

PHAT: MODELING PERIOD HETEROGENEITY FOR MULTIVARIATE TIME SERIES FORECASTING

Jiaming Ma¹, Qihe Huang¹, Guanjun Wang¹, Haofeng Ma³, Sheng Huang¹,
 Zhengyang Zhou^{1,2}, Pengkun Wang^{1,2}, Xu Wang^{1,2}, Binwu Wang^{1,2,*}, Yang Wang^{1,2,*}

¹University of Science and Technology of China (USTC), Hefei, Anhui, China

²Suzhou Institute for Advanced Research, USTC, Suzhou, Jiangsu, China

³The University of Nottingham Ningbo China, Ningbo, Zhejiang, China

ABSTRACT

While existing multivariate time series forecasting models have advanced significantly in modeling periodicity, they largely neglect the periodic heterogeneity common in real-world data, where variates exhibit distinct and dynamically changing periods. To effectively capture this periodic heterogeneity, we propose PHAT (Period Heterogeneity-Aware Transformer). Specifically, PHAT arranges multivariate inputs into a three-dimensional "periodic bucket" tensor, where the dimensions correspond to variate group characteristics with similar periodicity, time steps aligned by phase, and offsets within the period. By restricting interactions within buckets and masking cross-bucket connections, PHAT effectively avoids interference from inconsistent periods. We also propose a positive-negative attention mechanism, which captures periodic dependencies from two perspectives: periodic alignment and periodic deviation. Additionally, the periodic alignment attention scores are decomposed into positive and negative components, with a modulation term encoding periodic priors. This modulation constrains the attention mechanism to more faithfully reflect the underlying periodic trends. A mathematical explanation is provided to support this property. We evaluate PHAT comprehensively on **14** real-world datasets against **18** baselines, and the results show that it significantly outperforms existing methods, achieving highly competitive forecasting performance. Our sources is available at [GitHub](#).

1 INTRODUCTION

Multivariate Time Series (MTS) forecasting serves as a core enabling technology for critical applications such as energy demand prediction, traffic management, financial modeling, and healthcare monitoring (Ma et al., 2025b; Zhou et al., 2021; Zeng et al., 2023; Qiu et al., 2024b; Wu et al., 2021).

Periodicity, as a crucial intrinsic characteristic of time series data, plays a decisive role in enhancing forecasting performance through accurate modeling (Lin et al., 2024a; Zhou et al., 2022b; Lin et al., 2024b). To this end, researchers have developed a series of cutting-edge methods. One line of work modifies neural network architectures to adapt to periodicity (Luo & Wang, 2024). For instance, some studies leverage the strong capability of the Transformer architecture to model long-range dependencies (Zhou et al., 2021; 2022b; Nie et al., 2022). In addition, seasonal-trend decomposition techniques have been widely employed (Wang et al., 2024b; Hu et al., 2025; Ma et al., 2025d). These methods separate the original time series into seasonal and trend components, which are then modeled by parallel sub-networks. This strategy enables more efficient extraction and utilization of periodic information. Furthermore, recent studies have incorporated classical signal processing tools, particularly frequency domain analysis methods like the Fast Fourier Transform (FFT), to more precisely identify and model periodic patterns in time series data (Ye et al., 2024; Zhang et al., 2025).

Despite the substantial progress, two major limitations remain: ① Most existing models unify periodic modeling by treating variates as interchangeable channels for pooling and fusion, implicitly assuming a shared, static periodic length. This overlooks the pronounced heterogeneity in periodic behavior across variates. As shown in Figure 1, three variates from the ZafNoo dataset exhibit distinct period

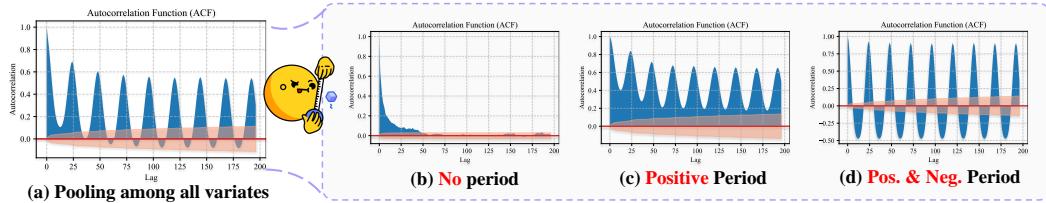


Figure 1: The visualization of period heterogeneity phenomenon on ZafNoo dataset. The orange region in autocorrelation function represents the 95% confidence interval for the null hypothesis from Bartlett’s Test (Arsham & Lovric, 2011). Correlation coefficients that fall within this interval are not statistically significant and cannot be rejected as noise (Chandler, 1987).

lengths. Forcing such diverse periodicities into the unified framework can lead models to learn spurious temporal dynamics. ② Mainstream Transformer architectures amplify positive correlations but suppress negative ones during attention normalization, overlooking inverse or complementary dynamics inherent in periodic signals, as shown in Figure 1 (d). Yet such negative correlations offer critical insights into system dynamic. Therefore, it is crucial to develop multivariate time series forecasting models capable of accurately capturing periodic heterogeneity.

In this paper, we propose a Period Heterogeneity-Aware Transformer (PHAT) for MTS forecasting. Specifically, PHAT restructures multivariate time series into a periodic bucket structure by grouping variates based on their periodic lengths. Within each bucket, each sequence is further reshaped into a 2D tensor—rows align time steps by phase, while columns capture offsets within the period. Interactions are restricted to variates within the same bucket, enabling the model to capture diverse periodic patterns, while cross-bucket links are masked to prevent interference between variates with differing periodicities. Second, PHAT introduces a Positive-Negative Self-Attention mechanism (PNA), which interprets periodic dependencies through two attention coefficients: phase alignment and periodic offset. The periodic offset coefficient is further decomposed into positive and negative components, with a modulation term encoding periodic priors. This term reduces the positive weights and increases the negative weights for distant phase-aligned points (and vice versa), enabling the attention mechanism to more faithfully capture the periodic structure.

Contributions. ① **Initial Exploration.** We relax the single-period assumption to handle complex time series with heterogeneous periodicities and introduce PHAT, the first method expressly developed to model such period heterogeneity. ② **Periodic Bucket.** PHAT introduces a "periodic bucket" structure to manage time series data with heterogeneous periodicity, facilitating the learning of periodic patterns. ③ **Novel Self-attention.** We further propose a positive-negative attention mechanism that represents periodic information with distinct positive and negative components, combined with a modulation term encoding periodic priors. Its favorable properties are demonstrated mathematically. ④ **Empirical Validation.** On 14 real-world datasets with 18 baselines, PHAT achieves SOTA performance on approximately **73.95% (71/96)** of the metrics while maintaining low computational complexity. Additionally, PHAT demonstrates strong robustness in complex periodic scenarios.

2 RELATED WORK

Periodicity plays a pivotal role in the predictability of time series data (Lin et al., 2025; Wang et al., 2024b; Ma et al., 2025d). Recent advancements have introduced a range of sophisticated techniques aimed at improving the ability to capture long-term dependencies, thereby enhancing the perception of periodic patterns. For example, classical seasonal–trend decomposition uses moving-average kernels for sliding aggregation to extract trend components (Zeng et al., 2023; Kingma & Ba, 2014), while CycleNet proposes a learnable Cycle Decomposition to capture periodicity (Lin et al., 2024a). Architecturally, Transformer-based models excel at modeling long-range dependencies: Autoformer connects sequences by periodically aggregating similar subsequences (Wu et al., 2021), while ModernTCN leverages very large convolutional kernels to substantially enlarge the receptive field and capture long-range temporal dependencies (Luo & Wang, 2024). Frequency-domain analysis and multi-scale modeling have also been used to strengthen periodic representations (Zhou et al., 2022b; Wang et al., 2024b).

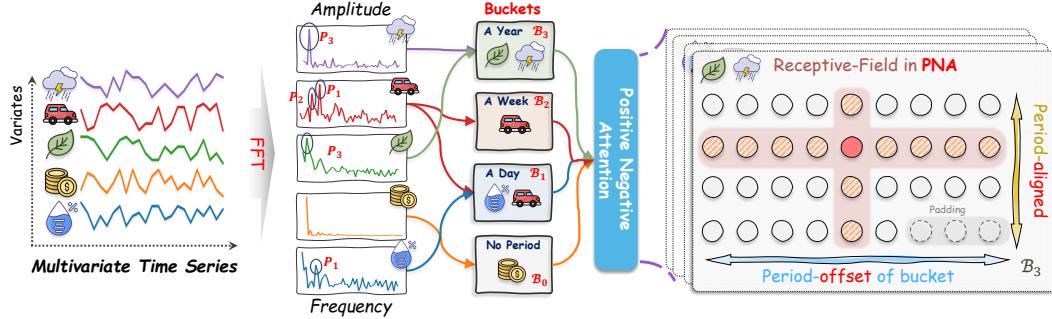


Figure 2: The overall architecture of PHAT. Each bucket contains the same periodic variates and uses PNA to capture the periodic attention mechanism.

However, existing methods typically handle multivariate series via pooling or adaptive fusion and model periodicity from a unified perspective, implicitly assuming all variates share the same periodic patterns. This overlooks widespread periodic heterogeneity and often ignores negatively correlated periodic components, which can provide complementary perspectives and valuable information gains. We provide a detailed discussion of multivariate time series forecasting in [Appendix A](#).

3 METHODOLOGY

Multivariate time-series data typically refer to either a collection of multiple temporal objects or a single object observed through multiple feature channels (Huang et al., 2024b; Wu et al., 2020; Ma et al., 2025e). Given a multivariate time series input $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T] \in \mathbb{R}^{C \times T}$ observed over the past T time steps, where $\mathbf{x}_t \in \mathbb{R}^C$ represents the observation at time step t across C variates. And we use $\mathbf{X}_i \in \mathbb{R}^T$ to represent the input sequence corresponding to variates i . The objective of the multivariate time-series forecasting is to forecast the subsequent L time steps $\mathbf{Y} = [\mathbf{x}_{T+1}, \mathbf{x}_{T+2}, \dots, \mathbf{x}_{T+L}] \in \mathbb{R}^{C \times L}$.

As shown in Figure 2, we propose PHAT for modeling periodic heterogeneity, a common yet previously overlooked property of real-world multivariate time series, thus motivating our PHAT. PHAT includes a novel period bucket structure to manage time series. Subsequently, PHAT integrates a positive-negative attention mechanism to precisely model the positive and negative correlations of periodicity. Finally, we perform weighted fusion of the generated periodic components based on their frequency saliency to produce the prediction results.

3.1 PERIOD BUCKET FOR TIME SERIES

3.1.1 PERIOD DETECTION

Fast Fourier Transform (FFT) (Duhamel & Vetterli, 1990) is a commonly used tool for analyzing the characteristics of time series, especially periodicity (Liu, 2025; Wu et al., 2022b; Fang et al., 2023). Therefore, we first apply the FFT to each variate in the input sequence \mathbf{X} . Subsequently, we retain only the spectral magnitudes corresponding to the Top- K significant frequency components, which are then converted into discrete period lengths. This process can be expressed as follows:

$$\mathbf{P} = \lfloor \frac{T}{\arg \text{Top}_K[\text{FFT}(\mathbf{X})]} + 0.5 \rfloor \in \mathbb{N}^{K \times C}, \quad (1)$$

where $\arg \text{Top}_K[\cdot]$ selects the indices of the K most salient frequencies. \mathbf{P} represents the set of period lengths from C variates.

3.1.2 TIME SERIES PERIOD BUCKET CONSTRUCTION AND REPRESENTATION

To improve periodicity modeling, we propose a novel period-bucket structure for reconstructing time series, consisting of two stages: bucketing and folding.

Bucketing This process groups variates by their dominant period lengths. Given $K * C$ periods, we first remove duplicate values, assuming there are \mathcal{N} distinct elements. For any period length P_i , we create a bucket matched to it, denoted as \mathcal{B}_i , and place the variates with a period length of P_i into the bucket. Because a variate can exhibit multiple periodicities, buckets are not necessarily disjoint. In addition, we assign Bucket-0 to manage variates that do not exhibit significant periodic behavior.

Folding This process further reshapes the sequence of each variate. For the input sequence of variate j in Bucket- b , denoted as \mathbf{X}_j , we first use a linear layer to align it with the future window: $\mathbf{X}_j \rightarrow \mathbf{X}_j \in \mathbb{R}^{L \times d}$. It is important to emphasize that the period length of variates in Bucket- b is P_b . Next, we segment \mathbf{X}_j into small fragments of length $\lfloor L/P_b \rfloor$, which may require zero-padding for \mathbf{X}_j . This process can be expressed mathematically as:

$$\begin{aligned} \bar{\mathbf{X}}_j &= \text{Unflatten}(\mathbf{X}_j^{\text{pad}}, P_b) \in \mathbb{R}^{P_b \times N_b}, \\ \mathbf{X}_j^{\text{pad}} &= \begin{cases} \text{Concat}(\mathbf{X}_j^\top, \mathbf{0}), & \text{if } L \bmod P_b > 0, \\ \mathbf{X}_j^\top, & \text{if } L \bmod P_b = 0. \end{cases} \in \mathbb{R}^{(P_b * N_b)} \end{aligned} \quad (2)$$

where $N_b = \lfloor L/P_b \rfloor$ represents the number of periods contained in the sequence. So for Bucket- b that contains of $|\mathcal{B}_b|$ time series, we can generate a 3D bucket-structured input denoted as $\bar{\mathbf{X}}^{(b)} \in \mathbb{R}^{|\mathcal{B}_b| \times P_b \times N_b}$, which adapts to the characteristics of periodic features: ① Periodic homogeneous variates are grouped in the first dimension; ② Time steps in the second dimension are periodic-offset within the same periodic; ③ Each time step in the last dimension is periodic-aligned.

Then, we perform interactions between variates to learn their dependencies as follows,

$$\mathbf{Z}^{(b)} = \left(\bar{\mathbf{X}}^{(b)} \right)^\top \mathbf{W}_i + \mathbf{b}_i \in \mathbb{R}^{P_b \times N_b \times d_h}, \quad (3)$$

where $\mathbf{W} \in \mathbb{R}^{|\mathcal{B}_b| \times d_h}$ and $\mathbf{b} \in \mathbb{R}^{d_h}$ are learnable parameters with the dimension size d_h . The processing of Bucket with periodicity \mathcal{B}_0 is represented in [Appendix B](#).

3.2 POSITIVE-NEGATIVE X-SHAPE ATTENTION FOR PERIODICITY MODELING

Conventional self-attention allows unrestricted token interactions, which is suboptimal for capturing the periodic structures inherent in time series. To address this, we propose Positive-Negative Attention for Periodicity Modeling (PNA), featuring three key designs: ① **X-Shaped Receptive Field**: PNA allocates attention along the rows and columns of the bucketed representation \mathbf{Z} , forming an X-shaped (cross-like) receptive field centered on each target element. This structure explicitly separates period-aligned (across-period) and period-offset (within-period) relationships, enhancing periodic pattern modeling. ② **Strong Periodic Inductive Bias**: Attention scores are modulated by a periodic-distance-dependent term, enforcing inductive bias that aligns attention weights with the underlying periodic trends. ③ **Decoupled Positive-Negative Correlation Modeling**: PNA separately model positive and negative periodic dependencies, mitigating interference between opposing correlations and enabling more expressive representation of complex periodic interactions.

Specifically, given Bucket- b representation $\mathbf{Z}^{(b)} \in \mathbb{R}^{P_b \times N_b \times d_h}$ (except for Bucket-0, which lacks periodic characteristics), we first obtain multiple components using the following formula:

$$[\mathbf{Q}_1; \mathbf{Q}_2] = \mathbf{Z}^{(b)} \mathbf{W}_q, \quad [\mathbf{K}_1; \mathbf{K}_2] = \mathbf{Z}^{(b)} \mathbf{W}_k, \quad \mathbf{V} = \mathbf{Z}^{(b)} \mathbf{W}_v, \quad \mathbf{\Lambda} = \sigma(\mathbf{Z}^{(b)} \mathbf{W}_g), \quad (4)$$

where $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{d_h \times d}$, $\mathbf{W}_\lambda^h \in \mathbb{R}^{d \times 1}$ are learnable projection parameters. \mathbf{Q}_1 and $\mathbf{Q}_2 \in \mathbb{R}^{P_b \times N_b \times 2d}$ are the query vectors, \mathbf{K}_1 and $\mathbf{K}_2 \in \mathbb{R}^{P_b \times N_b \times 2d}$ are the key vectors, and $\mathbf{V} \in \mathbb{R}^{P_b \times N_b \times 2d}$ is the value vector. $\mathbf{\Lambda}$ is the weighted strength filter. $\sigma(\cdot)$ is the sigmoid function. The calculation process of PNA can be written as:

$$\text{PNA}(\mathbf{Z}^{(b)}) = \text{Attention}([\mathbf{Q}_1; \mathbf{Q}_2], [\mathbf{K}_1; \mathbf{K}_2], \mathbf{V}, \mathbf{\Lambda}) = \bar{\mathbf{A}} \times_1 (\tilde{\mathbf{A}} \times_2 \mathbf{V}), \quad (5)$$

where \times_i means the multiplication in the i -th dimension of the matrix. Periodic-aligned attention $\bar{\mathbf{A}}$ encodes dependencies between time-steps that share the same phase across periods, while periodic-offset attention $\tilde{\mathbf{A}}$ captures dependencies among time-steps within the same period.

Period-offset Attention $\bar{\mathbf{A}}$ We compute the positive logits ζ and negative logits η of the period-offset attention with separate query and key in each head as follows,

$$\zeta = \mu \mathbf{Q}_1 \times_1 \mathbf{K}_1^\top \in \mathbb{R}^{P_b \times P_b \times N_b}, \quad \eta = \mu \mathbf{Q}_2 \times_1 \mathbf{K}_2^\top \in \mathbb{R}^{P_b \times P_b \times N_b}, \quad (6)$$

where $\mu = d^{-1/2}$ is the scale factor. The positive and negative logits are first adjusted by the periodic modulation terms and then separately normalized via softmax to produce attention coefficient matrices. These two matrices are subsequently fused to form the final attention matrix as follows,

$$\bar{\mathbf{A}} = \text{Softmax}(\tilde{\zeta}) - \mathbf{A} \odot \text{Softmax}(\tilde{\eta}) \in \mathbb{R}^{P_b \times P_b \times N_b}, \quad (7)$$

$$\tilde{\zeta}[m, n] = \zeta[m, n] - \underbrace{\sum_{s \in \Delta_{m,n}^{(b)}} \text{Softplus}(\zeta[m, s])}_{\text{Positive Modulation Term}}, \quad \tilde{\eta}[m, n] = \eta[m, n] - \underbrace{\sum_{s \in \nabla_{m,n}^{(b)}} \text{Softplus}(\eta[m, s])}_{\text{Negative Modulation Term}},$$

where $\zeta[m, n]$ denotes the attention coefficient between the m -th and n -th time steps within a signal of period P_b , corresponding to the (m, n) -th entry of the attention matrix ζ . We define the periodic relative distance between m -th and n -th time steps as $\delta_{m,n}^b$. For a fixed m and n , let $\Delta_{m,n}^p = \{s | \delta_{m,s}^b < \delta_{m,n}^b\} \cup \{m\}$ be the set of time steps whose periodic distance to m is smaller than that between m and n . We aggregate the attention coefficients $\zeta[m, s]$ for all $s \in \Delta_{m,n}^p$. After applying the Softplus(\cdot) activation (Zheng et al., 2015), this term encourages the model to produce attention weights that decay monotonically with increasing periodic distance, thereby reinforcing the inductive bias toward local periodic structure. The periodic distance is computed as follows:

$$\delta_{ij}^b = \min\{(i - j) \bmod \mathcal{B}_b, (j - i) \bmod \mathcal{B}_b\} \in [0, [\mathcal{B}_b/2]], \quad (8)$$

Conversely, the farther apart the periodic positions of the time steps, the greater their negative correlation. Therefore, the set $\nabla_{m,n}^b = \{s | \delta_{m,s}^b > \delta_{m,n}^b\} \cup \{m\}$ includes time steps with larger periodic relative distances. Then, the negative modulation term is computed as the sum of the attention coefficients of these time steps and applied to the attention coefficients.

In this manner, PNA ensures that as the periodic relative distance between two time steps increases, their positive correlation coefficient decreases while their negative correlation coefficient increases. **We provide a mathematical explanation of this favorable property of the periodic-offset attention, which can be found in Appendix C.2.**

The row sums of the generated $\bar{\mathbf{A}}$ are not strictly equal to 1, satisfying:

$$\sum_{j=1}^{\mathcal{B}_b} \bar{\mathbf{A}}[i, j, n] = 1 - \mathbf{A} < 1, \quad \forall i = \{1, \dots, \mathcal{B}_b\}; n \in \{1, \dots, N_b\}. \quad (9)$$

To address it, when generating the final attention, we inject a residual path with residual strength Λ^h to stabilize the information flow. Please refer to Equation 11.

Period-aligned Attention $\tilde{\mathbf{A}}$ To capture dependencies among time steps that are phase-aligned with the periodicity, we employ a simplified self-attention mechanism. This mechanism shares the same query, key, and value vectors as the positive component of the period-offset attention, as these phase-aligned time steps exhibit strong correlations. This process can be expressed as:

$$\tilde{\mathbf{A}} = \text{Softmax}(\mu \mathbf{Q}_1 \times_2 \mathbf{K}_1^\top) \in \mathbb{R}^{P_b \times N_b \times N_b}. \quad (10)$$

where μ is the learnable coefficient.

Multi-head Attention Output We introduce a multi-head mechanism to enhance the model's expressiveness. The final computation is as follows:

$$\text{Multi-Head}(\mathbf{Z}^{(b)}) = \text{Concat}(\text{head}^1, \text{head}^2, \dots, \text{head}^H) \mathbf{W}_O, \quad (11)$$

$$\text{head}^h = \gamma_h \text{Tanh}[\alpha_h(\text{PNA}(\mathbf{Z}^{(b)}) + \Lambda^h \odot \mathbf{Z}^{(b)})] + \beta_h, \quad (12)$$

where H is the number of heads, $\mathbf{W}_O \in \mathbb{R}^{d \times d}$ is learnable parameters, and $\alpha_h \in \mathbb{R}$, $\gamma_h, \beta_h \in \mathbb{R}^{2 \times d_h}$ are learnable normalization parameters. \odot means Hadamard Product. We use Dynamic Tanh (Zhu et al., 2025) to eliminate any remaining numerical instability.

A special case involves the variates in the zero bucket, which lack periodicity. For these variates, the folding operation is skipped. When computing period-aligned attention and period-offset attention, absolute distance is used instead of periodic distance. The details are represented in **Appendix B**.

3.3 BUCKET-WISE FORECASTING

Flatten&Align Let us define the output of the b -th bucket after PNA learning as $\bar{\mathbf{Z}}^{(b)} \in \mathbb{R}^{P_b \times N_b \times d}$. Then, we flatten its first two dimensions, and if zero-padding has been applied (i.e., $(P_b * N_b) > L$), we truncate the padded time steps: $\mathbb{R}^{P_b \times N_b \times d} \rightarrow \mathbb{R}^{L \times d}$. Finally, we align the channels to the number of variates originally in the bucket:

$$\tilde{\mathbf{Z}}^{(b)} = \bar{\mathbf{Z}}^{(b)} \mathbf{W}_i^{(b)} + \mathbf{b}^{(b)} \in \mathbb{R}^{|\mathcal{B}_b| \times L} \quad (13)$$

where $\mathbf{W}^{(b)} \in \mathbb{R}^{d \times |\mathcal{B}_b|}$ and $\mathbf{b}^{(b)} \in \mathbb{R}^{|\mathcal{B}_b|}$ are learnable parameters. $|\mathcal{B}_b|$ denotes the number of variates in the b -th bucket.

Frequency-based Multi-period Prediction To predict the future value of the c -th variate among C variates, we first determine which buckets contain this variate based on its K_c period lengths, and then extract the corresponding bucket representations for prediction, which can be expressed as:

$$\hat{\mathbf{Y}}_c = \sum_{b=1}^{K_c} \alpha_c^{(b)} \tilde{\mathbf{Z}}_c^{(b)} \in \mathbb{R}^L, \quad \text{s.t. } c \in \{\mathcal{B}_i\}_{i=1}^N, \quad (b = 1, \dots, N) \quad (14)$$

$$\alpha_c^{(b)} = \text{Avg}(\text{Softmax}(|\beta_c^{(b)}|)), \quad \beta_c^{(b)} = \text{Extract}(\text{FFT}(\mathbf{X}_c^{(b)})) \quad (15)$$

where $\beta_c^{(b)}$ denotes the spectral magnitude extracted in the frequency domain for period length b . After normalization, $\beta_c^{(b)}$ is converted into the weight $\alpha_c^{(b)}$, which is used to perform a weighted fusion of the corresponding bucket representation $\tilde{\mathbf{Z}}_c^{(b)}$, yielding the final prediction. Finally, we can generate the prediction for C variates $\hat{\mathbf{Y}} \in \mathbb{R}^{C \times L}$.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Protocol Settings All experiments are performed on an NVIDIA A100 GPU with 80 GB of memory using PyTorch. For a comprehensive evaluation, we disable the “Drop Last” batch sampling procedure (Li et al., 2024; Qiu et al., 2024a). Optimization is performed with Adam (Kingma & Ba, 2014). Mean squared error (MSE) and mean absolute error (MAE) are used for evaluation. Considering that different models vary in their sensitivity to input history, we treat the look back length T as a tunable hyperparameter and report the best performance, and also comparisons using a fixed input length in Section 4.3. The details of hyperparameters are summarized in Table 5.

Table 1: Statistics of used datasets.

| Datasets | NN5 | Exchange | FRED-MD | ETTh | ETTm | AQShunyi | AQWan | ILI | CzeLan | ZafNoo | NASDAQ | NYSE |
|--------------|---------|----------|----------|-------------|-------------|-------------|-------------|--------|---------|---------|--------|-------|
| # Samples | 791 | 7,588 | 728 | 14,440 | 57,600 | 35,064 | 35,064 | 966 | 19,934 | 19,225 | 1,244 | 1,243 |
| # Frequency | 1 day | 1 day | 1 month | 1 hour | 15 mins | 1 hour | 1 hour | 1 week | 30 mins | 30 mins | 1 day | 1 day |
| Split Ration | 7:1:2 | 7:1:2 | 7:1:2 | 6:2:2 | 6:2:2 | 6:2:2 | 6:2:2 | 7:1:2 | 7:1:2 | 7:1:2 | 7:1:2 | 7:1:2 |
| Domain | Banking | Economic | Economic | Electricity | Electricity | Environment | Environment | Health | Nature | Nature | Stock | Stock |

Datasets We evaluate our method on **15** datasets: 14 real-world benchmarks (NN5, Exchange, FRED-MD, ETTh1, ETTh2, ETTm1, ETTm2, AQShunyi, AQWan, ILI, CzeLan, ZafNoo, NASDAQ, NYSE) and one synthetic dataset. Synthetic dataset is constructed by concatenating sequences from ETTm1 (period 96) and ETTh1 (period 24) along the time axis to simulate varying periodicities (see Appendix D.3). To account for varying dataset sizes, we adopt two evaluation regimes: for datasets with fewer than 5,000 samples, the tunable range of look back length is $T \in \{36, 104\}$ and forecasting horizons are $L \in \{24, 36, 48, 60\}$; for larger datasets, the tunable range of look back length is $T \in \{96, 336, 512\}$ and forecasting horizons are $L \in \{96, 192, 336, 720\}$. Dataset statistics are summarized in Table 1.

Baselines We compare our model with **18** advanced multivariate time series forecasting baselines including TimeKAN (Huang et al., 2025), xPatch (Stitsyuk & Choi, 2025), Amplifier (Fei et al., 2025), CycleNet (Lin et al., 2024a), TimeMixer (Wang et al., 2024c), SparseTSF (Lin et al., 2024b), iTransformer (Liu et al., 2023), Pathformer (Chen et al., 2024), PDF (Dai et al., 2024), FITS (Xu et al., 2023), PatchTST (Nie et al., 2022), Crossformer (Zhang & Yan, 2023), NLinear (Zeng et al., 2023), TimesNet (Wu et al., 2022a), FEDformer (Zhou et al., 2022b), Triformer (Cirstea et al., 2022), FiLM (Zhou et al., 2022a) and Non-stationary Transformer (Liu et al., 2022b).

Table 2: Multivariate time series forecasting performance comparison. We report average MSE and MAE. Best results are **bold**, second-best are underlined.

| Methods | Ours 2025 | | TimeKan 2025 | | xPatch 2025 | | Amplifier 2025 | | CycleNet 2024 | | TimeMixer 2024 | | SparseTSF 2024 | | iTTransformer 2024 | | Pathformer 2024 | | PDF 2024 | | |
|----------|-----------|---------------|--------------|---------------|--------------|---------------|----------------|--------------|---------------|--------------|----------------|--------------|----------------|--------------|--------------------|--------------|-----------------|---------|--------------|--------------|--------------|
| Metrics | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | |
| ETTh | 96 | 0.316 | <u>0.356</u> | 0.324 | 0.37 | 0.327 | <u>0.364</u> | 0.334 | 0.377 | 0.327 | 0.369 | 0.327 | 0.376 | 0.322 | <u>0.364</u> | 0.342 | 0.377 | 0.326 | <u>0.364</u> | <u>0.318</u> | 0.366 |
| | 192 | 0.360 | <u>0.394</u> | <u>0.364</u> | 0.399 | 0.384 | 0.403 | 0.381 | 0.407 | 0.374 | 0.400 | 0.381 | 0.409 | 0.371 | <u>0.395</u> | 0.398 | 0.422 | 0.377 | 0.398 | 0.366 | 0.398 |
| | 336 | 0.384 | <u>0.412</u> | 0.394 | 0.421 | 0.420 | 0.427 | 0.413 | 0.435 | 0.401 | 0.422 | 0.402 | 0.432 | <u>0.387</u> | <u>0.408</u> | 0.419 | 0.439 | 0.408 | 0.421 | 0.396 | 0.421 |
| | 720 | <u>0.412</u> | <u>0.441</u> | 0.425 | 0.453 | 0.456 | 0.460 | 0.431 | 0.456 | 0.438 | 0.458 | 0.444 | 0.46 | <u>0.402</u> | <u>0.431</u> | 0.460 | 0.466 | 0.444 | 0.459 | 0.427 | 0.448 |
| ETTm | 96 | 0.224 | <u>0.288</u> | <u>0.225</u> | 0.3 | <u>0.225</u> | <u>0.291</u> | 0.233 | 0.303 | 0.229 | 0.298 | 0.229 | 0.301 | 0.241 | 0.306 | 0.238 | 0.31 | 0.227 | 0.293 | <u>0.225</u> | 0.296 |
| | 192 | <u>0.272</u> | <u>0.323</u> | 0.285 | 0.334 | 0.276 | <u>0.325</u> | 0.277 | 0.333 | 0.274 | 0.327 | 0.280 | 0.335 | 0.284 | 0.334 | 0.292 | 0.346 | 0.278 | <u>0.323</u> | <u>0.270</u> | 0.327 |
| | 336 | <u>0.311</u> | <u>0.349</u> | 0.316 | 0.358 | 0.320 | <u>0.352</u> | 0.321 | 0.359 | 0.319 | 0.354 | 0.323 | 0.359 | 0.323 | 0.357 | 0.328 | 0.367 | 0.322 | <u>0.352</u> | <u>0.312</u> | 0.357 |
| | 720 | 0.388 | <u>0.392</u> | <u>0.380</u> | 0.402 | 0.401 | 0.405 | 0.393 | 0.404 | 0.390 | 0.398 | 0.393 | 0.402 | 0.398 | 0.402 | 0.402 | 0.412 | 0.395 | <u>0.397</u> | <u>0.379</u> | 0.399 |
| NN5 | 24 | 0.681 | <u>0.551</u> | 0.769 | 0.605 | 0.853 | 0.663 | 1.400 | 0.910 | 0.754 | 0.588 | <u>0.723</u> | 0.574 | 1.679 | 1.026 | 0.727 | <u>0.568</u> | 0.769 | 0.602 | 0.736 | 0.579 |
| | 36 | 0.640 | <u>0.538</u> | 0.7 | 0.585 | 0.785 | 0.634 | 1.268 | 0.862 | 0.687 | 0.571 | <u>0.657</u> | 0.557 | 1.789 | 1.066 | 0.664 | <u>0.552</u> | 0.701 | 0.583 | 0.676 | 0.567 |
| | 60 | 0.613 | <u>0.532</u> | 0.659 | 0.565 | 0.716 | 0.601 | 1.202 | 0.842 | 0.656 | 0.564 | 0.630 | 0.55 | 1.982 | 0.907 | 0.633 | 0.543 | 0.668 | 0.573 | <u>0.616</u> | <u>0.538</u> |
| | 720 | 0.603 | <u>0.531</u> | 0.651 | 0.568 | 0.673 | 0.579 | 1.237 | 0.861 | 0.642 | 0.568 | 0.612 | 0.543 | 1.734 | 1.046 | 0.615 | 0.537 | 0.655 | 0.570 | <u>0.599</u> | <u>0.533</u> |
| Exchange | 96 | 0.083 | <u>0.203</u> | 0.089 | 0.208 | 0.087 | 0.204 | <u>0.084</u> | <u>0.203</u> | 0.093 | 0.216 | <u>0.084</u> | 0.207 | 0.116 | 0.243 | 0.086 | 0.205 | 0.088 | 0.208 | 0.083 | <u>0.200</u> |
| | 192 | <u>0.169</u> | <u>0.295</u> | 0.183 | 0.304 | 0.188 | 0.306 | 0.179 | 0.300 | 0.184 | 0.307 | 0.178 | 0.300 | 0.203 | 0.325 | 0.177 | 0.299 | 0.183 | 0.304 | <u>0.172</u> | <u>0.294</u> |
| | 336 | 0.316 | <u>0.406</u> | 0.338 | 0.421 | 0.332 | 0.422 | 0.337 | 0.419 | 0.329 | 0.416 | 0.376 | 0.451 | 0.360 | 0.437 | 0.331 | 0.417 | 0.354 | 0.429 | <u>0.323</u> | <u>0.411</u> |
| | 720 | <u>0.773</u> | <u>0.669</u> | 0.918 | 0.719 | 0.899 | 0.718 | 0.929 | 0.719 | 0.854 | 0.698 | 0.884 | 0.707 | 0.882 | 0.714 | 0.846 | 0.693 | 0.909 | 0.716 | <u>0.820</u> | <u>0.682</u> |
| FRED-MD | 24 | 21.495 | <u>0.903</u> | <u>25.178</u> | <u>0.875</u> | 52.290 | 1.030 | 42.125 | 1.192 | 31.558 | 0.945 | 29.132 | 0.93 | 66.664 | 1.55 | 28.581 | 0.917 | 43.434 | 1.142 | 30.949 | 0.936 |
| | 36 | 38.363 | <u>1.161</u> | <u>45.811</u> | <u>1.194</u> | 52.339 | 1.262 | 70.564 | 1.525 | 56.633 | 1.285 | 58.770 | 1.33 | 96.051 | 1.841 | 54.221 | 1.276 | 70.370 | 1.441 | 56.870 | 1.276 |
| | 48 | 52.065 | <u>1.453</u> | 76.144 | 1.542 | <u>73.426</u> | <u>1.495</u> | 102.150 | 1.825 | 124.077 | 2.043 | 90.283 | 1.632 | 139.333 | 2.175 | 89.574 | 1.607 | 96.827 | 1.728 | 88.240 | 1.615 |
| | 720 | 76.292 | <u>1.679</u> | <u>96.139</u> | <u>1.736</u> | 103.023 | 1.778 | 145.867 | 2.162 | 172.758 | 2.411 | 126.438 | 1.865 | 190.400 | 2.526 | 130.061 | 1.947 | 134.397 | 2.023 | 131.531 | 1.884 |
| AQShuny1 | 96 | 0.665 | 0.468 | 0.651 | 0.478 | 0.663 | <u>0.469</u> | 0.667 | 0.489 | 0.653 | 0.484 | 0.654 | 0.483 | 0.696 | 0.509 | <u>0.650</u> | 0.479 | 0.667 | 0.472 | 0.647 | 0.481 |
| | 192 | 0.707 | 0.488 | <u>0.689</u> | 0.497 | 0.708 | 0.492 | 0.702 | 0.505 | 0.699 | 0.502 | 0.700 | 0.498 | 0.719 | 0.521 | 0.693 | 0.498 | 0.707 | <u>0.491</u> | <u>0.690</u> | 0.499 |
| | 336 | 0.737 | 0.502 | <u>0.712</u> | 0.512 | 0.741 | 0.506 | 0.718 | 0.511 | 0.719 | 0.514 | 0.715 | 0.510 | 0.728 | 0.525 | 0.713 | 0.510 | 0.732 | <u>0.503</u> | <u>0.708</u> | 0.512 |
| | 720 | 0.795 | <u>0.529</u> | <u>0.763</u> | 0.531 | 0.790 | 0.532 | 0.765 | 0.537 | 0.774 | 0.540 | <u>0.756</u> | 0.534 | 0.784 | 0.552 | 0.766 | 0.537 | 0.783 | <u>0.515</u> | 0.765 | 0.537 |
| AQWan | 96 | 0.746 | <u>0.452</u> | 0.749 | 0.467 | 0.752 | <u>0.455</u> | 0.762 | 0.478 | 0.758 | 0.473 | <u>0.744</u> | 0.468 | 0.796 | 0.497 | 0.747 | 0.470 | 0.761 | 0.458 | 0.747 | 0.469 |
| | 192 | 0.808 | <u>0.475</u> | <u>0.791</u> | 0.488 | 0.803 | 0.479 | 0.806 | 0.493 | 0.804 | 0.492 | 0.804 | 0.488 | 0.820 | 0.509 | <u>0.787</u> | 0.486 | 0.801 | <u>0.478</u> | 0.794 | 0.489 |
| | 336 | 0.824 | 0.487 | <u>0.815</u> | 0.497 | 0.833 | 0.490 | 0.821 | 0.500 | 0.826 | 0.503 | <u>0.813</u> | 0.500 | 0.837 | 0.514 | <u>0.814</u> | 0.497 | 0.821 | <u>0.488</u> | 0.817 | 0.500 |
| | 720 | 0.907 | <u>0.517</u> | 0.888 | 0.527 | 0.906 | 0.522 | 0.894 | 0.530 | 0.899 | 0.522 | 0.907 | 0.543 | 0.888 | 0.522 | 0.888 | 0.529 | 0.888 | <u>0.506</u> | <u>0.885</u> | 0.527 |
| ILI | 24 | 1.318 | <u>0.705</u> | 2.176 | 0.928 | 2.320 | 1.201 | 2.950 | 1.220 | 2.195 | 1.023 | 1.804 | <u>0.820</u> | 4.917 | 1.660 | <u>1.783</u> | 0.846 | 2.086 | 0.922 | 1.801 | 0.874 |
| | 36 | 1.523 | <u>0.775</u> | 2.166 | 0.993 | 2.255 | 1.011 | 2.759 | 1.187 | 1.971 | 0.942 | 1.891 | 0.926 | 4.377 | 1.518 | 1.746 | <u>0.860</u> | 1.912 | 0.882 | <u>1.743</u> | 0.867 |
| | 48 | 4.137 | <u>0.775</u> | 2.011 | 0.928 | 2.221 | 1.004 | 2.821 | 1.199 | 1.950 | 0.951 | 1.752 | <u>0.866</u> | 4.412 | 1.412 | <u>2.116</u> | 0.898 | 1.985 | 0.905 | 1.843 | 0.926 |
| | 60 | 1.461 | <u>0.777</u> | 2.010 | 0.967 | 2.269 | 1.013 | 2.857 | 1.196 | 1.996 | 0.966 | <u>1.831</u> | 0.930 | 4.497 | 1.521 | 2.183 | 0.963 | 1.999 | 0.929 | 1.845 | <u>0.925</u> |
| Czclan | 96 | 0.439 | <u>0.383</u> | 0.427 | 0.398 | <u>0.433</u> | 0.396 | 0.443 | 0.404 | 0.406 | 0.409 | 0.441 | 0.396 | 0.510 | 0.452 | 0.439 | 0.408 | 0.435 | <u>0.391</u> | 0.438 | 0.401 |
| | 192 | 0.503 | 0.429 | <u>0.485</u> | 0.433 | 0.498 | 0.437 | <u>0.495</u> | 0.442 | 0.503 | 0.445 | 0.498 | 0.444 | 0.542 | 0.475 | 0.505 | 0.443 | 0.501 | <u>0.432</u> | 0.498 | 0.440 |
| | 336 | 0.547 | 0.456 | <u>0.533</u> | 0.464 | <u>0.541</u> | 0.463 | 0.545 | 0.469 | 0.543 | 0.467 | 0.543 | 0.466 | 0.586 | 0.493 | 0.555 | 0.473 | 0.551 | <u>0.461</u> | 0.545 | 0.463 |
| | 720 | 0.592 | 0.477 | 0.606 | 0.503 | 0.595 | <u>0.478</u> | 0.590 | 0.496 | <u>0.584</u> | 0.490 | 0.588 | 0.498 | 0.616 | 0.514 | 0.591 | 0.501 | 0.596 | 0.483 | <u>0.580</u> | 0.486 |
| ZalNoo | 24 | 0.416 | <u>0.473</u> | 0.541 | 0.517 | <u>0.503</u> | <u>0.500</u> | 0.753 | 0.630 | 0.631 | 0.570 | 0.720 | 0.612 | 1.363 | 0.876 | 0.570 | 0.540 | 0.572 | 0.568 | 0.611 | 0.541 |
| | 36 | 0.610 | <u>0.600</u> | 1.013 | 0.732 | 0.865 | <u>0.665</u> | 1.179 | 0.784 | 0.886 | 0.675 | 0.951 | 0.699 | 1.223 | 0.854 | <u>0.691</u> | 0.600 | 0.836 | 0.683 | 0.862 | 0.683 |
| | 48 | 0.949 | <u>0.731</u> | 1.202 | 0.779 | <u>1.128</u> | 0.794 | 1.425 | 0.881 | 1.321 | 0.854 | 1.214 | 0.795 | 1.310 | 0.884 | 1.188 | <u>0.773</u> | 1.227 | 0.798 | 1.235 | 0.803 |
| | 60 | 0.961 | <u>0.779</u> | 1.414 | 0.846 | 1.445 | 0.836 | 1.352 | 0.862 | 1.535 | 0.888 | 1.136 | <u>0.817</u> | <u>1.110</u> | 0.818 | 1.325 | 0.820 | 1.421 | 0.855 | 1.426 | 0.841 |
| NYSE | 24 | 0.161 | <u>0.251</u> | 0.224 | 0.304 | <u>0.210</u> | <u>0.296</u> | 0.280 | 0.362 | 0.237 | 0.314 | 0.253 | 0.369 | 0.512 | 0.497 | 0.225 | 0.302 | 1.102 | 0.814 | 0.261 | 0.331 |
| | 36 | 0.298 | <u>0.352</u> | <u>0.322</u> | <u>0.360</u> | 0.366 | 0.387 | 0.640 | 0.599 | 0.395 | 0.411 | 0.638 | 0.621 | 0.669 | 0.566 | 0.392 | 0.409 | 1.598 | 1.002 | 0.385 | 0.406 |
| | 48 | 0.445 | <u>0.427</u> | <u>0.475</u> | <u>0.447</u> | 0.693 | 0.558 | 0.761 | 0.599 | 0.555 | 0.484 | 0.747 | 0.663 | 0.937 | 0.691 | 0.529 | 0.480 | 1.910 | 1.116 | 0.689 | 0.562 |
| | 60 | 0.619 | <u>0.525</u> | <u>0.660</u> | <u>0.534</u> | 0.886 | 0.646 | 0.950 | 0.685 | 0.759 | 0.584 | 1.027 | 0.807 | 1.219 | 0.814 | 0.687 | 0.557 | 2.123 | 1.188 | 0.810 | 0.616 |
| # Top1 | 71 | 6 | 2 | 0 | 0 | 4 | 0 | 4 | 3 | 2 | 0 | 4 | 3 | 2 | 3 | 9 | 9 | 9 | 9 | | |
| # Top2 | | | | | | | | | | | | | | | | | | | | | |

4.3 COMPARISON WITH FIXED INPUT LENGTH SETTINGS

The handling of input windows in time series forecasting tasks has always been flexible. In our experiments, we treated the historical window length as a hyperparameter, allowing models to flexibly select the length based on their own characteristics. We believe that this approach provides a more comprehensive evaluation. In contrast, maintaining a fixed input window length avoids fairness concerns arising from unequal usage of historical information across models. Here, we do not assess the advantages or disadvantages of these two settings. Instead, we further supplement a comparison under the fixed look back window setting, with the results presented below. As shown in Table 3, Our model still achieves the best performance.

Table 3: Comparison in the same input setting. ETTh1 is 96 to 96 and the others are 24 to 36.

| Dataset | Ours | TimeKAN | xPatch | Amplifer | CycleNet | TimeMixer | SparseTSF | iTransformer | Pathformer | PDF |
|---------|------------------|---------|--------|----------|----------|-----------|-----------|--------------|------------|-------|
| ETTh1 | MSE 0.375 | 0.378 | 0.379 | 0.443 | 0.400 | 0.381 | 0.400 | 0.386 | 0.386 | 0.378 |
| | MAE 0.389 | 0.402 | 0.401 | 0.428 | 0.415 | 0.397 | 0.403 | 0.405 | 0.392 | 0.396 |
| ILI | MSE 2.062 | 2.857 | 2.704 | 3.455 | 2.955 | 2.015 | 5.498 | 2.112 | 2.906 | 2.168 |
| | MAE 0.876 | 1.053 | 1.064 | 1.296 | 1.195 | 0.863 | 1.777 | 0.885 | 1.154 | 0.899 |
| NASDAQ | MSE 0.606 | 0.610 | 0.616 | 0.752 | 0.967 | 0.619 | 1.363 | 0.615 | 0.643 | 0.624 |
| | MAE 0.534 | 0.542 | 0.539 | 0.629 | 0.729 | 0.546 | 0.876 | 0.546 | 0.578 | 0.550 |

4.4 ABLATION STUDIES

We design several ablation variants of PHAT to validate the contribution of its key components. Specifically, “**w/o POA**” removes the period-offset attention branch from PNA, while “**w/o PAA**” eliminates the period-aligned attention branch. “**w/o Attn**” cancels the self-attention mechanism entirely, reducing it to a single feed-forward net. “**w/o Bucket**” disables the period-based variate grouping, treating each variate independently rather than within its detected period bucket.

As shown in Figure 3 (a) (where taller bars indicate lower MSE for intuitive display), the “**w/o Bucket**” variant exhibits significantly larger prediction errors, indicating that interactions between variates with different periodic characteristics increase the difficulty of capturing complex periodic patterns. The poor performance of the “**w/o PAA**” variant highlights the importance of cross-period synchronization signals at the same phase, especially on datasets with weaker periodicity, such as NN5 and CzeLan. “**w/o POA**” variant achieves extremely poor predictive performance. We further conducted ablation experiments on the combination of periodic offset attention, and the results are shown in Table 10. We found that modeling dependencies between time steps from both positive and negative perspectives is beneficial. Additionally, the modulators configured for each perspective help generate attention coefficients that align more closely with periodic trends. Detailed comparison of combination ablation study on additional datasets is in Appendix D.7.

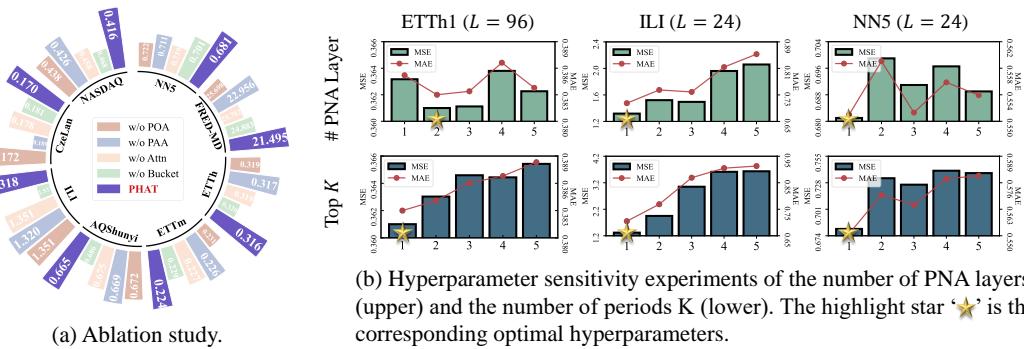


Figure 3: Ablation study and hyperparameters sensitivity experiments of phats

4.5 HYPERPARAMETER SENSITIVITY EXPERIMENTS

We further investigate the sensitivity of two key hyperparameters: the number of PNA layers and the number of cycle lengths K selected for each variate, as shown in Figure 3. Experimental results

(Figure 3 (b-upper)) show that our model achieves strong performance with just one or two layers, thanks to the periodic bucket structure and the PNA mechanism, which effectively capture periodic patterns. In contrast, adding more Transformer layers brings little gain and may even hurt performance due to overfitting. As for the hyperparameter K , Figure 3(b-lower) shows that on datasets with simple periodicity, a single dominant period ($K=1$) is often sufficient. Using more periods tends to introduce noise and degrade prediction accuracy.

4.6 COMPLEXITY ANALYSIS

We evaluate PHAT against advanced baselines on two datasets. As summarized in Table 4, PHAT achieves the best forecasting accuracy while substantially reducing model complexity. Compared with Transformer-based methods (e.g., PatchTST, PDF) and recently advanced methods (e.g. TimeKAN, xPatch), PHAT cuts parameter counts by more than an order of magnitude and reduces MACs and FLOPs by over 98%. Even compared to lightweight architectures like TimeKAN, PHAT retains a Transformer-style backbone while reducing the parameter count by up to **91.27%** with competitive inference latency. Because the computational cost of our attention scales with the square of the detected period length rather than with the full input sequence, so long sequences remain inexpensive to process. In practice, PHAT also typically requires only a single layer to capture temporal dependencies effectively.

Table 4: Computational Complexity Comparison. # Para: All learnable parameters requiring gradient descent. MACs: multiply–accumulate operations. # FLOPs: floating point operations. K: Kilo (10^3). M: Million (10^6). B: Billion (10^9). Inf: actual inference latency (s).

| Models | MAE | # Para | # MACs | # FLOPs | Inf lat | Models | MAE | # Para | # MACs | # FLOPs | Inf lat | | | | | |
|--------------------|-------------|--------|---------|---------|---------|--------------|---------------|--------------|--------------|----------------|-------------|--------------|---------------|--------------|---------------|----------------|
| ETTm1 ($L = 96$) | FEDformer | 0.463 | 3.4 M | 1.7 B | 1.3 B | 20.044 s | FEDformer | 0.488 | 16.8 M | 111.9 B | 95.7 B | 0.480 s | | | | |
| | TimesNet | 0.378 | 2.4 M | 72.2 B | 72.2 B | 3.348 s | TimesNet | 0.495 | 665.9 K | 77.0 B | 76.9 B | 1.022 s | | | | |
| | Crossformer | 0.367 | 2.1 M | 13.5 B | 14.3 B | 1.806 s | Crossformer | 0.514 | 11.6 M | 60.4 B | 61.5 B | 4.031 s | | | | |
| | PDF | 0.340 | 2.0 M | 4.4 B | 4.6 B | 4.234 s | PDF | 0.462 | 534.2 K | 1.6 B | 1.6 B | 0.351 s | | | | |
| | TimeMixer | 0.345 | 397.3 K | 3.2 B | 3.1 B | 3.967 s | TimeMixer | 0.484 | 2.3 M | 34.0 B | 34.0 B | 0.891 s | | | | |
| | TimeKAN | 0.346 | 52.8 K | 205.4 M | 455.0 M | 6.542 s | TimeKAN | 0.463 | 374.4 K | 1.4 B | 1.7 B | 1.934 s | | | | |
| PHAT | | | | | | 0.330 | 33.4 K | 2.9 M | 7.0 M | 1.671 s | PHAT | 0.458 | 32.7 K | 3.7 M | 16.5 M | 0.462 s |

4.7 VISUALIZATION OF PERIODIC POSITIVE-NEGATIVE COMPONENTS IN DATA, FEATURE, AND ATTENTION LEVELS

In this subsection, we further demonstrate that positive-negative periodic correlation indeed exist at all multiple levels. We provide comprehensive experimental analysis at three levels—sequence, feature, and attention—to support our motivation on ETTh1 dataset of variates with daily period and NNS with weekly period.

① Sequence Level. We visualize the autocorrelation function (ACF) of the original time series data \mathbf{X} , which clearly demonstrates the presence of both positive and negative correlations between time steps.

② Feature Level. We computed the autocorrelation function for each time step in the bucket representation $\mathbf{Z}^{(b)}$. As shown in Figure 4 (a), across each hidden feature dimension, the resulting high-dimensional features continue to exhibit negative correlation, indicating that the positive and negative correlations present in the original signals persist after the complex-valued linear projection.

③ Attention Level. We aim to demonstrate two points: (1) attention coefficients do contain negative components; and (2) the softmax operation tends to smooth out or suppress these negative components. As shown in Figure 4 (b) and (c), we visualized PHAT model’s attention logits before and after the softmax, and compared them with a control variant in which our positive–negative attention is replaced by standard self-attention. The results show that before softmax, both models’ logits exhibit clear positive and negative components; after softmax, the negative components of the standard self-attention nearly vanish, whereas our attention preserves negative correlations.

In contrast, our PNA mechanism preserves these negative correlations by separately modulating the positive and negative pathways, which is beneficial for capturing the complete periodic structure.

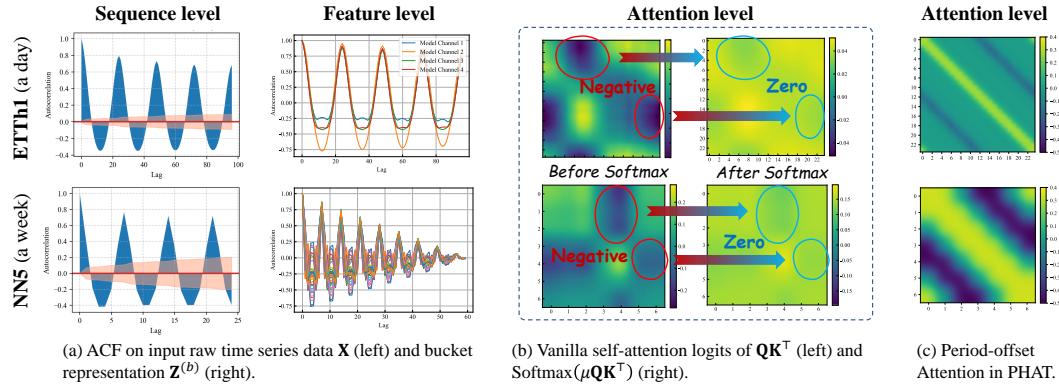


Figure 4: Multi-level visualizations on raw-sequence level, feature level and attention level.

4.8 ATTENTION VISUALIZATION

As shown in Figure 5, we illustrate the autocorrelation coefficients between different time steps within a time series (left side) for various datasets, while the right side visualizes the corresponding cycle-shifted attention weights generated by our method. The attention matrix exhibits a clear periodic structure: bright bands represent positive attention, dark bands represent negative attention, and green areas indicate zero attention. We can observe that the trend between our generated attention matrix and the periodic correlation of time series data is consistent, demonstrating the ability to capture underlying periodic structures. For time series with weaker negative periodicity (such as ZafNoo), there are fewer negative dark regions in the attention. For non-periodic time series (such as those from the ETTh2 and NASDAQ datasets), our cycle-shifted attention mechanism naturally degrades to focus on each individual time step. Overall, our method effectively captures the periodic dynamics of time series.

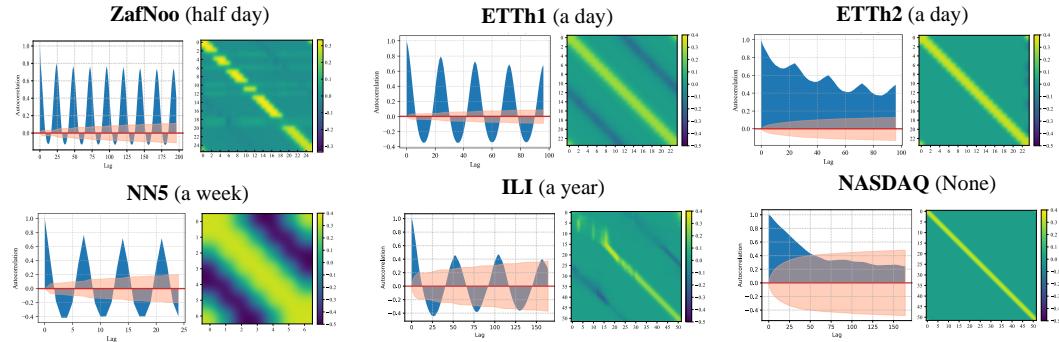


Figure 5: Autocorrelation function (left) and periodic-offset attention weights (right). Highlighted regions denote positive components; dark regions denote negative components.

5 CONCLUSION

In this paper, we identify the challenge of periodic heterogeneity, where the periodic characteristics of variates within a dataset differ. To address this, we propose a Periodic Heterogeneity-Aware Transformer (PHAT). PHAT employs a periodic bucket structure to manage multivariate time series with heterogeneous periodicity. By introducing the periodic bucket mechanism and the novel Positive-Negative Attention (PNA) module, PHAT is able to seamlessly and accurately model heterogeneous periodic patterns while capturing both positive and negative phase dependencies. Extensive experiments conducted on 14 real-world datasets demonstrate that PHAT achieves state-of-the-art performance across most evaluation metrics.

ETHICS STATEMENT

Our work only focuses on the time series forecasting problem, so there is no potential ethical risk.

REPRODUCIBILITY STATEMENT

In the main text, we have strictly formalized the model architecture with equations. All the implementation details are included in the Appendix, including dataset descriptions, metrics, model, and experiment configurations. The code is provided at an anonymous link.

ACKNOWLEDGMENT

This paper is partially supported by the National Natural Science Foundation of China (No.12227901) and the Natural Science Foundation of Jiangsu Province (BK20250482). The AI-driven experiments, simulations and model training were performed on the robotic AI-Scientist platform of Chinese Academy of Science.

REFERENCES

Hossein Arsham and Miodrag Lovric. Bartlett's test. *International encyclopedia of statistical science*, 2(2):20–23, 2011.

David Chandler. Introduction to modern statistical. *Mechanics*. Oxford University Press, Oxford, UK, 5(449):11, 1987.

Peng Chen, Yingying Zhang, Yunyao Cheng, Yang Shu, Yihang Wang, Qingsong Wen, Bin Yang, and Chenjuan Guo. Pathformer: Multi-scale transformers with adaptive pathways for time series forecasting. *arXiv preprint arXiv:2402.05956*, 2024.

Panayiotis Christou, Shichu Chen, Xupeng Chen, and Parijat Dube. Test time learning for time series forecasting. *arXiv preprint arXiv:2409.14012*, 2024.

Razvan-Gabriel Cirstea, Chenjuan Guo, Bin Yang, Tung Kieu, Xuanyi Dong, and Shirui Pan. Triformer: Triangular, variable-specific attentions for long sequence multivariate time series forecasting–full version. *arXiv preprint arXiv:2204.13767*, 2022.

Tao Dai, Beiliang Wu, Peiyuan Liu, Naiqi Li, Jigang Bao, Yong Jiang, and Shu-Tao Xia. Periodicity decoupling framework for long-term series forecasting. In *The Twelfth International Conference on Learning Representations*, 2024.

Serge Dégerine and Sophie Lambert-Lacroix. Characterization of the partial autocorrelation function of nonstationary time series. *Journal of Multivariate Analysis*, 87(1):46–59, 2003.

Pierre Duhamel and Martin Vetterli. Fast fourier transforms: a tutorial review and a state of the art. *Signal processing*, 19(4):259–299, 1990.

Yuchen Fang, Yanjun Qin, Haiyong Luo, Fang Zhao, and Kai Zheng. Stwave++: A multi-scale efficient spectral graph attention network with long-term trends for disentangled traffic flow forecasting. *IEEE Transactions on Knowledge and Data Engineering*, 36(6):2671–2685, 2023.

Jingru Fei, Kun Yi, Wei Fan, Qi Zhang, and Zhendong Niu. Amplifier: Bringing attention to neglected low-energy components in time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 11645–11653, 2025.

Pradeep Hewage, Ardhendu Behera, Marcello Trovati, Ella Pereira, Morteza Ghahremani, Francesco Palmieri, and Yonghuai Liu. Temporal convolutional neural (tcn) network for an effective weather forecasting using time-series data from the local weather station. *Soft Computing*, 24(21):16453–16482, 2020.

Yifan Hu, Peiyuan Liu, Peng Zhu, Dawei Cheng, and Tao Dai. Adaptive multi-scale decomposition framework for time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 17359–17367, 2025.

Qihe Huang, Lei Shen, Ruixin Zhang, Jiahuan Cheng, Shouhong Ding, Zhengyang Zhou, and Yang Wang. Hdmixer: Hierarchical dependency with extendable patch for multivariate time series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pp. 12608–12616, 2024a.

Qihe Huang, Lei Shen, Ruixin Zhang, Shouhong Ding, Binwu Wang, Zhengyang Zhou, and Yang Wang. Crossgnn: Confronting noisy multivariate time series via cross interaction refinement. *Proc. of NeurIPS*, 2024b.

Songtao Huang, Zhen Zhao, Can Li, and Lei Bai. Timekan: Kan-based frequency decomposition learning architecture for long-term time series forecasting. *arXiv preprint arXiv:2502.06910*, 2025.

Taesung Kim, Jinhee Kim, Yunwon Tae, Cheonbok Park, Jang-Ho Choi, and Jaegul Choo. Reversible instance normalization for accurate time-series forecasting against distribution shift. In *International conference on learning representations*, 2021.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Shiyang Li, Xiaoyong Jin, Yao Xuan, Xiyou Zhou, Wenhui Chen, Yu-Xiang Wang, and Xifeng Yan. Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. *Advances in neural information processing systems*, 32, 2019.

Zhe Li, Xiangfei Qiu, Peng Chen, Yihang Wang, Hanyin Cheng, Yang Shu, Jilin Hu, Chenjuan Guo, Aoying Zhou, Qingsong Wen, et al. Foundts: Comprehensive and unified benchmarking of foundation models for time series forecasting. *arXiv preprint arXiv:2410.11802*, 2024.

Yuxuan Liang, Haomin Wen, Yuqi Nie, Yushan Jiang, Ming Jin, Dongjin Song, Shirui Pan, and Qingsong Wen. Foundation models for time series analysis: A tutorial and survey. In *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*, pp. 6555–6565, 2024.

Shengsheng Lin, Weiwei Lin, Xinyi Hu, Wentai Wu, Ruichao Mo, and Haocheng Zhong. Cyclenet: Enhancing time series forecasting through modeling periodic patterns. *Advances in Neural Information Processing Systems*, 37:106315–106345, 2024a.

Shengsheng Lin, Weiwei Lin, Wentai Wu, Haojun Chen, and Junjie Yang. Sparsesf: Modeling long-term time series forecasting with 1k parameters. *arXiv preprint arXiv:2405.00946*, 2024b.

Shengsheng Lin, Weiwei Lin, Wentai Wu, Haojun Chen, and CL Philip Chen. Sparsesf: Lightweight and robust time series forecasting via sparse modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

Shizhan Liu, Hang Yu, Cong Liao, Jianguo Li, Weiyao Lin, Alex X Liu, and Schahram Dustdar. Pyraformer: Low-complexity pyramidal attention for long-range time series modeling and forecasting. In # PLACEHOLDER_PARENT_METADATA_VALUE#, 2022a.

Yong Liu, Haixu Wu, Jianmin Wang, and Mingsheng Long. Non-stationary transformers: Exploring the stationarity in time series forecasting. *Advances in neural information processing systems*, 35: 9881–9893, 2022b.

Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. itransformer: Inverted transformers are effective for time series forecasting. *arXiv preprint arXiv:2310.06625*, 2023.

Ziqi Liu. Freqmoe: Enhancing time series forecasting through frequency decomposition mixture of experts. *arXiv preprint arXiv:2501.15125*, 2025.

Donghao Luo and Xue Wang. Moderntcn: A modern pure convolution structure for general time series analysis. In *The twelfth international conference on learning representations*, pp. 1–43, 2024.

Jiaming Ma, Zhiqing Cui, Binwu Wang, Pengkun Wang, Zhengyang Zhou, Zhe Zhao, and Yang Wang. Causal learning meet covariates: Empowering lightweight and effective nationwide air quality forecasting. In *International Joint Conference on Artificial Intelligence*, 2025a.

Jiaming Ma, Binwu Wang, Qihe Huang, Guanjun Wang, Pengkun Wang, Zhengyang Zhou, and Yang Wang. Mofo: Empowering long-term time series forecasting with periodic pattern modeling. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025b.

Jiaming Ma, Binwu Wang, Guanjun Wang, Kuo Yang, Zhengyang Zhou, Pengkun Wang, Xu Wang, and Yang Wang. Less but more: Linear adaptive graph learning empowering spatiotemporal forecasting. *Advances in Neural Information Processing Systems*, 2025c.

Jiaming Ma, Binwu Wang, Pengkun Wang, Zhengyang Zhou, Xu Wang, and Yang Wang. Bist: A lightweight and efficient bi-directional model for spatiotemporal prediction. *Proceedings of the VLDB Endowment*, 18(6):1663–1676, 2025d.

Jiaming Ma, Binwu Wang, Pengkun Wang, Zhengyang Zhou, Xu Wang, and Yang Wang. Robust spatio-temporal centralized interaction for ood learning. In *Forty-second International Conference on Machine Learning*, 2025e.

Jiaming Ma, Binwu Wang, Pengkun Wang, Zhengyang Zhou, Yudong Zhang, Xu Wang, and Yang Wang. Mobimixer: A multi-scale spatiotemporal mixing model for mobile traffic prediction. *IEEE Transactions on Mobile Computing*, 2025f.

Jiaming Ma, Guanjun Wang, Sheng Huang, Kuo Yang, Binwu Wang, Pengkun Wang, and Yang Wang. Spatiotemporal causal decoupling model for air quality forecasting. *arXiv preprint arXiv:2505.20119*, 2025g.

Jiaming Ma, Siyuan Mu, Ruilin Tang, Haofeng Ma, Qihe Huang, Zhengyang Zhou, Pengkun Wang, Binwu Wang, and Yang Wang. To see far, look close: Evolutionary forecasting for long-term time series. *arXiv preprint arXiv:2601.23114*, 2026a.

Jiaming Ma, Binwu Wang, Pengkun Wang, Xu Wang, Zhengyang Zhou, and Yang Wang. A general relearner: Empowering spatiotemporal prediction by re-learning input-label residual. *arXiv preprint arXiv:2602.02563*, 2026b.

Kasun Mendis, Manjusri Wickramasinghe, and Pasindu Marasinghe. Multivariate time series forecasting: A review. In *Proceedings of the 2024 2nd Asia Conference on Computer Vision, Image Processing and Pattern Recognition*, pp. 1–9, 2024.

Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*, 2022.

Xihao Piao, Zheng Chen, Taichi Murayama, Yasuko Matsubara, and Yasushi Sakurai. Fredformer: Frequency debiased transformer for time series forecasting. In *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*, pp. 2400–2410, 2024.

Xiangfei Qiu, Jilin Hu, Lekui Zhou, Xingjian Wu, Junyang Du, Buang Zhang, Chenjuan Guo, Aoying Zhou, Christian S Jensen, Zhenli Sheng, et al. Tfb: Towards comprehensive and fair benchmarking of time series forecasting methods. *arXiv preprint arXiv:2403.20150*, 2024a.

Xiangfei Qiu, Xingjian Wu, Yan Lin, Chenjuan Guo, Jilin Hu, and Bin Yang. Duet: Dual clustering enhanced multivariate time series forecasting. *arXiv preprint arXiv:2412.10859*, 2024b.

Xiangfei Qiu, Xingjian Wu, Hanyin Cheng, Xvyuan Liu, Chenjuan Guo, Jilin Hu, and Bin Yang. Dbloss: Decomposition-based loss function for time series forecasting. *arXiv preprint arXiv:2510.23672*, 2025.

Lu Ren, Lan Du, David B Dunson, et al. Logistic stick-breaking process. *Journal of Machine Learning Research*, 12(1), 2011.

Sima Siami-Namini, Neda Tavakoli, and Akbar Siami Namin. The performance of lstm and bilstm in forecasting time series. In *2019 IEEE International conference on big data (Big Data)*, pp. 3285–3292. IEEE, 2019.

Artyom Stityuk and Jaesik Choi. xpatch: Dual-stream time series forecasting with exponential seasonal-trend decomposition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 20601–20609, 2025.

Peiwang Tang and Weitai Zhang. Unlocking the power of patch: Patch-based mlp for long-term time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 12640–12648, 2025.

Hao Wang, Licheng Pan, Zhichao Chen, Degui Yang, Sen Zhang, Yifei Yang, Xinggao Liu, Haoxuan Li, and Dacheng Tao. Fredf: Learning to forecast in frequency domain. *arXiv preprint arXiv:2402.02399*, 2024a.

Shiyu Wang, Jiawei Li, Xiaoming Shi, Zhou Ye, Baichuan Mo, Wenze Lin, Shengtong Ju, Zhixuan Chu, and Ming Jin. Timemixer++: A general time series pattern machine for universal predictive analysis. *arXiv preprint arXiv:2410.16032*, 2024b.

Shiyu Wang, Haixu Wu, Xiaoming Shi, Tengge Hu, Huakun Luo, Lintao Ma, James Y Zhang, and Jun Zhou. Timemixer: Decomposable multiscale mixing for time series forecasting. *arXiv preprint arXiv:2405.14616*, 2024c.

Qingsong Wen, Tian Zhou, Chaoli Zhang, Weiqi Chen, Ziqing Ma, Junchi Yan, and Liang Sun. Transformers in time series: A survey. *arXiv preprint arXiv:2202.07125*, 2022.

Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems*, 34:22419–22430, 2021.

Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. Timesnet: Temporal 2d-variation modeling for general time series analysis. *arXiv preprint arXiv:2210.02186*, 2022a.

Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. Timesnet: Temporal 2d-variation modeling for general time series analysis. In *The Eleventh International Conference on Learning Representations*, 2022b.

Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang. Connecting the dots: Multivariate time series forecasting with graph neural networks. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 753–763, 2020.

Zhijian Xu, Ailing Zeng, and Qiang Xu. Fits: Modeling time series with 10^k parameters. *arXiv preprint arXiv:2307.03756*, 2023.

Hengyu Ye, Jiadong Chen, Shijin Gong, Fuxin Jiang, Tieying Zhang, Jianjun Chen, and Xiaofeng Gao. Atfnet: Adaptive time-frequency ensembled network for long-term time series forecasting. *arXiv preprint arXiv:2404.05192*, 2024.

Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. Are transformers effective for time series forecasting? In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pp. 11121–11128, 2023.

Yunhao Zhang and Junchi Yan. Crossformer: Transformer utilizing cross-dimension dependency for multivariate time series forecasting. In *The eleventh international conference on learning representations*, 2023.

Zhaodian Zhang, Guangpo Tian, Fenghua Guo, and Pengfei Wang. Tfp-mixer: A lightweight time and frequency combining model for multivariate long-term time series forecasting. *Applied Intelligence*, 55(7):1–18, 2025.

Hao Zheng, Zhanlei Yang, Wenju Liu, Jizhong Liang, and Yanpeng Li. Improving deep neural networks using softplus units. In *2015 International joint conference on neural networks (IJCNN)*, pp. 1–4. IEEE, 2015.

Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 11106–11115, 2021.

Tian Zhou, Ziqing Ma, Qingsong Wen, Liang Sun, Tao Yao, Wotao Yin, Rong Jin, et al. Film: Frequency improved legendre memory model for long-term time series forecasting. *Advances in neural information processing systems*, 35:12677–12690, 2022a.

Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International conference on machine learning*, pp. 27268–27286. PMLR, 2022b.

Jiachen Zhu, Xinlei Chen, Kaiming He, Yann LeCun, and Zhuang Liu. Transformers without normalization. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 14901–14911, 2025.

CONTENTS

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 2 | Related Work | 2 |
| 3 | Methodology | 3 |
| 3.1 | Period Bucket for Time Series | 3 |
| 3.1.1 | Period Detection | 3 |
| 3.1.2 | Time Series Period Bucket Construction and Representation | 3 |
| 3.2 | Positive-Negative X-shape Attention for Periodicity Modeling | 4 |
| 3.3 | Bucket-wise Forecasting | 6 |
| 4 | Experiments | 6 |
| 4.1 | Experimental Setup | 6 |
| 4.2 | Forecasting Performance Comparison | 7 |
| 4.3 | Comparison with Fixed Input Length Settings | 8 |
| 4.4 | Ablation Studies | 8 |
| 4.5 | Hyperparameter Sensitivity Experiments | 8 |
| 4.6 | Complexity Analysis | 9 |
| 4.7 | Visualization of Periodic Positive-negative Components in Data, Feature, and Attention Levels | 9 |
| 4.8 | Attention Visualization | 10 |
| 5 | Conclusion | 10 |
| A | Related Work | 18 |
| B | The Bucket without Periodicity | 18 |
| C | Mathematical Justification | 19 |
| C.1 | A Mathematics Definition of Period Heterogeneity | 19 |
| C.2 | A statistical Interpretation of Period-offset Attention (Another View) | 19 |
| C.3 | The Variance of Period-offset Attention under Period Heterogeneity | 21 |
| D | Experiments | 21 |
| D.1 | Experiments Details | 21 |
| D.2 | Dataset Periodicity Test | 21 |
| D.3 | Synthetic Dataset | 22 |
| D.4 | Detailed Forecasting Comparison | 23 |
| D.5 | Comparison of Periodic Heterogeneity Modeling | 23 |
| D.6 | Comparison of Transformer Variants | 24 |

| | | |
|----------|-------------------------|-----------|
| D.7 | Detailed Ablation Study | 25 |
| E | Discussions | 25 |
| F | Use of LLM | 26 |

A RELATED WORK

Multivariate Time Series Forecasting Deep learning has become the dominant paradigm for multivariate time series (MTS) forecasting thanks to its ability to model nonlinear dynamics and long-range dependencies (Mendis et al., 2024; Liang et al., 2024; Ma et al., 2025f;c). Researchers have developed numerous specialized and general-purpose methods for time series modeling architectures (Qiu et al., 2024b; Ma et al., 2026b; Kim et al., 2021), optimization techniques (Qiu et al., 2025; Wang et al., 2024a; Ma et al., 2025e), and inference strategies (Christou et al., 2024; Ma et al., 2026a). Early solutions relied on recurrent architectures (e.g., LSTM, GRU) or convolutional models such as TCN (Siami-Namini et al., 2019; Hewage et al., 2020). More recently, the Transformer has emerged as a popular backbone, offering flexible global dependency modeling for sequence-to-sequence tasks and coming to dominate the field (Wen et al., 2022). For example, ModernTCN (Luo & Wang, 2024) uses large convolutional kernels to substantially widen the receptive field and capture broader temporal dependencies, while Pyraformer (Liu et al., 2022a) combines TCN layers with a Transformer backbone to model both local and global patterns.

At the same time, lightweight MLP-based models have gained traction by achieving strong forecasting performance with lower computational cost. DLinear (Zeng et al., 2023) decomposes series into trend and seasonal components via a moving-average kernel and models each with separate linear modules; PatchMLP (Tang & Zhang, 2025) introduces a patch-based approach and channel mixing to enhance inter-variate communication; and HDMixer (Huang et al., 2024a) employs flexible patch lengths to capture short-term dependencies within patches and long-term dependencies across patches while better modeling cross-variate interactions. Overall, Transformer and lightweight MLP approaches present complementary trade-offs between representational capacity and computational efficiency.

Transformers for Time Series. Although the Transformer architecture has shown significant promise, it poses challenges such as high computational complexity (Zhang & Yan, 2023; Wu et al., 2021; Ma et al., 2025g). To mitigate this, various approaches have been proposed to modify the self-attention mechanism, incorporating techniques like distillation and temporal decomposition to reduce computational demands (Zhou et al., 2021; 2022b; Ma et al., 2025a). LogTrans (Li et al., 2019) introducing a convolutional LogSparse attention that ensured long-distance interactions while reducing the number of interactions, is an early and influential attempt to apply Transformer. Informer (Zhou et al., 2021) proposes ProbSparse attention while combining with a distillation mechanism to select the most representative query vectors to compute the attention scores. More recently, researchers have integrated signal processing methods with Transformers, leveraging techniques such as signal frequency analysis and decomposition (Piao et al., 2024; Zeng et al., 2023). These advancements have enabled models to more effectively capture and analyze the structural characteristics of time series data. However, the softmax operation in Transformers filters out negatively correlated signals—preserving only positive correlations—which leads to a loss of information.

B THE BUCKET WITHOUT PERIODICITY

Not all variates in a multivariate time series exhibit clear periodic behavior. To account for such cases, we introduce a dedicated bucket \mathcal{B}_0 , which collects all variates whose Fourier spectra do not reveal statistically significant periodicity. Although \mathcal{B}_0 does not correspond to any explicit cycle, it is still processed through the proposed PNA to maintain architectural consistency and allow residual correlations to be captured. **For notational convenience, we continue to use the terms *period-offset* and *period-aligned* to denote the two orthogonal attention directions, even though no true period exists.** Concretely, learning for bucket \mathcal{B}_0 differs from periodic buckets in two key aspects:

① No segmentation into periodic fragments. Since \mathcal{B}_0 has no detected period length, we do not segment its sequences into smaller fragments. Instead, each variate j is without any padding and folding operation preserving the full temporal structure without periodic folding.

② Offset attention with absolute distance. In the absence of periodic structure, the offset axis is defined by absolute temporal distance as follows,

$$\delta_{ij}^0 = |j - i| \in [0, L - 1], \quad (16)$$

where the period-offset attention is also coupled with positive attention and negative attention with the same modulation as in periodic buckets. This allows PHAT to still capture local autocorrelations and longer-range and potential temporal recurrences.

An important consequence is that the **period-aligned attention matrix degenerates into the identity** since the period-aligned attention matrix $\tilde{\mathbf{A}} \in \mathbb{R}^{L \times 1 \times 1}$ after softmax normalization along the last axis, all entries are equal to 1. Thus, the PNA attention in \mathcal{B}_0 is equal to,

$$\text{PNA}(\mathbf{Z}^{(0)}) = \overline{\mathbf{A}} \times_1 (\tilde{\mathbf{A}} \times_2 \mathbf{V}) = \overline{\mathbf{A}} \times_1 \mathbf{V} \in \mathbb{R}^{L \times 1 \times d}, \quad (17)$$

Overall, this design ensures that variates without periodicity are handled consistently within the PHAT framework: they bypass unnecessary folding while still benefiting from offset-based attention to capture temporal dependencies. By unifying both periodic and non-periodic buckets under the same framework, PHAT achieves robustness across diverse real-world datasets where not all signals are strictly cyclical.

C MATHEMATICAL JUSTIFICATION

In this section, we provide mathematical justification to clarify the motivation of PHAT and to address concerns raised during the review process. Section C.1 definite the Period Heterogeneity by mathematical formulation. Section C.2 shows that period-offset attention exhibits periodic behavior aligned with the autocorrelation periodicity of time series. Section C.3 demonstrates that decomposition of positive and negative paths yields strictly lower variance than vanilla attention, offering a clearer learning-theoretic rationale for the effectiveness of our method.

C.1 A MATHEMATICS DEFINITION OF PERIOD HETEROGENEITY

Here we formally defines period heterogeneity.

Definition: Period Heterogeneity. Period heterogeneity refers to the phenomenon where different variates exhibit distinct periodic lengths and periodic correlation patterns.

① Different periodic lengths. There exist two variates c_i and c_j whose dominant periods differ, i.e.,

$$P_{c_i} \neq P_{c_j}, \quad (18)$$

where P_{c_i} and P_{c_j} denote their corresponding fundamental period lengths.

② Different periodic correlations. Although periodic patterns are commonly characterized by the presence of correlations across specific time lags, existing models typically overlook the sign of such correlations. In particular, real-world sequences may exhibit negative correlations between certain time steps within a period. Let $\rho_c(\delta) = \text{ACF}(x_{c,t}, x_{c,t+\delta})$ denote the autocorrelation function of variate c at lag δ within a single period. Then this kind of period heterogeneity exists if there exist a variate c^* such that

$$\text{sign}(\rho_{c^*}([P/2])) \ll 0. \quad (19)$$

This definition reminds us that modeling period heterogeneity not only need to differences in period lengths among variates, but also to the diversity of period correlation patterns—ranging from purely positive correlations to positive and negative correlation patterns.

C.2 A STATISTICAL INTERPRETATION OF PERIOD-OFFSET ATTENTION (ANOTHER VIEW)

In this subsection, we will show that period-offset attention aligns with the inductive bias of auto-correlation in real-world time series by separating positive and negative paths and assigning them different modulation types. In addition, through a statistical stick-breaking perspective (Ren et al., 2011), we provide an interpretable view of the mechanism and show that the upper bound of the attention logits is monotonic with respect to the period distance.

Proposition. Period-offset Attention as a Stick-breaking Allocation Let $\zeta[m, n]$ and $\eta[m, n]$ denote the raw positive and negative logit between time steps m and n within a detected period of length P_b , respectively. Define the modulated positive and negative logits of the period-offset attention as

$$\tilde{\zeta}[m, n] = \zeta[m, n] - \sum_{s \in \Delta_{m, n}^{(b)}} \text{Softplus}(\zeta[m, s]), \quad (20)$$

$$\tilde{\eta}[m, n] = \eta[m, n] - \sum_{s \in \nabla_{m, n}^{(b)}} \text{Softplus}(\eta[m, s]), \quad (21)$$

where $\Delta_{m, n}^{(b)} = \{s \mid \delta_{m, s}^b < \delta_{m, n}^b\} \cup \{n\}$ and $\nabla_{m, n}^{(b)} = \{s \mid \delta_{m, s}^b > \delta_{m, n}^b\} \cup \{n\}$ is the set of offsets closer and farther to m than n under corresponding distance $\delta_{m, s}^b$. Then, before the normalization step in softmax, the exponentiated logits satisfy the following stick-breaking decomposition (Ren et al., 2011):

$$\exp(\tilde{\zeta}[m, n]) = \sigma(\zeta[m, n]) \prod_{\{s \mid \delta_{m, s} < \delta_{m, n}\}} (1 - \sigma(\zeta[m, s])), \quad (22)$$

$$\exp(\tilde{\eta}[m, n]) = \sigma(\eta[m, n]) \prod_{\{s \mid \delta_{m, s} > \delta_{m, n}\}} (1 - \sigma(\eta[m, s])), \quad (23)$$

where $\sigma(\cdot)$ denotes the sigmoid function.

Proof of Proposition: The logits from positive branch of period-offset attention is,

$$\begin{aligned} & \tilde{\zeta}[m, n] \\ &= \zeta[m, n] - \sum_{s \in \Delta_{m, n}^{(b)}} \text{Softplus}(\zeta[m, s]) \\ &= \zeta[m, n] - \sum_{s \in \Delta_{m, n}^{(b)}} \ln(1 + \exp(\zeta[m, s])) \\ &= \zeta[m, n] - \ln(1 + \exp(\zeta[m, n])) - \sum_{s \in \Delta_{m, n}^{(b)} - \{n\}} \ln(1 + \exp(\zeta[m, s])) \\ &= \ln\left(\frac{\exp(\zeta[m, n])}{1 + \exp(\zeta[m, n])}\right) + \sum_{s \in \Delta_{m, n}^{(b)} - \{n\}} \ln\left(\frac{1}{1 + \exp(\zeta[m, s])}\right) \\ &= \ln\left(\frac{1}{1 + \exp^{-1}(\zeta[m, n])}\right) + \sum_{s \in \Delta_{m, n}^{(b)} - \{n\}} \ln\left(1 - \frac{\exp(\zeta[m, s])}{1 + \exp(\zeta[m, s])}\right) \\ &= \ln(\sigma(\zeta[m, n])) + \sum_{s \in \Delta_{m, n}^{(b)} - \{n\}} \ln\left(1 - \frac{1}{1 + \exp^{-1}(\zeta[m, s])}\right) \\ &= \ln(\sigma(\zeta[m, n])) + \sum_{s \in \Delta_{m, n}^{(b)} - \{n\}} \ln(1 - \sigma(\zeta[m, s])), \end{aligned} \quad (24)$$

where $\sigma(\cdot)$ is the sigmoid function and set $\Delta_{m, n}^{(b)} - \{n\} = \{s \mid \delta_{m, s}^b < \delta_{m, n}^b\}$ means all of the closer time step from m than n . Hence before the normalization operator in Softamax(\cdot), the postie logits are activated by exponential function satisfying,

$$\begin{aligned} \exp(\tilde{\zeta}[m, n]) &= \exp(\ln(\sigma(\zeta[m, n])) * \exp(\sum_{\{s \mid \delta_{m, s} < \delta_{m, n}\}} \ln(1 - \sigma(\zeta[m, s])))) \\ &= \exp(\ln(\sigma(\zeta[m, n])) * \prod_{\{s \mid \delta_{m, s} < \delta_{m, n}\}} \exp(\ln(1 - \sigma(\zeta[m, s])))) \\ &= \sigma(\zeta[m, n]) * \prod_{\{s \mid \delta_{m, s} < \delta_{m, n}\}} (1 - \sigma(\zeta[m, s])). \end{aligned} \quad (25)$$

The proof of the situation of negative logits is similar. \square

This identity reveals that the unnormalized attention coefficient assigned to time step n is exactly a stick-breaking allocation. Attention coefficient is allocated sequentially in the order of periodic distance: nearby offsets consume mass first, while farther offsets only receive the leftover. Unlike standard dot-product attention, which distributes weights independently, our period-offset formulation embeds a period-aware stick-breaking bias, yielding structured sparsity, interpretability, and inductive alignment with the autocorrelation structure of time series.

Corollary: Local dominance of periodic distance

From Proposition 1, it follows that for any reference time step m , the allocation weight to time step n in period-offset attention is strictly upper-bounded by the leftover stick mass after all closer offsets have been considered:

$$\exp(\tilde{\zeta}[m, n]) < \prod_{\{s | \delta_{m, s} < \delta_{m, n}\}} (1 - \sigma(\zeta[m, s])), \quad (26)$$

$$\exp(\tilde{\eta}[m, n]) < \prod_{\{s | \delta_{m, s} > \delta_{m, n}\}} (1 - \sigma(\eta[m, s])), \quad (27)$$

This corollary formalizes the intuition that local periodic dependencies are always prioritized: nearby (far away) time steps consume most of the positive (negative) attention budget, and more distant steps are suppressed unless strongly supported by their logits. This guarantees that period-offset attention respects the inductive bias of autocorrelation in real-world time series, where closer (farther) periodic positions tend to carry stronger positive (negative) correlations.

C.3 THE VARIANCE OF PERIOD-OFFSET ATTENTION UNDER PERIOD HETEROGENEITY

As the request from reviewer tCVf, we add a brief learning-theoretic argument to further strengthen the theoretical contribution. Let the vanilla attention between steps i and j be $\mathbf{A}_{ij} = \text{Softmax}(\mu \mathbf{Q}_i \mathbf{K}_j^\top)$ with variance $\mathbb{V}[\mathbf{A}_{ij}] = \sigma^2$. Period-offset attention decomposes this interaction into positive and negative paths, leading to

$$\mathbb{V}[\bar{\mathbf{A}}_{ij}] = (1 - \Lambda_{ij}) \sigma^2, \quad (28)$$

where $\Lambda_{ij} \in (0, 1)$ when periodic heterogeneity induces a non-zero negative path. Thus, period-offset attention yields strictly lower variance than vanilla attention. This provides a clear learning-theoretic intuition: the positive and negative decomposition reduces attention-logit variance in heterogeneous periodic settings, producing more stable and accurate estimates than vanilla attention.

D EXPERIMENTS

D.1 EXPERIMENTS DETAILS

We report the detailed hyperparameters of our model in Table 5. All experiments were conducted on a single NVIDIA A100 GPU with 80 GB of memory using PyTorch. Forecasting performance was evaluated using mean squared error (MSE) and mean absolute error (MAE). For a comprehensive assessment, we disabled the “Drop Last” batch sampling procedure (Li et al., 2024; Qiu et al., 2024a) and trained models with the Adam optimizer (Kingma & Ba, 2014). Because models vary in sensitivity to input history, the look back length T was treated as a tunable hyperparameter and each model’s best performance is reported. Hyperparameter settings are summarized in Table 5.

D.2 DATASET PERIODICITY TEST

We further visualize temporal correlations across multiple datasets. We use the Autocorrelation Function (Dégerine & Lambert-Lacroix, 2003) to analyze the autoregressiveness between time steps. As shown in Figure 6, most datasets exhibit complex periodic heterogeneity. For example, in the ETT1 dataset, variate 3 shows only positively correlated periodicity, while variate 4 demonstrates

Table 5: Optimal hyperparameters on our experiments.

| Hyperparameters | NN5 | Exchange | FRED-MD | ETTh1 | ETTh2 | ETTm1 | ETTm2 | AQShunyi | AQWan | IL1 | CzeLan | ZafNoo | NASDAQ | NYSE | |
|-----------------|------------------|----------|---------|-------|-------|-------|--------|----------|--------|-------|--------|--------|--------|------|-----|
| L=24/96 | Batch Size | 16 | 16 | 16 | 16 | 32 | 8 | 16 | 32 | 32 | 8 | 16 | 16 | 256 | 32 |
| | Learning Rate | 0.01 | 0.001 | 0.01 | 0.001 | 0.001 | 0.01 | 0.001 | 0.001 | 0.001 | 0.01 | 0.001 | 0.01 | 0.01 | 0.1 |
| | Model Dimension | 32 | 4 | 32 | 4 | 4 | 4 | 8 | 4 | 8 | 16 | 24 | 4 | 4 | 4 |
| | Layers | 1 | 1 | 1 | 2 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| | Head Number | 2 | 2 | 2 | 4 | 4 | 4 | 2 | 2 | 2 | 2 | 4 | 2 | 4 | 4 |
| L=36/192 | Batch Size | 256 | 32 | 16 | 16 | 16 | 16 | 32 | 32 | 64 | 16 | 16 | 16 | 64 | 64 |
| | Learning Rate | 0.01 | 0.01 | 0.01 | 0.001 | 0.001 | 0.01 | 0.001 | 0.0001 | 0.01 | 0.0001 | 0.0001 | 0.01 | 0.01 | 0.1 |
| | Model Dimension | 128 | 8 | 96 | 4 | 12 | 8 | 8 | 8 | 8 | 8 | 4 | 6 | 24 | 6 |
| | Layers | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| | Head Number | 4 | 4 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 4 | 2 | 2 | 2 |
| L=48/336 | Batch Size | 256 | 64 | 16 | 16 | 16 | 16 | 32 | 32 | 64 | 16 | 16 | 32 | 64 | 64 |
| | Learning Rate | 0.01 | 0.01 | 0.01 | 0.001 | 0.001 | 0.0001 | 0.0001 | 0.0001 | 0.01 | 0.001 | 0.0001 | 0.1 | 0.1 | 0.1 |
| | Model Dimension | 64 | 8 | 32 | 4 | 16 | 8 | 8 | 16 | 16 | 8 | 4 | 4 | 8 | 8 |
| | Layers | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| | Head Number | 4 | 2 | 2 | 2 | 4 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 4 |
| L=60/720 | Batch Size | 128 | 256 | 32 | 32 | 128 | 16 | 16 | 32 | 32 | 8 | 16 | 16 | 256 | 64 |
| | Learning Rate | 0.01 | 0.1 | 0.1 | 0.001 | 0.001 | 0.0001 | 0.0001 | 0.0001 | 0.01 | 0.0001 | 0.0001 | 0.1 | 0.1 | 0.1 |
| | Model Dimension | 32 | 16 | 32 | 4 | 12 | 4 | 4 | 8 | 8 | 4 | 4 | 4 | 32 | 2 |
| | Layers | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| | Head Number | 8 | 2 | 2 | 2 | 4 | 4 | 2 | 4 | 2 | 2 | 2 | 4 | 2 | 2 |
| | Look Back Window | 104 | 96 | 36 | 512 | 512 | 336 | 336 | 336 | 512 | 104 | 512 | 512 | 104 | 36 |

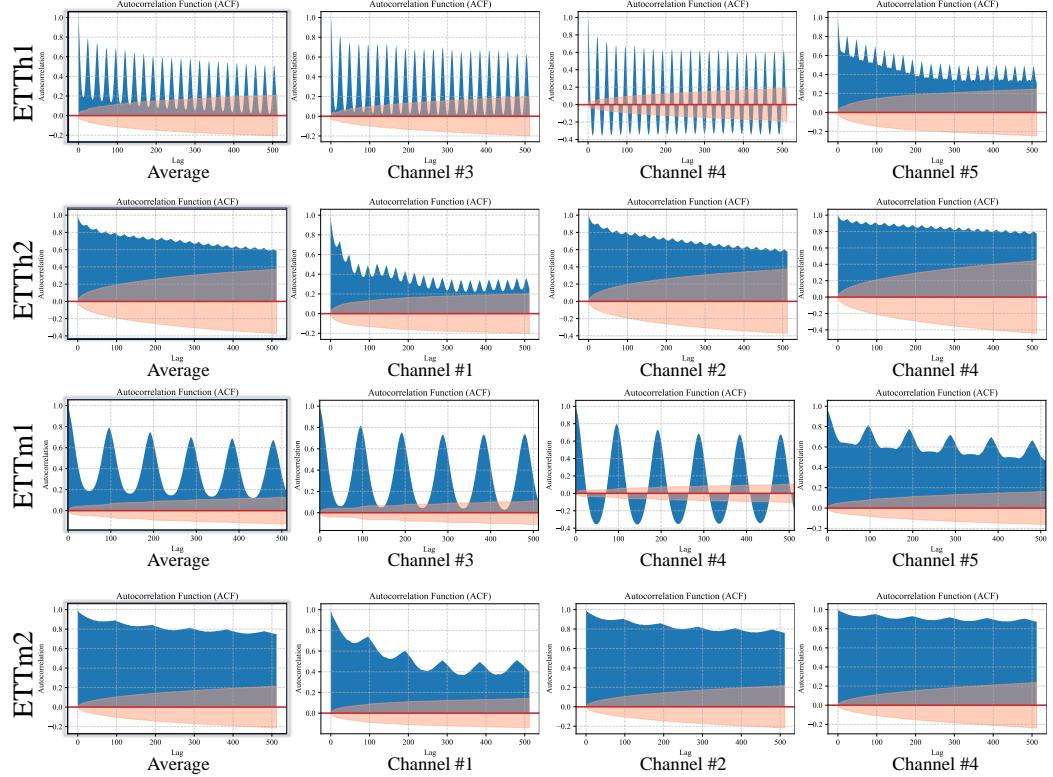


Figure 6: The autocorrelation visualization of ETT datasets.

a bilateral periodic pattern with both positive and negative correlations, and variate 5 exhibits no significant periodicity. In contrast, the ETTh2 dataset lacks clear periodicity across the entire dataset.

D.3 SYNTHETIC DATASET

We systematically created a synthetic dataset by combining ETTh1 (with a period of 24) and ETTm1 (with a period of 96), which record data from the same variates but at different time granularities. Each dataset was divided into four equal parts along the time axis. Then, we alternately sampled

Table 6: Multivariate forecasting results on 14 datasets. We report MSE and MAE. Best results are **bold**, second-best are underlined. The corresponding results on ETTh1 and ETTh2 (denoted ETTh), ETTm1 and ETTm2 (denoted ETTm) are averaged for better presentation.

| Methods | Ours 2025 | FITS 2024 | PatchTST 2023 | Crossformer 2023 | NLinear 2023 | TimesNet 2023 | FEDformer 2022 | Triformer 2022 | FiLM 2022 | Stationary 2022 | | | | | | | | | | | |
|----------|--------------|----------------------------|---------------------------|---------------------|---------------------------|---------------------------|-------------------|---------------------------|-----------------------------|--------------------|---------------------------|--------------|---------------------------|---------|--------------|---------------------------|--------|---------------------------|---------------------------|-------|-------|
| Metrics | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | | | | | | | | | | | |
| ETTh | 96 | 0.316 0.356 | 0.327 | 0.371 | <u>0.326</u> 0.367 | 0.570 | 0.519 | 0.331 | 0.371 | 0.362 | 0.391 | 0.358 | 0.400 | 0.668 | 0.543 | 0.326 | 0.370 | 0.469 | 0.456 | | |
| | 192 | 0.360 0.394 | <u>0.366</u> 0.399 | 0.379 | 0.405 | 0.566 | 0.523 | 0.384 | 0.404 | 0.422 | 0.428 | 0.418 | 0.436 | 0.867 | 0.609 | 0.382 | 0.409 | 0.497 | 0.479 | | |
| | 336 | 0.384 0.412 | <u>0.385</u> 0.416 | 0.404 | 0.430 | 0.587 | 0.543 | 0.400 | 0.419 | 0.456 | 0.461 | 0.424 | 0.462 | 0.909 | 0.630 | 0.403 | 0.430 | 0.495 | 0.482 | | |
| | 720 | 0.412 0.441 | <u>0.409</u> 0.442 | 0.432 | 0.459 | 0.945 | 0.698 | 0.423 | 0.447 | 0.478 | 0.472 | 0.479 | 0.488 | 1.025 | 0.690 | 0.444 | 0.4645 | 0.625 | 0.558 | | |
| ETTm | 96 | 0.224 0.288 | 0.234 | 0.300 | <u>0.227</u> 0.299 | 0.305 | 0.379 | 0.232 | 0.298 | 0.265 | 0.322 | 0.340 | 0.386 | 0.313 | 0.366 | 0.233 | 0.299 | 0.313 | 0.352 | | |
| | 192 | 0.272 0.323 | 0.278 | <u>0.328</u> | 0.275 | 0.331 | 0.372 | 0.413 | 0.287 | 0.335 | 0.323 | 0.357 | 0.436 | 0.438 | 0.430 | 0.432 | 0.280 | <u>0.328</u> | 0.416 | 0.412 | |
| | 336 | 0.311 0.349 | 0.320 | <u>0.355</u> | 0.319 | 0.359 | 0.501 | 0.516 | 0.323 | 0.356 | 0.368 | 0.386 | 0.492 | 0.472 | 0.559 | 0.498 | 0.326 | 0.357 | 0.505 | 0.453 | |
| | 720 | 0.388 0.392 | 0.390 | <u>0.397</u> | 0.389 | 0.402 | 0.752 | 0.613 | 0.396 | 0.400 | 0.444 | 0.428 | 0.536 | 0.501 | 1.209 | 0.666 | 0.393 | 0.400 | 0.595 | 0.506 | |
| NN5 | 24 | 0.681 0.551 | 0.870 | 0.663 | 0.740 | 0.596 | 0.741 | 0.591 | 0.758 | 0.592 | 0.739 0.579 | 0.785 | 0.618 | 1.382 | 0.929 | 0.846 | 0.651 | 1.274 | 0.900 | | |
| | 36 | 0.640 0.538 | 0.814 | 0.655 | 0.694 | 0.595 | 0.703 | 0.589 | <u>0.693</u> 0.577 | 0.717 | 0.585 | 0.727 | 0.606 | 1.352 | 0.920 | 0.883 | 0.702 | 1.318 | 0.93 | | |
| | 48 | 0.613 0.532 | 0.780 | 0.644 | 0.667 | 0.585 | 0.669 | 0.575 | 0.688 | 0.587 | 0.647 | 0.558 | 0.623 0.555 | 1.348 | 0.918 | 0.969 | 0.741 | 1.277 | 0.905 | | |
| | 60 | 0.603 0.531 | 0.781 | 0.650 | 0.653 | 0.582 | 0.683 | 0.587 | 0.679 | 0.587 | 0.633 | <u>0.547</u> | 0.630 | 1.346 | 0.918 | 0.633 | 0.556 | 1.313 | 0.927 | | |
| FRED-MD | 96 | 0.083 0.203 | 0.082 0.199 | 0.087 | 0.204 | 0.231 | 0.356 | 0.085 | 0.204 | 0.112 | 0.242 | 0.138 | 0.268 | 0.201 | 0.335 | 0.087 | 0.210 | 0.121 | 0.247 | | |
| | 192 | 0.169 0.295 | <u>0.173</u> 0.295 | 0.177 | 0.300 | 0.460 | 0.509 | 0.175 | 0.297 | 0.209 | 0.334 | 0.273 | 0.379 | 0.453 | 0.495 | 0.182 | 0.308 | 0.220 | 0.337 | | |
| | 336 | 0.316 0.406 | 0.317 | <u>0.406</u> | 0.297 0.399 | 1.034 | 0.825 | 0.320 | 0.409 | 0.358 | 0.437 | 0.485 | 0.703 | 0.630 | 0.318 | 0.409 | 0.352 | 0.437 | | | |
| | 720 | 0.773 0.669 | 0.825 | 0.684 | 0.843 | 0.692 | 1.576 | 1.021 | 0.838 | 0.690 | 0.944 | 0.736 | 1.158 | 0.828 | 1.395 | 0.915 | 0.815 | 0.725 0.656 | | | |
| AQShunyi | 24 | 21.495 0.903 | 56.779 | 1.374 | 32.808 | 0.962 | 385.599 | 3.559 | <u>32.125</u> 0.931 | 43.219 | 1.265 | 66.090 | 1.623 | 395.947 | 5.049 | 40.183 | 1.145 | 47.852 | 1.238 | | |
| | 36 | 38.363 1.161 | 97.396 | 1.774 | 61.035 | 1.345 | 398.728 | 3.716 | <u>58.332</u> 1.258 | 69.554 | 1.531 | 94.359 | 1.863 | 412.165 | 5.210 | 90.434 | 1.670 | 68.140 | 1.493 | | |
| | 48 | 52.065 1.453 | 145.471 | 2.183 | 91.835 | 1.648 | 414.353 | 3.939 | 82.184 1.609 | 95.071 | 1.810 | 129.798 | 2.135 | 423.926 | 5.234 | 131.081 | 2.119 | 92.906 | 1.736 | | |
| | 60 | 76.292 1.679 | 196.613 | 2.523 | 127.018 | 1.958 | 422.864 | 4.093 | 109.625 1.882 | 116.341 | 1.976 | 173.616 | 2.435 | 432.464 | 5.299 | 180.367 | 2.397 | 117.756 | 1.929 | | |
| AQWan | 96 | 0.665 | 0.468 | 0.655 | 0.485 | 0.464 | 0.478 | <u>0.552</u> | 0.484 | 0.653 | 0.486 | 0.658 | 0.488 | 0.706 | 0.525 | 0.665 | 0.492 | 0.664 | 0.486 | | |
| | 192 | 0.707 | 0.488 | 0.701 | 0.503 | 0.688 | <u>0.498</u> | 0.674 | 0.499 | 0.701 | 0.506 | 0.707 | 0.512 | 0.729 | 0.531 | 0.681 | 0.501 | 0.705 | 0.504 | | |
| | 336 | 0.737 | 0.502 | 0.720 | 0.515 | <u>0.710</u> | 0.513 | 0.704 | 0.515 | 0.722 | 0.519 | 0.786 | 0.538 | 0.824 | 0.569 | 0.731 | 0.525 | 0.725 | 0.517 | | |
| | 720 | 0.795 | 0.529 | 0.774 | 0.540 | 0.768 | 0.539 | <u>0.747</u> | 0.518 | 0.777 | 0.545 | 0.756 | 0.528 | 0.794 | 0.561 | 0.742 | 0.534 | 0.782 | 0.544 | | |
| IL1 | 96 | 0.746 0.452 | 0.757 | 0.473 | 0.745 0.468 | 0.750 | <u>0.465</u> | 0.758 | 0.475 | 0.787 | 0.486 | 0.796 | 0.508 | 0.762 | 0.474 | 0.766 | 0.475 | 0.872 | 0.508 | | |
| | 192 | 0.808 | 0.475 | 0.806 | 0.492 | 0.793 | 0.490 | 0.762 0.479 | 0.809 | 0.496 | 0.778 | 0.849 | 0.825 | 0.517 | 0.786 | 0.484 | 0.809 | 0.494 | 0.860 | 0.524 | |
| | 336 | 0.824 | 0.487 | 0.826 | 0.504 | 0.819 | 0.502 | 0.802 | 0.504 | 0.830 | 0.508 | 0.815 | 0.505 | 0.863 | 0.537 | 0.802 0.495 | 0.831 | 0.505 | 0.864 | 0.532 | |
| | 720 | 0.907 | 0.517 | 0.900 | 0.532 | 0.890 | 0.533 | 0.829 0.512 | 0.906 | 0.538 | 0.869 | 0.519 | 0.907 | 0.552 | 0.852 | 0.519 | 0.906 | 0.536 | 0.897 | 0.526 | |
| CzclLan | 24 | 1.318 0.705 | 2.182 | 1.002 | <u>1.932</u> 0.872 | 2.981 | 1.096 | 1.998 | 0.919 | 2.131 | 0.958 | 2.398 | 1.02 | 6.052 | 1.730 | 2.256 | 0.996 | 2.394 | 1.066 | | |
| | 36 | 1.523 0.775 | 2.330 | 1.051 | 1.869 0.866 | 3.549 | 1.196 | 1.920 | 0.916 | 2.612 | 0.974 | 2.410 | 1.005 | 6.111 | 1.743 | 2.133 | 0.992 | 2.226 | 1.031 | | |
| | 48 | 1.437 0.775 | 2.761 | 1.184 | 1.891 0.883 | 3.851 | 1.288 | 1.895 | 0.924 | 1.916 | 0.897 | 2.591 | 1.033 | 6.218 | 1.774 | 2.034 | 0.969 | 2.525 | 1.003 | | |
| | 60 | 1.461 0.777 | 2.929 | 1.217 | 1.914 0.896 | 4.692 | 1.450 | 1.964 | 0.947 | 1.995 | 0.905 | 2.539 | 1.070 | 7.000 | 1.893 | 1.974 | 0.929 | 2.410 | 1.010 | | |
| NASDAQ | 96 | 0.170 0.204 | 0.187 | 0.243 | <u>0.176</u> | 0.232 | 0.581 | 0.443 | 0.178 | 0.228 | 0.179 | 0.239 | 0.231 | 0.311 | 0.818 | 0.519 | 0.180 | 0.232 | 0.240 | 0.296 | |
| | 192 | 0.206 0.227 | 0.214 | 0.263 | <u>0.205</u> | 0.263 | 0.705 | 0.503 | 0.210 | 0.252 | 0.216 | 0.28 | 0.283 | 0.349 | 0.962 | 0.589 | 0.212 | 0.255 | 0.280 | 0.315 | |
| | 336 | 0.236 0.254 | 0.247 | 0.292 | 0.236 | 0.286 | 0.971 | 0.596 | <u>0.243</u> 0.280 | 0.265 | 0.313 | 0.298 | 0.363 | 1.161 | 0.659 | 0.243 | 0.281 | 0.311 | 0.341 | | |
| | 720 | 0.285 | 0.293 | 0.291 | 0.329 | 0.270 | 0.316 | 1.566 | 0.762 | 0.290 | 0.326 | 0.283 | 0.338 | 0.426 | 0.449 | 1.496 | 0.741 | 0.282 0.312 | 0.328 | 0.379 | |
| NYSE | 96 | 0.439 | 0.383 | 0.449 | 0.412 | 0.429 0.405 | 0.430 | 0.418 | 0.447 | 0.410 | 0.478 | 0.419 | 0.476 | 0.450 | 0.441 | 0.419 | 0.451 | 0.411 | 0.524 | 0.449 | |
| | 192 | 0.503 | 0.429 | 0.511 | 0.447 | 0.494 | 0.479 | 0.449 | 0.503 | 0.447 | 0.491 0.445 | 0.544 | 0.479 | 0.493 | 0.451 | 0.508 | 0.448 | 0.562 | 0.475 | | |
| | 336 | 0.547 | 0.456 | 0.544 | 0.468 | 0.538 | 0.475 | 0.505 0.464 | 0.545 | 0.470 | 0.551 | 0.480 | 0.628 | 0.523 | 0.523 | 0.467 | 0.549 | 0.471 | 0.614 | 0.502 | |
| | 720 | 0.592 | 0.477 | 0.585 | 0.491 | 0.573 | 0.486 | 0.560 | 0.494 | 0.589 | 0.497 | 0.626 | 0.511 | 0.652 | 0.562 | 0.564 | 0.491 | 0.598 | 0.504 | 0.692 | 0.542 |
| | 24 | 0.416 0.473 | 0.709 | 0.645 | 0.649 | 0.567 | 1.149 | 0.745 | 0.557 | 0.522 | 0.587 | 0.533 | 0.537 0.481 | 2.737 | 1.334 | 0.767 | 0.645 | 0.655 | 0.607 | | |
| | 36 | 0.610 0.600 | 1.058 | 0.821 | 1.414 | 0.885 | 0.869 | 0.668 | <u>0.792</u> | 0.664 | 0.808 | 0.628 | 3.387 | 1.534 | 1.379 | 0.835 | 0.991 | 0.695 | | | |
| | 48 | 0.949 0.731 | 1.255 | 0.834 | 1.169 | 0.793 | 2.108 | 1.136 | 1.152 | 0.770 | 1.216 | 0.783 | 1.137 0.746 | 3.425 | 1.555 | 1.179 | 0.829 | 1.260 | 0.814 | | |
| | 60 | 0.961 0.779 | 1.153 | 0.818 | 1.247 | 0.843 | 2.276 | 1.201 | 1.284 | 0.809 | 1.220 | 0.768 | 1.251 | 0.783 | 3.313 | 1.537 | 1.303 | 0.853 | 1.119 0.819 | | |
| | # Top1 | 68 | 4 | 8 | 10 | 1 | 1 | 1 | 1 | 3 | 2 | 0 | 0 | 1 | | | | | | | |
| | # Top2 | 83 | 14 | 29 | 16 | 20 | 10 | 9 | 7 | 5 | 3 | | | | | | | | | | |

each segment to construct the **Synthetic**, represented as ETTh1[: $\frac{1}{4}$], ETTm1[: $\frac{1}{4}:\frac{1}{2}$], ETTh1[: $\frac{1}{2}:\frac{3}{4}$], ETTm1[: $\frac{3}{4}:\frac{1}{2}$]. As a result, the synthetic dataset exhibits mixed periodic lengths of (24, 96, 24, 96).

D.4 DETAILED FORECASTING COMPARISON

We present additional benchmark comparison results in Table 6 and provide fine-grained experimental results for ETTh in Table 7. Among the advanced baseline models, Crossformer achieves strong predictive performance due to its cross-dimension attention mechanism, which facilitates efficient modeling of long-term dependencies. PatchTST, on the other hand, introduces self-supervised learning and adopts a channel-independent strategy to capture complex temporal dynamics. Overall, PHAT achieves state-of-the-art (SOTA) performance, demonstrating exceptional adaptability to periodic heterogeneity. It outperforms other models on the majority of metrics, ranking in the top two positions for 86% of the evaluation metrics, thereby showcasing a significant performance advantage.

Table 7: Detailed forecasting results on ETTh1, ETTh2, ETTm1 and ETTm2 datasets. We report MSE and MAE. Best results are **bold**, second-best are underlined.

| Datasets | | ETTh1 | | | | ETTh2 | | | | ETTm1 | | | | ETTm2 | | | | |
|----------|---------------|---------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| # | Top2 | Methods | 96 | 192 | 336 | 720 | 96 | 192 | 336 | 720 | 96 | 192 | 336 | 720 | 96 | 192 | 336 | 720 |
| 27 | PHAT (Ours) | MSE | <u>0.361</u> | <u>0.393</u> | 0.414 | 0.441 | <u>0.272</u> | <u>0.328</u> | <u>0.354</u> | <u>0.384</u> | <u>0.287</u> | 0.328 | <u>0.360</u> | 0.422 | <u>0.162</u> | <u>0.216</u> | 0.262 | 0.355 |
| | | MAE | <u>0.383</u> | <u>0.410</u> | <u>0.427</u> | 0.458 | <u>0.330</u> | <u>0.379</u> | <u>0.398</u> | <u>0.425</u> | <u>0.330</u> | <u>0.362</u> | 0.381 | <u>0.412</u> | <u>0.246</u> | 0.285 | <u>0.317</u> | <u>0.372</u> |
| 5 | TimeKan | MSE | 0.369 | 0.402 | 0.419 | 0.442 | 0.279 | <u>0.326</u> | 0.368 | 0.408 | <u>0.286</u> | 0.331 | <u>0.354</u> | 0.400 | 0.163 | 0.238 | 0.277 | 0.359 |
| | | MAE | 0.396 | 0.417 | 0.430 | 0.463 | 0.343 | <u>0.380</u> | 0.411 | 0.443 | 0.346 | 0.368 | 0.386 | 0.417 | 0.254 | 0.299 | 0.330 | 0.387 |
| 8 | xPatch | MSE | 0.378 | 0.420 | 0.467 | 0.516 | 0.275 | 0.348 | 0.373 | 0.395 | 0.288 | 0.331 | 0.367 | 0.447 | <u>0.162</u> | 0.202 | 0.272 | <u>0.354</u> |
| | | MAE | 0.395 | 0.426 | 0.451 | 0.493 | <u>0.332</u> | <u>0.380</u> | 0.403 | 0.427 | <u>0.335</u> | 0.360 | <u>0.383</u> | 0.428 | <u>0.247</u> | 0.289 | 0.320 | 0.381 |
| 1 | Amplifier | MSE | 0.373 | 0.414 | 0.442 | 0.455 | 0.295 | 0.348 | 0.383 | 0.407 | 0.292 | <u>0.327</u> | 0.365 | 0.427 | 0.174 | 0.226 | 0.276 | 0.358 |
| | | MAE | 0.399 | 0.420 | 0.446 | 0.467 | 0.354 | 0.393 | 0.424 | 0.444 | 0.348 | 0.365 | 0.386 | 0.419 | 0.257 | 0.300 | 0.331 | 0.388 |
| 4 | CycleNet | MSE | 0.374 | 0.406 | 0.431 | 0.450 | 0.279 | 0.342 | 0.371 | 0.426 | 0.299 | 0.334 | 0.368 | 0.417 | 0.159 | 0.214 | 0.269 | 0.363 |
| | | MAE | 0.396 | 0.415 | 0.430 | 0.464 | 0.341 | 0.385 | 0.413 | 0.451 | 0.348 | 0.367 | 0.386 | 0.414 | <u>0.247</u> | <u>0.286</u> | 0.322 | 0.382 |
| 0 | TimeMixer | MSE | 0.372 | 0.413 | 0.438 | 0.486 | 0.281 | 0.349 | 0.366 | 0.401 | 0.293 | 0.335 | 0.368 | 0.426 | 0.165 | 0.225 | 0.277 | 0.36 |
| | | MAE | 0.401 | 0.430 | 0.450 | 0.484 | 0.351 | 0.387 | 0.413 | 0.436 | 0.345 | 0.372 | 0.386 | 0.417 | 0.256 | 0.298 | 0.332 | 0.387 |
| 10 | SparseTSF | MSE | <u>0.361</u> | 0.394 | <u>0.415</u> | 0.419 | 0.283 | 0.347 | 0.358 | <u>0.384</u> | 0.316 | 0.348 | 0.373 | 0.434 | 0.166 | 0.220 | 0.273 | 0.361 |
| | | MAE | <u>0.386</u> | <u>0.406</u> | 0.419 | <u>0.440</u> | 0.341 | 0.381 | <u>0.396</u> | <u>0.422</u> | 0.355 | 0.376 | 0.387 | 0.422 | 0.256 | 0.292 | 0.327 | 0.381 |
| 0 | iTTransformer | MSE | 0.386 | 0.424 | 0.449 | 0.495 | 0.297 | 0.372 | 0.388 | 0.424 | 0.300 | 0.341 | 0.374 | 0.429 | 0.175 | 0.242 | 0.282 | 0.375 |
| | | MAE | 0.405 | 0.440 | 0.460 | 0.487 | 0.348 | 0.403 | 0.417 | 0.444 | 0.353 | 0.380 | 0.396 | 0.430 | 0.266 | 0.312 | 0.337 | 0.394 |
| 5 | Pathformer | MSE | 0.372 | 0.408 | 0.438 | 0.450 | 0.279 | 0.345 | 0.378 | 0.437 | 0.290 | 0.337 | 0.374 | 0.428 | 0.164 | 0.219 | <u>0.267</u> | 0.361 |
| | | MAE | 0.392 | 0.415 | 0.434 | 0.463 | 0.336 | <u>0.380</u> | 0.408 | 0.455 | <u>0.335</u> | 0.363 | 0.384 | 0.416 | 0.250 | 0.288 | <u>0.319</u> | <u>0.377</u> |
| 7 | PDF | MSE | 0.360 | <u>0.392</u> | 0.418 | 0.456 | 0.276 | 0.339 | 0.374 | 0.398 | <u>0.286</u> | <u>0.321</u> | <u>0.354</u> | 0.408 | 0.163 | 0.219 | 0.269 | <u>0.349</u> |
| | | MAE | 0.391 | 0.414 | 0.435 | 0.462 | 0.341 | 0.382 | 0.406 | 0.433 | 0.340 | 0.364 | <u>0.383</u> | 0.415 | 0.251 | 0.290 | 0.330 | 0.382 |
| 7 | FITS | MSE | 0.376 | 0.400 | 0.419 | <u>0.435</u> | 0.277 | 0.331 | <u>0.350</u> | 0.382 | 0.303 | 0.337 | 0.368 | 0.420 | 0.165 | 0.219 | 0.272 | 0.359 |
| | | MAE | 0.396 | 0.418 | 0.435 | 0.458 | 0.345 | <u>0.379</u> | <u>0.396</u> | <u>0.425</u> | 0.345 | 0.365 | 0.384 | <u>0.413</u> | 0.254 | 0.291 | 0.326 | 0.381 |
| 1 | PatchTST | MSE | 0.377 | 0.409 | 0.431 | 0.457 | <u>0.274</u> | 0.348 | 0.377 | 0.406 | 0.289 | 0.329 | 0.362 | 0.416 | 0.165 | 0.221 | 0.276 | 0.362 |
| | | MAE | 0.397 | 0.425 | 0.444 | 0.477 | 0.337 | 0.384 | 0.416 | 0.441 | 0.343 | 0.368 | 0.390 | 0.423 | 0.255 | 0.293 | 0.327 | 0.381 |
| 0 | Crossformer | MSE | 0.411 | 0.409 | 0.433 | 0.501 | 0.728 | 0.723 | 0.740 | 0.386 | 0.314 | 0.374 | 0.413 | 0.753 | 0.296 | 0.369 | 0.588 | 0.75 |
| | | MAE | 0.435 | 0.438 | 0.457 | 0.514 | 0.603 | 0.607 | 0.628 | 0.882 | 0.367 | 0.410 | 0.432 | 0.613 | 0.391 | 0.416 | 0.600 | 0.612 |
| 1 | NLinear | MSE | 0.385 | 0.422 | 0.431 | 0.439 | 0.276 | 0.345 | 0.368 | 0.406 | 0.301 | 0.355 | 0.372 | 0.430 | 0.163 | 0.218 | 0.273 | 0.361 |
| | | MAE | 0.403 | 0.426 | <u>0.429</u> | <u>0.452</u> | 0.338 | 0.382 | 0.408 | 0.441 | 0.343 | 0.379 | 0.385 | 0.418 | 0.252 | 0.290 | 0.326 | 0.382 |
| 0 | TimesNet | MSE | 0.389 | 0.440 | 0.523 | 0.521 | 0.334 | 0.404 | 0.389 | 0.434 | 0.340 | 0.392 | 0.423 | 0.475 | 0.189 | 0.254 | 0.313 | 0.413 |
| | | MAE | 0.412 | 0.443 | 0.487 | 0.495 | 0.370 | 0.413 | 0.435 | 0.448 | 0.378 | 0.404 | 0.426 | 0.453 | 0.265 | 0.310 | 0.345 | 0.402 |
| 0 | FEDformer | MSE | 0.379 | 0.420 | 0.458 | 0.474 | 0.337 | 0.415 | 0.389 | 0.483 | 0.463 | 0.575 | 0.618 | 0.612 | 0.216 | 0.297 | 0.366 | 0.459 |
| | | MAE | 0.419 | 0.444 | 0.466 | 0.488 | 0.380 | 0.428 | 0.457 | 0.488 | 0.463 | 0.516 | 0.544 | 0.551 | 0.309 | 0.360 | 0.400 | 0.450 |
| 0 | Triformer | MSE | 0.399 | 0.444 | 0.492 | 0.549 | 0.936 | 1.290 | 1.325 | 1.500 | 0.349 | 0.387 | 0.426 | 0.482 | 0.276 | 0.473 | 0.692 | 1.936 |
| | | MAE | 0.425 | 0.449 | 0.479 | 0.529 | 0.660 | 0.768 | 0.781 | 0.850 | 0.388 | 0.410 | 0.446 | 0.476 | 0.344 | 0.453 | 0.549 | 0.856 |
| 0 | FiLM | MSE | 0.370 | 0.405 | 0.434 | 0.463 | 0.282 | 0.358 | 0.372 | 0.425 | 0.301 | 0.339 | 0.374 | 0.423 | 0.165 | 0.220 | 0.277 | 0.363 |
| | | MAE | 0.394 | 0.416 | 0.435 | 0.474 | 0.346 | 0.401 | 0.425 | 0.455 | 0.343 | 0.365 | 0.385 | 0.414 | 0.254 | 0.291 | 0.329 | 0.386 |
| 0 | Stationary | MSE | 0.591 | 0.615 | 0.632 | 0.828 | 0.347 | 0.379 | 0.358 | 0.422 | 0.415 | 0.494 | 0.577 | 0.636 | 0.21 | 0.338 | 0.432 | 0.554 |
| | | MAE | 0.524 | 0.540 | 0.551 | 0.658 | 0.387 | 0.418 | 0.413 | 0.457 | 0.410 | 0.451 | 0.490 | 0.535 | 0.294 | 0.373 | 0.416 | 0.476 |

ILI and CzeLan show significant differences in periodicity between variates. The NASDAQ dataset displays no periodicity, whereas the Synthetic dataset demonstrates mixed periodicity. Our analysis shows that these periodic learners struggle to handle periodic heterogeneity effectively. On datasets without clear periodic patterns, such as NASDAQ, PDF and SparseTSF exhibit significant performance errors. In contrast, our model remains robust to periodic heterogeneity.

Table 8: Comparison with periodicity-modeling model. The results are average MAE among all horizon. ↑ is the relative percentage increasing regarding PHAT.

| Datasets | ETTh | ETTm | ILI | CzeLan | NASDAQ | Synthetic |
|-------------|--------------|--------------|--------------|--------------|--------------|--------------|
| TimesNet | 0.438 | ↑9.50% | 0.373 | ↑10.35% | 0.933 | ↑23.08% |
| PDF | 0.408 | ↑2.00% | 0.345 | ↑2.07% | 0.898 | ↑18.46% |
| CycleNet | 0.412 | ↑3.00% | 0.344 | ↑1.77% | 0.971 | ↑28.10% |
| SparseTSF | 0.400 | ↑0.00% | 0.350 | ↑3.55% | 1.556 | ↑105.27% |
| PHAT (Ours) | 0.400 | 0.338 | 0.758 | 0.245 | 0.646 | 0.243 |

D.6 COMPARISON OF TRANSFORMER VARIANTS

We further evaluated the effectiveness of the proposed positive-negative attention mechanism (PNA) by comparing it with other attention variants. As shown in Table 9, we found that Crossformer improves representation capability by modeling the dependencies between variates across channels. However, the interaction of variates with different periodic characteristics may introduce confusion. In contrast, the proposed attention mechanism, PNA, achieved the lowest prediction error. This is because it comprehensively models both the positive and negative correlations of periodicity, demonstrating strong periodic representation capabilities.

Table 9: Comparison with various self-attention mechanisms The results are average MAE among all horizon.

| Datasets | ETTh | ETTm | ILI |
|-----------------------|--------------|--------------|--------------|
| + CrossFormer | 0.441 | 0.380 | 0.742 |
| + Vanilla Transformer | 0.424 | 0.366 | 0.768 |
| + PatchFormer | 0.435 | 0.373 | 0.825 |
| PHAT (Ours) | 0.400 | 0.338 | 0.758 |

D.7 DETAILED ABLATION STUDY

In this section, we conduct combined ablation studies across additional datasets and provide further analysis. Our results show that jointly modeling both positive and negative periodic correlation is beneficial, as it yields a more complete representation of the underlying periodic characteristic. We also observe that using only positive periodicity performs slightly better than using only negative correlations, likely because positive correlations capture the most direct and dominant periodic dynamics. For datasets where negative periodic correlations are weak or largely absent (e.g., ETTh2, ETTm2, and AQShunyi), whether or not negative paths are included—or whether their attention logits are modulated—has minimal impact on performance. This behavior arises naturally from the gating mechanism Λ , which adaptively determines the extent to which negative correlations should be modeled based on dataset characteristics. As a result, the model avoids unnecessary overfitting when negative periodicity is not present.

Table 10: Details of combination Ablation Study of the period-offset attention in PNA. Pos (Neg): Positive (Negative) components. Mod: Modulation.

| Variants | | ETTh1 | | ETTh2 | | ETTm1 | | ETTm2 | | NN5 | | FRED-MD | | AQShunyi | | ILI | | CzeLan | | NASDAQ | |
|----------|------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|---------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Pos. | Neg. | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE |
| - | - | 0.369 | 0.390 | 0.275 | 0.336 | 0.305 | 0.340 | 0.169 | 0.253 | 0.688 | 0.552 | 25.698 | 0.980 | 0.666 | 0.468 | 1.551 | 0.775 | 0.185 | 0.222 | 0.440 | 0.490 |
| - | ✓ | 0.368 | 0.386 | 0.273 | 0.331 | 0.302 | 0.339 | 0.168 | 0.254 | 0.688 | 0.551 | 25.305 | 0.954 | 0.665 | 0.468 | 1.531 | 0.765 | 0.176 | 0.221 | 0.433 | 0.483 |
| ✓ | - | 0.363 | 0.384 | 0.272 | 0.330 | 0.297 | 0.332 | 0.167 | 0.253 | 0.684 | 0.553 | 25.520 | 0.967 | 0.665 | 0.468 | 1.485 | 0.759 | 0.175 | 0.215 | 0.427 | 0.476 |
| ✓ | ✓ | 0.361 | 0.383 | 0.272 | 0.330 | 0.287 | 0.330 | 0.162 | 0.246 | 0.681 | 0.551 | 21.495 | 0.903 | 0.665 | 0.468 | 1.318 | 0.705 | 0.170 | 0.204 | 0.416 | 0.473 |
| PMod. | | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE |
| - | - | 0.368 | 0.390 | 0.275 | 0.334 | 0.309 | 0.345 | 0.166 | 0.243 | 0.688 | 0.554 | 23.869 | 0.942 | 0.669 | 0.468 | 1.560 | 0.774 | 0.177 | 0.284 | 0.420 | 0.479 |
| - | ✓ | 0.365 | 0.386 | 0.272 | 0.330 | 0.305 | 0.339 | 0.167 | 0.242 | 0.685 | 0.553 | 24.740 | 0.943 | 0.666 | 0.468 | 1.535 | 0.764 | 0.174 | 0.252 | 0.419 | 0.478 |
| ✓ | - | 0.362 | 0.385 | 0.272 | 0.330 | 0.295 | 0.334 | 0.166 | 0.243 | 0.685 | 0.554 | 23.876 | 0.907 | 0.665 | 0.468 | 1.446 | 0.750 | 0.177 | 0.234 | 0.416 | 0.479 |
| ✓ | ✓ | 0.361 | 0.383 | 0.272 | 0.330 | 0.287 | 0.330 | 0.162 | 0.246 | 0.681 | 0.551 | 21.495 | 0.903 | 0.665 | 0.468 | 1.318 | 0.705 | 0.170 | 0.204 | 0.416 | 0.473 |

E DISCUSSIONS

Although both our model and patch-based Transformers rely on operations such as unflattening or folding the sequence, their design philosophies diverge significantly.

Patch-based methods, such as PatchTST (Zeng et al., 2023), Crossformer (Zhang & Yan, 2023), and PDF (Dai et al., 2024), divide the temporal sequence into contiguous segments, reinterpreting one dimension of each patch as a feature axis. This effectively performs temporal down-sampling by aggregating multiple time steps into a single token. While this approach improves computational efficiency, it comes at the cost of mixing distinct temporal roles, leading to the loss of fine-grained alignment among the original time steps.

In contrast, our period folding approach preserves the full temporal resolution by avoiding any form of down-sampling. Instead, the sequence is reorganized into two orthogonal axes—period-offset and period-aligned—based on the detected periodic structure. Attention is applied along both axes, enabling the model to explicitly capture rich intra-period and inter-period dependencies while maintaining the integrity of the original temporal sequence.

A related idea can be observed in TimesNet (Wu et al., 2022a), which also employs a folding operation. However, TimesNet extracts features from the folded representation using convolutional kernels. These kernels aggregate information across multiple positions, which can blur inherent periodic boundaries and disrupt latent cyclic correlations. This misalignment between convolutional receptive fields and true periodic structures limits TimesNet’s ability to fully capture cycle-aware dependencies. In contrast, our model applies X-shaped attention directly to the folded representation, preserving

periodic fidelity and ensuring that the periodic structure is consistently leveraged throughout the learning process.

In summary, our fold-and-attend approach stands apart by preserving the full temporal sequence and explicitly structuring it along periodic axes. Unlike patch-based Transformers, which sacrifice temporal resolution for efficiency, or TimesNet, which disrupts periodic integrity with convolutional receptive fields, our design respects periodicity at its core. This allows our model to generalize effectively across diverse temporal patterns while offering enhanced interpretability in capturing heterogeneous periodic dependencies.

F USE OF LLM

We only use LLMs as a language optimization tool to polish sentences, improving their readability and fluency.