hat{X}_{2D}^i$  are fed into the following mechanisms:

**Positional Embedding (PosEnc):** PosEnc encodes positional information to maintain the temporal and spatial structure within the 2D tensors.

$$\hat{X}_{2D}^i = PosEnc(\hat{X}_{2D}^i) \quad (5)$$

**Multi-Head Self-Attention (MHSA):** MHSA Captures intricate dependencies within the 2D tensors by considering relationships across spatial and temporal dimensions.

$$\hat{X}_{2D}^i = MHSA(\hat{X}_{2D}^i) \quad (6)$$

**Feed-Forward Networks (FFN):** FFN applies non-linear transformations to enhance the ability of the model to learn complex patterns in the data.

$$\hat{X}_{2D}^i = FFN(\hat{X}_{2D}^i) \quad (7)$$

**Residual Connection (Res):** A residual connection is applied to preserve the original information from the 2D tensors and facilitate gradient flow during training.

$$\hat{X}_{2D}^i = \hat{X}_{2D}^i + Res(\hat{X}_{2D}^i) \quad (8)$$

**Normalization (Norm):** Norm normalizes the tensor to stabilize training and improve convergence.

$$\hat{X}_{2D}^i = Norm(\hat{X}_{2D}^i) \quad (9)$$

The normalized outputs are then transformed back into a 1D representation and the desired prediction length by applying a linear transformation as:

$$\hat{X}_{1D}^{PDB} = Linear(\hat{X}_{2D}^i) \quad (10)$$

Where  $\hat{X}_{1D}^{PDB}$ , with shape  $[B, P, N]$ , denotes the processed 1D representation obtained from the PDB, which serves as the input to subsequent layers for generating the final forecasted values.

### First and Second Derivative Heatmaps

To effectively capture sharp fluctuations and turning points within the time series data, we compute the first and second derivatives. Given the input tensor  $X_{1D}$  with dimensions  $[B, S, N]$ , the first derivative tensor  $D_1$  can be calculated as:

$$D_1(t) = X_{1D}(t+1) - X_{1D}(t) \quad (11)$$

To maintain the original sequence length, the first derivative tensor is padded at the beginning with zeros. Next, the second derivative tensor  $D_2$  is computed from the first derivative tensor as:

$$D_2(t) = D_1(t+1) - D_1(t) \quad (12)$$

Similarly, the second derivative tensor is padded with zeros at the beginning. To construct the final 2D heatmap tensor  $H_{2D}$  the padded first and second derivative tensors are stacked along a new dimension as:

$$H_{2D}(t, D) = Stack(D_1(t), D_2(t)) \quad (13)$$




Figure 2: 2D heatmap representation of a 1D sequence with a length of 96.

For example, Figure 2 demonstrates a 2D heatmap representation of a 1D sequence with a length of 96. These derivatives are 2D representations where the X-axis represents time, and the Y-axis is conceptually divided into two parts, with the lower half representing the first derivative (indicating sharp fluctuations) and the upper half representing the second derivative (indicating turning points). The colors in the heatmap correspond to the magnitude of these derivatives, with more intense colors indicating greater changes. Then, to enhance the ability of the model to extract meaningful temporal features, the 2D derivative heatmap tensor is processed through a series of convolutional layers.

$$\hat{H}_{2D}(t, D) = Conv2D(H_{2D}(t, D)) \quad (14)$$

Finally, the extracted features are transformed back into a 1D representation and the desired prediction length by applying a learned set of weights and performing a weighted summation across the derivative dimension as:

$$\hat{X}_{1D}^{FSDH} = Sum(\hat{H}_{2D}(t, D).weights) \quad (15)$$

Where  $\hat{X}_{1D}^{FSDH}$ , with shape  $[B, P, N]$ , denotes the processed 1D representation obtained from the FSDH block, which serves as the input to subsequent layers for generating the final forecasted values.

### Aggregation Block

After obtaining the 1D representations  $\hat{X}_{1D}^{PDB}$  from the PDB and  $\hat{X}_{1D}^{FSDH}$  from the FSDH, these two components are aggregated to create a combined feature representation for the final forecasting task. This is achieved by performing an element-wise summation of the two 1D representations:

$$\hat{X}_{1D} = \hat{X}_{1D}^{PDB} + \hat{X}_{1D}^{FSDH} \quad (16)$$

Data	Metric	Times2D (Ours)	N-HITS (2023)	PatchTST (2023)	DLinear (2023)	TimesNet (2023)	MICN (2023)	SCINet (2023)	FEDformer (2023)	Stationary (2022)	FiLM (2022)	LightTS (2022)	Autoformer (2021)	Informer (2021)
Yearly	SMAPE	<b>13.418</b>	<u>13.422</u>	13.477	15.11	15.378	25.022	18.605	13.728	13.717	17.431	14.247	13.974	14.727
	MASE	<b>2.990</b>	3.056	<u>3.019</u>	3.565	3.554	7.162	4.471	3.048	3.078	4.043	3.109	3.134	3.418
	OWA	<b>0.787</b>	0.795	<u>0.792</u>	0.911	0.918	1.667	1.132	0.803	0.807	1.042	0.827	0.822	0.881
Quarterly	SMAPE	<u>10.362</u>	<b>10.185</b>	10.380	10.597	10.465	15.214	14.871	10.792	10.958	12.925	11.364	11.338	11.360
	MASE	<u>1.221</u>	<b>1.180</b>	1.233	1.253	1.227	1.963	2.054	1.283	1.325	4.664	1.328	1.365	1.401
	OWA	<u>0.915</u>	<b>0.893</b>	0.921	0.938	0.923	1.407	1.424	0.958	0.981	1.193	1.000	1.012	1.024
Monthly	SMAPE	<b>12.918</b>	13.059	<u>12.959</u>	13.258	13.513	16.943	14.925	14.260	13.917	15.407	14.014	13.958	14.062
	MASE	<b>0.959</b>	1.013	<u>0.970</u>	1.003	1.039	1.442	1.131	1.102	1.097	1.298	1.053	1.103	1.141
	OWA	<b>0.898</b>	0.929	<u>0.905</u>	0.931	0.957	1.265	1.027	1.012	0.998	1.144	0.981	1.002	1.024

Table 1: Short-term forecasting results on the M4 dataset with prediction lengths of 6 for yearly, 8 for quarterly, and 18 for monthly. **Bold** and underlined values indicate the best and second-best performance, respectively.

The aggregated representation,  $\hat{X}_{1D} \in R^{B \times P \times N}$ , encapsulates various temporal variations captured by the PDB and FSDH. This combined representation is subsequently passed through a loss function during the training phase to optimize the model parameters for accurate forecasting.

## Experimental Results

### Datasets

We evaluated the Times2D model using several publicly available datasets commonly used in time series forecasting. For long-term forecasting, we utilized datasets like the electricity transformer temperature (ETT) and transportation traffic data. These datasets cover periods ranging from July 2016 to July 2018 with frequencies of 15 minutes (ETTm1 and ETTm2) to 1 hour (ETTh1, ETTh2, and Traffic). For short-term forecasting, we employed sub-sets of the M4 dataset, which include time series from finance, economics, demographics, and industry, with frequencies ranging from monthly to yearly (Jin et al. 2023).

### Evaluation Metrics

To comprehensively evaluate the performance of the Times2D, we used mean square error (MSE) and mean absolute error (MAE) for long-term forecasting, and symmetric mean absolute percentage error (SMAPE), mean absolute scaled error (MASE), and overall weighted average (OWA) for short-term forecasting. These metrics can be calculated as follows:

$$\text{MSE} = \frac{1}{H} \sum_{i=1}^H (\hat{\mathbf{Y}}_i - \mathbf{Y}_i)^2 \quad (17)$$

$$\text{MAE} = \frac{1}{H} \sum_{i=1}^H |\hat{\mathbf{Y}}_i - \mathbf{Y}_i| \quad (18)$$

$$\text{SMAPE} = \frac{200}{H} \sum_{i=1}^H \frac{|\hat{\mathbf{Y}}_i - \mathbf{Y}_i|}{|\hat{\mathbf{Y}}_i| + |\mathbf{Y}_i|} \quad (19)$$

$$\text{MASE} = \frac{\frac{1}{H} \sum_{i=1}^H |\hat{\mathbf{Y}}_i - \mathbf{Y}_i|}{\frac{1}{H-s} \sum_{j=s+1}^H |\mathbf{Y}_j - \mathbf{Y}_{j-s}|} \quad (20)$$

$$\text{OWA} = \frac{1}{2} \left( \frac{\text{SMAPE}}{\text{SMAPE}_{naive}} + \frac{\text{MASE}}{\text{MASE}_{naive}} \right) \quad (21)$$

Where  $\hat{\mathbf{Y}}_i$  and  $\mathbf{Y}_i$  denote the predicted and actual values for the  $i$ -th observation,  $H$  represents the forecasting horizon,  $s$  is the seasonal period length, and  $\text{SMAPE}_{naive}$  and  $\text{MASE}_{naive}$  refer to the SMAPE and MASE values obtained using a naive forecasting method (Jin et al. 2023).

### Baselines

The performance of Times2D is compared with multiple state-of-the-art algorithms in the time series forecasting domain. These models are categorized into different model architectures. Transformer-based models include PatchTST (2023), FEDformer (2022b), Stationary (2022b), Autoformer (2021), Informer (2021), CARD (2024), SageFormer (2024), and Crossformer (2023). MLP-based models encompass N-HITS (2023), DLinear (2023), and LightTS (2022). TCN-based models include TimesNet (2023), MICN (2023), and SCINet (2022a). Additionally, the comparison includes GPT4TS from the category of large language models (LLMs).

### Setups

To maintain consistency with the baseline models, the training, validation, and test sets were normalized using the mean and standard deviation derived from the training data. Prediction horizons  $H = \{96, 192, 336, 720\}$  are used for long-term forecasting tasks. In addition, for short-term forecasting tasks, the prediction lengths were chosen as follows: 6 for yearly, 8 for quarterly, and 18 for monthly.

### Key Results

The results demonstrate the superior performance of the Times2D model compared to baseline forecasting models.

**Short-term forecasting:** Table 1 presents the short-term forecasting performance of Times2D across the M4 datasets. For the yearly data, Times2D achieved the lowest SMAPE at 13.418, surpassing advanced models like N-HITS (13.422)

Models	$H$	Time2D (Ours)		CARD (2024)		SageFormer (2024)		N-HITS (2023)		GPT4TS (2023)		PatchTST (2023)		DLinear (2023)		TimesNet (2023)		Crossformer (2023)	
		MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTh1	96	<b>0.362</b>	<b>0.395</b>	0.383	0.391	0.377	0.397	0.378	0.393	0.376	<u>0.397</u>	<u>0.370</u>	0.399	0.375	0.399	0.384	0.402	0.423	0.448
	192	<b>0.405</b>	0.423	0.435	0.420	0.423	0.425	0.427	0.436	0.416	<u>0.418</u>	<u>0.413</u>	<b>0.416</b>	<b>0.405</b>	<b>0.416</b>	0.436	0.429	0.471	0.474
	336	<u>0.423</u>	<b>0.433</b>	0.479	0.442	0.459	0.445	0.458	0.484	0.442	<b>0.433</b>	<b>0.422</b>	0.436	0.439	0.443	0.491	0.469	0.570	0.546
	720	<b>0.445</b>	0.465	0.471	<b>0.461</b>	0.465	0.466	0.472	0.561	0.477	0.456	<u>0.447</u>	<u>0.466</u>	0.472	0.490	0.521	0.500	0.653	0.621
	Avg	<b>0.410</b>	<b>0.429</b>	0.442	<b>0.429</b>	0.431	0.433	0.433	0.468	0.465	0.455	<u>0.413</u>	<b>0.429</b>	0.422	0.437	0.458	0.450	0.529	0.522
ETTh2	96	<b>0.270</b>	<u>0.334</u>	0.281	<b>0.330</b>	0.286	0.338	0.274	0.345	0.285	0.342	<u>0.274</u>	<u>0.336</u>	0.289	0.353	0.340	0.374	0.745	0.584
	192	<b>0.334</b>	<b>0.378</b>	0.363	0.381	0.368	0.394	0.353	0.401	0.354	0.389	<u>0.339</u>	<u>0.379</u>	0.383	0.418	0.402	0.414	0.877	0.656
	336	<b>0.329</b>	<u>0.385</u>	0.411	0.418	0.413	0.429	0.382	0.425	<u>0.373</u>	0.407	<b>0.329</b>	<b>0.380</b>	0.448	0.465	0.452	0.452	1.043	0.731
	720	<b>0.376</b>	<b>0.422</b>	0.416	0.431	0.427	0.449	0.625	0.557	0.406	0.441	<u>0.379</u>	<u>0.422</u>	0.605	0.551	0.462	0.468	1.104	0.763
	Avg	<b>0.327</b>	<b>0.379</b>	0.368	<u>0.390</u>	0.374	0.403	0.408	0.507	0.381	0.412	<u>0.330</u>	<b>0.379</b>	0.431	0.446	0.414	0.427	0.942	0.684
ETTh1	96	<b>0.279</b>	<b>0.343</b>	0.316	0.347	0.324	0.362	0.302	0.350	0.292	0.346	<u>0.290</u>	<u>0.342</u>	0.299	0.343	0.338	0.375	0.404	0.426
	192	<b>0.313</b>	<b>0.361</b>	0.363	0.370	0.368	0.387	0.374	0.383	0.332	0.372	<u>0.332</u>	<u>0.369</u>	0.335	<u>0.365</u>	0.327	0.387	0.450	0.451
	336	<b>0.347</b>	<b>0.380</b>	0.392	0.390	0.401	0.408	0.369	0.402	<u>0.366</u>	0.394	<u>0.366</u>	<u>0.392</u>	0.369	<u>0.386</u>	0.410	0.411	0.532	0.515
	720	<b>0.410</b>	<b>0.413</b>	0.458	0.425	0.457	0.441	0.431	0.441	0.417	0.421	<u>0.416</u>	<u>0.420</u>	0.425	0.421	0.478	0.450	0.666	0.589
	Avg	<b>0.337</b>	<b>0.374</b>	0.383	0.384	0.388	0.400	0.369	0.394	0.388	0.403	<u>0.351</u>	0.380	0.357	<u>0.378</u>	0.400	0.406	0.513	0.495
ETTh2	96	<b>0.166</b>	0.259	0.169	<b>0.248</b>	0.173	0.255	0.176	0.255	0.173	0.262	<b>0.166</b>	<u>0.255</u>	<u>0.167</u>	0.269	0.187	0.267	0.287	0.366
	192	<b>0.220</b>	<u>0.296</u>	0.234	<b>0.292</b>	0.239	0.299	0.245	0.305	0.229	0.301	<u>0.223</u>	<u>0.296</u>	0.224	0.303	0.249	0.309	0.414	0.492
	336	<b>0.272</b>	<u>0.330</u>	0.294	0.339	0.299	0.388	0.295	0.346	0.286	0.341	<u>0.274</u>	<b>0.329</b>	0.281	0.342	0.321	0.351	0.597	0.542
	720	<b>0.342</b>	<b>0.377</b>	0.390	0.388	0.395	0.395	0.401	0.413	0.378	0.401	<u>0.362</u>	<u>0.385</u>	0.397	0.421	0.408	0.403	1.730	1.042
	Avg	<b>0.250</b>	<b>0.315</b>	0.272	<u>0.317</u>	0.277	0.322	0.279	0.329	0.284	0.339	<u>0.255</u>	<b>0.315</b>	0.267	0.333	0.291	0.333	0.757	0.610
Weather	96	<b>0.146</b>	0.197	0.150	<b>0.188</b>	0.162	0.206	0.158	<u>0.195</u>	0.162	0.212	<u>0.149</u>	0.198	0.176	0.237	0.172	0.220	0.195	0.271
	192	<b>0.191</b>	0.240	0.202	<b>0.238</b>	0.211	0.250	0.211	<u>0.247</u>	0.204	0.248	<u>0.194</u>	0.241	0.220	0.282	0.219	0.261	0.209	0.277
	336	<b>0.243</b>	<b>0.282</b>	0.260	<b>0.282</b>	0.271	0.294	0.274	0.300	0.254	0.286	<u>0.245</u>	<b>0.282</b>	0.265	0.319	0.280	0.306	0.273	0.332
	720	<b>0.313</b>	<b>0.333</b>	0.343	0.353	0.345	0.343	0.401	0.413	0.326	0.337	<u>0.314</u>	<u>0.334</u>	0.333	0.362	0.365	0.359	0.379	0.401
	Avg	<b>0.223</b>	<u>0.263</u>	0.239	<b>0.261</b>	0.247	0.273	0.261	0.288	0.237	0.270	<u>0.225</u>	<u>0.263</u>	0.248	0.300	0.259	0.287	0.264	0.320

Table 2: Long-term forecasting results with prediction horizons of  $H = \{96, 192, 336, 720\}$  for all experiments. **Bold** and underlined values indicate the best and second-best performance, respectively.

and PatchTST (13.477). This demonstrates the capability of Times2D to effectively capture long-term seasonal patterns. In the quarterly data, Times2D outperforms competing models such as TimesNet (10.465) and FEDformer (10.792) with a SMAPE of 10.362, emphasizing its robustness in managing seasonal variations. In addition, Times2D outperforms PatchTST (12.959) and N-HITS (13.059) in the monthly forecasting task, demonstrating its capability to forecast at higher temporal frequencies.

**Long-term forecasting:** The long-term forecasting results are provided in Table 2. Times2D demonstrates superior performance in comparison to state-of-the-art models across multiple datasets and forecasting horizons  $H = \{96, 192, 336, 720\}$ . Specifically, Times2D achieves an average reduction of 5% in MSE and 4% in MAE across the datasets compared to advanced models like PatchTST and N-HITS. Times2D consistently provides lower error rates, particularly for longer forecasting horizons, with up to 8% improvement in MSE and 7% in MAE for the ETTh2 dataset at  $H = 720$  compared to SageFormer, the best-second model. This highlights the strength of Times2D in capturing complex temporal dependencies over extended periods. The robustness of Times2D is further demonstrated by its superior performance over transformer-based models such as FEDformer and DLinear, achieving an average

6% reduction in MSE across the Weather dataset. This highlights the effectiveness of Times2D in handling diverse data types. Additionally, comparisons with recent models like GPT4TS and CARD reveal that Times2D exceeds their performance, achieving a 7% reduction in MAE on the ETTh1 dataset.

### Computational Efficiency Analysis

In this study, all simulations were conducted using PyTorch on a single NVIDIA L40S 46 GB GPU. The computational efficiency of the Times2D model was evaluated and compared with several state-of-the-art models, including TimesNet, PatchTST, Autoformer, Crossformer, and FEDFormer. To maintain consistency with the baseline models, we used average time per iteration and average RAM memory (MB) usage as key metrics across different prediction lengths  $H = \{384, 768, 1536\}$ . The results of this evaluation are summarized in Table 3. Times2D demonstrates a strong balance between computational speed and resource efficiency. The model maintains a consistently low average time per step across all prediction lengths, indicating its ability to scale effectively without significant performance degradation. In terms of memory usage, Times2D shows minimal variation, with RAM consumption ranging from 4329 MB to 4340 MB as prediction lengths increase. This low memory footprint highlights the suitability of Times2D for resource-

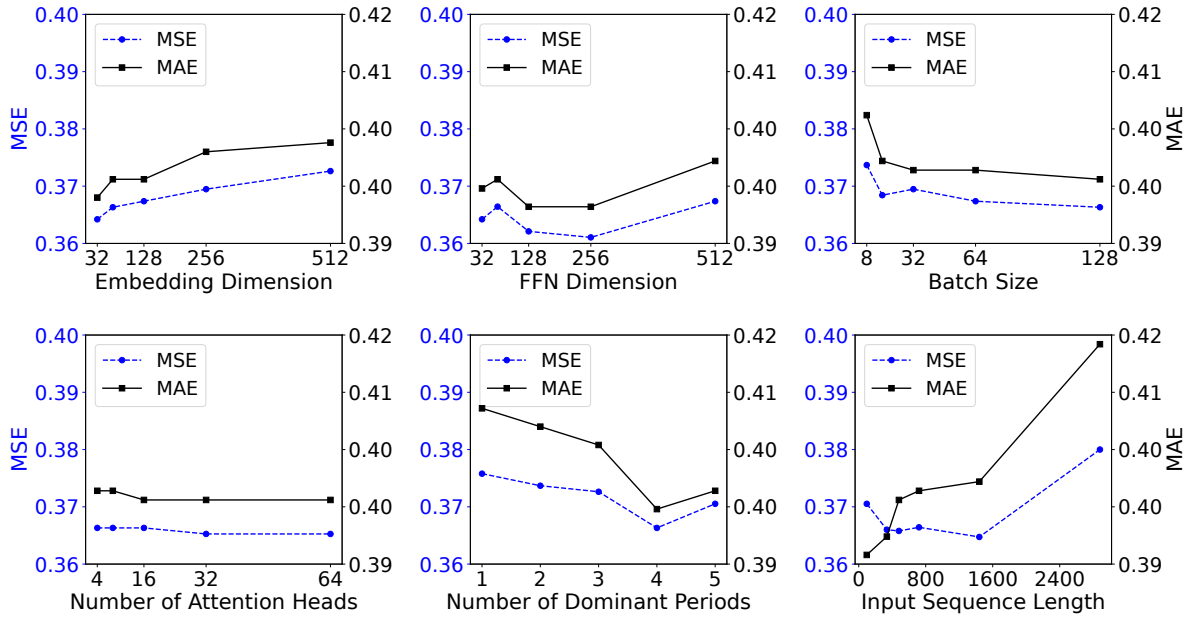


Figure 3: Sensitivity analysis of key hyperparameters for Times2D on the ETTh1 dataset with a prediction horizon of  $H=96$ .

Efficiency	RAM Memory (MB)			Running Time (S/iter)		
	384	768	1536	384	768	1536
H						
Times2D	4329	4340	4324	0.045	0.050	0.052
PatchTST	6565	6805	8586	0.014	0.016	0.024
TimesNet	3908	3911	3903	0.018	0.025	0.079
FEDformer	18758	22341	28596	0.105	0.106	0.132
Crossformer	5674	5985	7272	0.022	0.023	0.028
Autoformer	6566	9914	16655	0.052	0.053	0.957

Table 3: RAM memory usage (MB) and running time (S/iter) for prediction horizons  $H = \{384, 768, 1536\}$ .

constrained environments. In contrast, other models exhibited more pronounced increases in both computation time and memory usage with longer prediction lengths, underscoring the efficiency of Times2D in handling large-scale time series data.

### Sensitivity Analysis of Hyperparameters

In this section, we examine key hyperparameters of the model, including embedding size, feedforward network size, batch size, attention heads, number of dominant periods, and input forecasting length. To do that, we conducted this analysis on ETTh1 data with a fixed forecasting horizon of 96, where only one hyperparameter changes while others remain fixed in each scenario. As shown in Figure 3, variations in the number of attention heads slightly affect model performance. The model performs better with input lengths up to 1440. Beyond this, both MSE and MAE increase signifi-

cantly, indicating that very long input sequences are ineffective. Lower embedding dimensions (32-64) and a feed-forward network size in the range of 128-256 provide better performance, while larger sizes gradually increase both MSE and MAE, raising the risk of overfitting. A batch size of 128 balances computational efficiency and predictive accuracy, with smaller batches resulting in higher errors and larger batches offering no significant improvement.

### Conclusions

In this paper, we introduced Times2D, an innovative and efficient framework for general time series forecasting that transforms 1D time series into 2D representations. This approach effectively captures complex temporal patterns, including multi-periodicity, sharp fluctuations, and turning points, which are often challenging for traditional 1D models. The proposed Periodic Decomposition Block (PDB) addresses multi-periodicity by converting the time series into a 2D tensor in the frequency domain, while the First and Second Derivative Heatmaps (FSDH) emphasize sharp changes and turning points. Extensive experiments on real-world datasets demonstrate that Times2D achieves superior performance and computational efficiency in both short-term and long-term forecasting, outperforming other state-of-the-art models. In the future, Times2D could be extended to address the detection of rare events, focusing on scenarios where anomalous or infrequent patterns affect forecasting accuracy significantly. Additionally, The model strengths and limitations can be examined in greater detail, rather than solely relying on aggregate metrics like MSE and MAE.

## Acknowledgments

The authors declare financial support was received for the research and authorship of this article. This research was conducted with funding from NSF through award Number 2225341.

## References

- Ahmadi, A.; Zargaran, Z.; Mohebi, A.; and Taghavi, F. 2014. Hybrid model for weather forecasting using ensemble of neural networks and mutual information. In *2014 IEEE geoscience and remote sensing symposium*, 3774–3777. IEEE.
- 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.
- 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 *The Twelfth International Conference on Learning Representations*.
- Fu, Y.; Wu, D.; and Boulet, B. 2022. Reinforcement learning based dynamic model combination for time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 6639–6647.
- He, H.; Zhang, Q.; Bai, S.; Yi, K.; and Niu, Z. 2022. CATN: Cross attentive tree-aware network for multivariate time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 4030–4038.
- Huang, L.; Yuan, J.; Chen, S.; and Li, X. 2024. MDG: A Multi-Task Dynamic Graph Generation Framework for Multivariate Time Series Forecasting. *IEEE Transactions on Emerging Topics in Computational Intelligence*.
- 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*.
- Li, Y.; Xu, J.; and Anastasiu, D. C. 2023. An extreme-adaptive time series prediction model based on probability-enhanced lstm neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 8684–8691.
- 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 *#PLACEHOLDER\_PARENT\_METADATA\_VALUE#*.
- Luo, D.; and Wang, X. 2024. Modernctn: A modern pure convolution structure for general time series analysis. In *The Twelfth International Conference on Learning Representations*.
- Massart, E.; and Abrol, V. 2022. Coordinate descent on the orthogonal group for recurrent neural network training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 7744–7751.
- Oreshkin, B. N.; Carpov, D.; Chapados, N.; and Bengio, Y. 2019. N-BEATS: Neural basis expansion analysis for interpretable time series forecasting. *arXiv preprint arXiv:1905.10437*.
- Tan, C.; Li, S.; Gao, Z.; Guan, W.; Wang, Z.; Liu, Z.; Wu, L.; and Li, S. Z. 2023. Openstl: A comprehensive benchmark of spatio-temporal predictive learning. *Advances in Neural Information Processing Systems*, 36: 69819–69831.
- 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 *The eleventh international conference on learning representations*.
- 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.
- 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, 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, Z.; Meng, L.; and Gu, Y. 2024. SageFormer: Series-Aware Framework for Long-Term Multivariate Time Series Forecasting. *IEEE Internet of Things Journal*.
- 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.