

FreqCycle: A Multi-Scale Time-Frequency Analysis Method for Time Series Forecasting

Boya Zhang, Shuaijie Yin, Huiwen Zhu, Xing He*

Shanghai Jiao Tong University, Shanghai, China
{zby112200, yinshuaijie, shu_yuan, xhe}@sjtu.edu.cn

Abstract

Mining time-frequency features is critical for time series forecasting. Existing research has predominantly focused on modeling low-frequency patterns, where most time series energy is concentrated. The overlooking of mid to high frequency continues to limit further performance gains in deep learning models. We propose **FreqCycle**, a novel framework integrating: (i) a **Filter-Enhanced Cycle Forecasting (FECF)** module to extract low-frequency features by explicitly learning shared periodic patterns in the time domain, and (ii) a **Segmented Frequency-domain Pattern Learning (SFPL)** module to enhance mid to high frequency energy proportion via learnable filters and adaptive weighting. Furthermore, time series data often exhibit coupled multi-periodicity, such as intertwined weekly and daily cycles. To address coupled multi-periodicity as well as long lookback window challenges, we extend FreqCycle hierarchically into **MFreqCycle**, which decouples nested periodic features through cross-scale interactions. Extensive experiments on seven diverse domain benchmarks demonstrate that FreqCycle achieves state-of-the-art accuracy while maintaining faster inference speeds, striking an optimal balance between performance and efficiency.

Code — <https://github.com/boya-zhang-ai/FreqCycle>

1 Introduction

Time series forecasting (TSF) plays a pivotal role in numerous domains, including energy optimization (Wang et al. 2021), weather forecasting (Wu et al. 2023; Fan et al. 2023), and economic analysis (Ariyo, Adewumi, and Ayo 2014), where it serves as a priori guidance for proactive planning and decision-making. The ubiquity and significance of TSF tasks have spurred a proliferation of methodological advances in this field.

Traditional forecasting approaches, such as ARIMA (Ariyo, Adewumi, and Ayo 2014) and LSTM (Jerome, Ishii, and Chen 2018), were predominantly employed in early TSF research.

The rapid development of deep learning has given rise to a diverse array of neural network based forecasting models. Recurrent Neural Network, RNN-based methods (e.g.,




Figure 1: Shared periodic patterns present in the ETTm2 dataset.

SegRNN (Lin et al. 2023), WITRAN (Jia et al. 2023)) are frequently hampered by gradient vanishing or explosion issues, which limit their ability to model long-term dependencies. Convolutional Neural Network, CNN-based approaches (e.g., TCN (Bai, Kolter, and Koltun 2018), SCINet (Liu et al. 2022)) leverage parallel computation but often exhibit restricted receptive fields, potentially overlooking global temporal relationships. Transformer-based techniques (e.g., Autoformer (Wu et al. 2021), PatchTST (Nie et al. 2023), iTransformer (Liu et al. 2023)) excel at modeling global dependencies through self-attention mechanisms, albeit at the cost of high computational complexity. Meanwhile, Multilayer Perceptron MLP-based models (e.g., DLinear (Zeng et al. 2023), CycleNet (Lin et al. 2024a)) offer simplicity and training efficiency, demonstrating notable efficacy in low-frequency component prediction.

MLP-based methods have garnered significant attention due to their low complexity and competitive performance. However, current MLP-based approaches exhibit several limitations: (i) Deep learning models aim to construct complex architectures to extract periodic patterns from long-range dependencies. However, we opt for a more direct and effective approach by explicitly learning shared daily and weekly periodic patterns through parametric modeling. Figure 1 shows the daily periodic patterns in the ETT dataset and the observable weekly patterns after average pooling. (ii) From a frequency-domain perspective, although simple models (e.g., DLinear (Zeng et al. 2023)) excel at learning low-frequency periodic components, their fundamental ar-

*Corresponding Author

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




Figure 2: Frequency spectrum of Traffic dataset before and after SFPL enhancement

chitecture—built upon Linear Layers or MLPs—inherently limits their ability to effectively represent mid to high frequency components, which capture short-term fluctuations and aperiodic features. As shown in Figure 2, these mid to high frequency components contribute very little energy proportion, urgently requiring effective methods to enhance the energy contribution of these critical frequency components. (iii) Existing architectures are not specifically designed to address the multi-scale nested periodicity (such as coupled daily-weekly cycles) inherent in time series data. Naively employing MLP-based models could result in compromised performance or reduced efficiency, thereby demanding the construction of an improved framework that can adaptively model multi-level periodic structures in a coordinated manner.

To address the challenges in time series forecasting, this paper makes the following key contributions:

- We introduce Filter-Enhanced Cycle Forecasting (FECF), a method that explicitly learns shared periodic patterns, avoiding complex architectures to learn long-range dependencies.
- We propose the Segmented Frequency Pattern Learning (SFPL) method, which enhances the energy contribution of mid to high frequency components. This approach effectively resolves the long-standing limitation of insufficient key frequency modeling in conventional methods, improving prediction accuracy for aperiodic elements.
- To address the critical challenge of coupled multi-periodicity under long lookback windows, we developed MFreqCycle - an architecture that decouples and jointly models nested periodic features through a multi-scale approach.
- We conduct extensive experiments on seven time series forecasting benchmarks. The results demonstrate that our model achieves superior performance in both effectiveness and efficiency.

2 Related Work

MLP-based Time Series Model

Several studies have investigated the application of MLP-based deep learning networks for time series forecasting. N-BEATS (Oreshkin et al. 2019) employs stacked MLP lay-

ers with dual residual learning to process input data and generate iterative future predictions. LightTS (Zhang et al. 2022) adopts a lightweight sampling-oriented MLP architecture to reduce complexity and computational overhead while maintaining prediction accuracy. DLinear (Zeng et al. 2023) introduces a set of simple single-layer Linear models to capture temporal relationships between input and output sequences. SparseTSF (Lin et al. 2024b) utilizes cross-cycle sparse prediction techniques to decouple periodic and trend components. CycleNet (Lin et al. 2024a) demonstrates remarkable performance in day-ahead forecasting by explicitly modeling periodic patterns through an integrated framework combining cyclical decomposition and MLP-based residual prediction. TimeMixer (Wang et al. 2024a) introduces a novel architecture featuring Past-Decomposable-Mixing for multiscale representation learning and Future-Multipredictor-Mixing to ensemble multi-scale forecasting skills. These studies collectively validate the effectiveness of MLP-based architectures in time series forecasting tasks, motivating our further development of MLP-based deep learning prediction algorithms in this work.

Time Series Modeling with Frequency Learning

Recent years have witnessed increasing integration of frequency techniques into deep learning models for time series forecasting, capitalizing on their computational efficiency and energy compaction properties (Yi et al. 2023a). Specifically, Autoformer (Wu et al. 2021) replaces self-attention with an autocorrelation mechanism incorporating Fast Fourier Transform (FFT), while FEDformer (Zhou et al. 2022b) proposes a DFT-based frequency-enhanced attention that computes attention weights through query-key spectral analysis. FiLM (Zhou et al. 2022a) employs Fourier analysis to preserve historical patterns while filtering noise, and FreTS (Yi et al. 2023b) introduces frequency-domain MLPs to learn cross-channel and temporal dependencies. FITS (Xu, Zeng, and Xu 2024) applies low-pass filtering before performing complex-valued Linear transformations in the frequency domain. FilterNet (Yi et al. 2024) introduces plain shaping filters for baseline frequency processing and contextual shaping filters for adaptive spectrum modulation. Notably, most existing frequency-based methods primarily focus on low-frequency components through conventional or adaptive filtering. In contrast, our work proposes a novel frequency learning architecture that specifically enhances mid to high frequency energy representation.

3 Methodology

In this section, we first formalize the TSF problem and introduce the time-frequency transformation theory along with its inherent time-frequency duality, thereby establishing the theoretical foundation for our methodology. We then present the two core components of FreqCycle: (i) the Filter-Enhanced Cycle Forecasting (FECF) technique, and (ii) the Segmented Frequency-domain Pattern Learning (SFPL) module. To address the critical challenge of coupled periodic patterns in real-world time series, we further develop MFreqCycle by incorporating cross-scale feature fusion mechanisms into the FreqCycle architecture, enabling decoupling

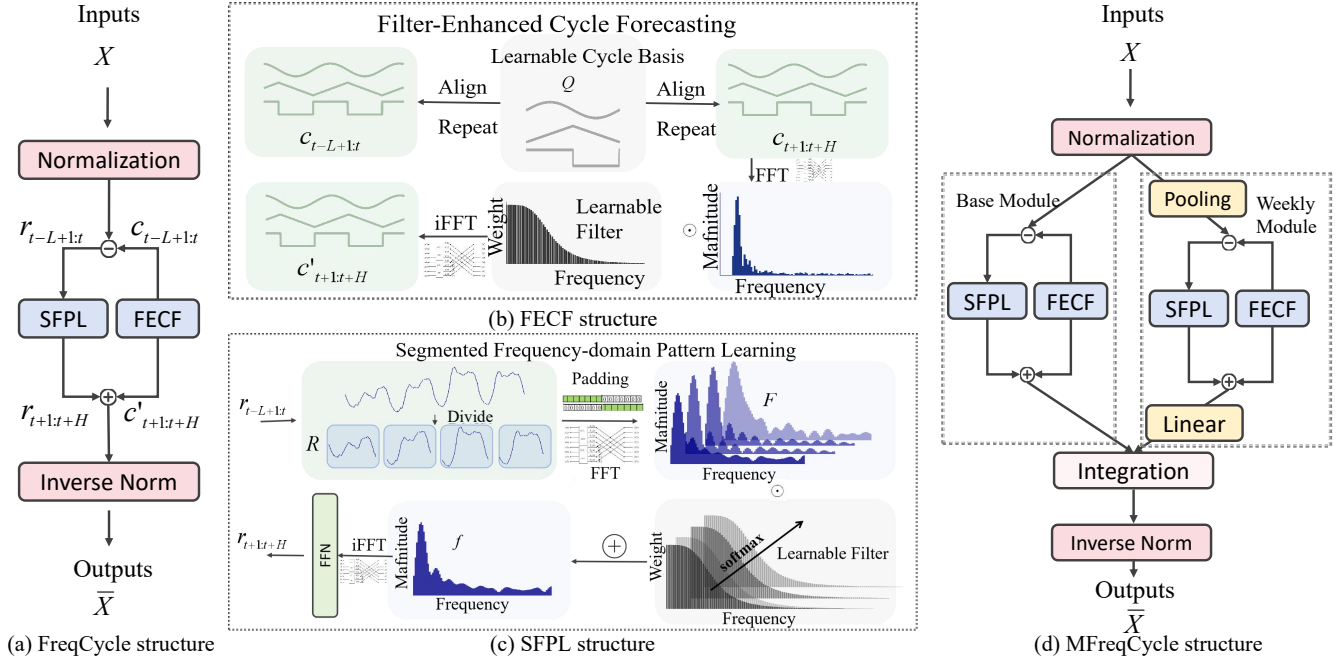


Figure 3: The architecture of FreqCycle. (a) The overall architecture of FreqCycle. (b) FECF block is employed to explicitly learn shared periodic patterns in the time domain; (c) SFPL extracts critical frequency information in the spectral domain via learnable filters and adaptive weighting. (d) MFreqCycle explicitly decouples and models nested periodic features through cross-scale interactions.

and joint modeling of nested periodic features. The overall structure is illustrated in Figure 3.

Preliminaries

1) Problem Definition: Given a time series $X \in \mathbb{R}^{L \times D}$ with L historical observations and D variables, the goal is to predict future values $\bar{X} \in \mathbb{R}^{H \times D}$ where H is the forecast horizon. The mapping function can be formalized as:

$$f : x_{t-L+1:t} \in \mathbb{R}^{L \times D} \rightarrow \bar{x}_{t+1:t+H} \in \mathbb{R}^{H \times D}$$

2) Discrete Fourier Transform: In TSF problems, time series data are discrete samples that vary over time. The Discrete Fourier Transform (DFT) is a widely used method for converting discrete data from the time domain to the frequency domain. For a discrete sequence $x[n]$ of length N , its DFT is defined as:

$$X[k] = \sum_{n=0}^{N-1} x[n] e^{-j \frac{2\pi}{N} kn}, \quad k = 0, 1, \dots, N-1 \quad (1)$$

where $j = \sqrt{-1}$ represents the imaginary unit.

The Inverse Discrete Fourier Transform (iDFT) is the reverse operation of DFT, which reconstructs the time-domain signal from its frequency-domain representation:

$$x[n] = \frac{1}{N} \sum_{k=0}^{N-1} X[k] e^{j \frac{2\pi}{N} kn}, \quad n = 0, 1, \dots, N-1 \quad (2)$$

In practical applications, Fast Fourier Transform (FFT) and inverse Fast Fourier Transform (iFFT) are commonly

employed instead of DFT and iDFT to significantly reduce computational complexity and improve operational efficiency. This efficient time-frequency transformation serves as the fundamental basis for implementing various frequency-domain operations on time series data.

3) Duality of Time and Frequency Domain: For two discrete-time sequences $x[n]$ and $h[n]$, the convolution operation in the time domain is computed as follows:

$$y[n] = (x * h)[n] = \sum_{m=0}^{N-1} x[m] h[n-m] \quad (3)$$

The convolution operation in the time domain is equivalent to point-wise multiplication in the frequency domain:

$$\mathcal{F}\{x * h\} = X[k] \cdot H[k] \quad (4)$$

Conversely, multiplication in the time domain corresponds to convolution in the frequency domain. Based on the time-frequency duality property between convolution and multiplication, numerous researchers (Zhao et al. 2024; Yi et al. 2024) have adopted frequency-domain pointwise multiplication operations to replace time-domain circular convolution for enhanced computational efficiency.

Another industrially significant time-frequency duality manifests in the deterministic relationship between temporal sampling interval Δt and signal bandwidth f_{\max} , as well as the inverse dependence of frequency resolution Δf on total observation duration T .

Filter-Enhanced Cycle Forecasting

The Filter-Enhanced Cyclical Forecasting (FECF) technique represents an improvement over the original Residual Cycle Forecasting framework (Lin et al. 2024a), specifically designed to enhance the learning of low-frequency periodic patterns in time series data.

Given input sequences with D channels and a predefined base period length W , the method first generates a learnable cyclical basis $Q \in \mathbb{R}^{W \times D}$ (initialized as a zero matrix) that is globally shared across the temporal dimension. Through periodic replication operations (Lin et al. 2024a), this basis produces a cyclical component $c_{t-L+1:t}$ matching the original sequence length L , as well as a cyclical component $c_{t+1:t+H}$ matching the predict length H . To prioritize low-frequency information capture, the generated cyclical component $c_{t+1:t+H}$ undergoes processing via adaptive filtering. Specifically, the adaptive filtering enhanced periodic learning module focuses on amplifying low-frequency information while attenuating mid- and high-frequency components that may degrade periodic pattern learning performance. The workflow is implemented as follows:

$$\begin{aligned} c'_{t+1:t+H} &= \text{IFFT}(\text{Filter}(\text{FFT}(c_{t+1:t+H}))) \\ \text{Filter}(\xi) &= \xi \odot \theta_c \end{aligned} \quad (5)$$

where $\text{FFT}(\cdot)$ denotes the Fast Fourier Transform along the temporal dimension, $\text{IFFT}(\cdot)$ represents the Inverse Fast Fourier Transform, θ_c corresponds to the learnable filter parameters for the concatenated periodic components, and \odot indicates element-wise multiplication (Hadamard product).

Both the periodic basis Q and the adaptive filters are jointly trained through gradient backpropagation to effectively extract the underlying low-frequency periodic patterns.

Notably, the base period length W is determined by the inherent characteristics of the dataset. For time series with 1-hour intervals, $W = 24$ for daily cycles or $W = 168$ for weekly patterns. This architecture combines explicit periodic modeling with frequency-aware filtering, demonstrating superior performance in capturing dominant cyclical trends while maintaining computational efficiency through its parameter-sharing mechanism.

In conclusion, FECF procedure is as follow:

1) Extract Residual Component: Remove the periodic component $c_{t-L+1:t}$ from the input $x_{t-L+1:t}$ obtaining the residual component:

$$r_{t-L+1:t} = x_{t-L+1:t} - c_{t-L+1:t} \quad (6)$$

2) Residual Prediction via SFPL Module: Feed the residual component into the SFPL module to predict future residuals:

$$r_{t+1:t+H} = \text{SFPL}(r_{t-L+1:t}) \quad (7)$$

3) Reconstruct Prediction: Combine the predicted residual with the filtered periodic component:

$$\bar{x}_{t+1:t+H} = r_{t+1:t+H} + c'_{t+1:t+H} \quad (8)$$

Segmented Frequency-domain Pattern Learning

The SFPL technique is specifically designed to enhance the energy contribution of mid to high frequency components, enabling the prediction of aperiodic components and short-term fluctuations. Within the FreqCycle framework, SFPL operates on residual components $r_{t-L+1:t}$ obtained after removing low-frequency periodic elements from the input. SFPL procedure is as follow:

1) The input $r_{t-L+1:t}$ is divided into s sub-segments via a sliding window, with zero-padding applied to both ends of each sub-segment. These processed sub-segments are then stacked to obtain segmented time feature $R = \{r_1, r_2, \dots, r_s\} \in \mathbb{R}^{s \times L \times D}$.

2) Then FFT is applied to R along the temporal dimension to obtain stacked frequency representation $F \in \mathbb{C}^{s \times \lfloor \frac{L}{2} + 1 \rfloor \times D}$, after which frequency-domain feature extraction is performed on F . To consolidate critical frequency-domain information, an adaptive weighted summation is conducted across the sub-segments in F , ultimately yielding the refined frequency-domain data $f \in \mathbb{C}^{\lfloor \frac{L}{2} + 1 \rfloor \times D}$ with enhanced features.

$$\begin{aligned} \Theta_F &= \{\theta_1, \dots, \theta_s\} \\ \theta'_1, \dots, \theta'_s &= \text{softmax}(\theta_1, \dots, \theta_s) \\ \Theta'_F &= \{\theta'_1, \dots, \theta'_s\} \\ F' &= F \odot \Theta'_F \\ f &= \sum_{i=1}^s F'(i) \end{aligned} \quad (9)$$

where Θ_F corresponds to the learnable parameters for the segmented frequency components. $F'(i)$ referring to each sub-element of F along the segmented dimension.

3) The refined data f undergoes iFFT and is then passed through FFN layer to produce the residual prediction output $r_{t+1:t+H}$ of SFPL.

The segmentation operation is an important technique in SFPL. By dividing the signal into short sub-segments and performing FFT on each segment individually, it essentially constitutes a variant of the Short-Time Fourier Transform (STFT). Each sub-segment achieves higher frequency-domain locality by shortening the time-domain window (Nawab, Quatieri, and Lim 1983), thereby enabling more precise localization of transient frequency components. The resulting frequency-domain representation f , obtained through learnable filtering and adaptive weighting integration, becomes more discriminative with enhanced mid to high frequency energy. We provide a more detailed introduction to the theory of STFT and its strategies for enhancing mid to high frequency energy proportion in Appendix A.

MFreqCycle

MFreqCycle is a multi-scale parallel framework that achieves high-precision forecasting by collaboratively extracting temporal features across different time scales. The features include periodic patterns (such as daily and weekly cycles) and non-periodic components. Current experiments have demonstrated the feasibility of incorporating weekly cycles on several datasets, while longer-period patterns like

Models	FreqCycle		DLinear		CycleNet		Amplifier		FreTS		FilterNet		FITS		iTransformer		PatchTST		TimeMixer	
Metrics	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTm1	0.372	0.389	0.403	0.407	0.386	0.395	<u>0.381</u>	<u>0.394</u>	0.408	0.416	0.384	0.398	0.415	0.408	0.407	0.410	0.387	0.400	0.381	0.395
ETTm2	0.263	0.311	0.350	0.401	<u>0.272</u>	<u>0.315</u>	0.276	0.323	0.321	0.368	0.276	0.322	0.286	0.328	0.288	0.332	0.281	0.326	0.275	0.323
ETTh1	0.428	0.427	0.456	0.452	0.432	0.427	<u>0.430</u>	<u>0.428</u>	0.475	0.463	0.440	0.432	0.451	0.440	0.454	0.448	0.469	0.455	0.447	0.440
ETTh2	0.371	0.399	0.559	0.515	0.383	0.404	0.359	0.391	0.472	0.465	0.378	0.404	0.383	0.408	0.383	0.407	0.387	0.407	<u>0.364</u>	<u>0.395</u>
Weather	<u>0.243</u>	0.270	0.265	0.317	0.254	0.279	0.243	<u>0.271</u>	0.250	0.270	0.245	0.272	0.249	0.276	0.258	0.278	0.259	0.281	0.240	0.271
ECL	0.168	0.259	0.212	0.300	<u>0.170</u>	<u>0.260</u>	0.171	0.265	0.189	0.278	0.201	0.285	0.217	0.295	0.178	0.270	0.205	0.290	0.182	0.272
Traffic	<u>0.448</u>	0.261	0.625	0.383	0.485	0.313	0.482	0.315	0.618	0.390	0.521	0.340	0.627	0.376	0.428	<u>0.282</u>	0.481	0.304	0.484	0.297

Table 1: Time series forecasting comparison. We set the lookback window size L as 96 and the prediction length as $H \in \{96, 192, 336, 720\}$. The best results are highlighted in bold and the second best are underlined. Results are averaged from all prediction lengths. Full results for all datasets are listed in Table 2 of Appendix B.

annual cycles remain unverified due to dataset limitations. The architecture comprises two key modules: a base cycle learning and frequency enhancement module, which focuses on capturing the smallest significant cycle (e.g., daily cycle) and its associated non-periodic features, and a weekly cycle learning and frequency enhancement module, which models macro-scale periodic patterns (e.g., weekly cycle). These multi-scale predictions are then adaptively integrated through a prediction fusion layer. To improve computational efficiency, the framework employs different input window lengths when extracting temporal features based on base-cycle and weekly-cycle durations.

To capture macro-scale periodic patterns, the system requires longer input sequences spanning, which significantly increases computational demands. However, for short-term forecasting, the model only needs to extract key trend information from data. To balance computational efficiency with model performance, the weekly cycle learning and frequency enhancement module incorporates three submodules: a pooling layer, a feature learning module, and a Linear layer. The distinctive feature extraction architecture based on pooling and Linear projection represents the key difference between the weekly cycle module and the base cycle module, while the feature learning submodule follows the same processing pipeline as FreqCycle. This design enables efficient extraction of essential weekly patterns without compromising predictive accuracy.

After processing through the base and weekly modules, the model generates multi-scale prediction outputs including the base period output $\bar{x}_{t+1:t+H}^0$ and weekly period output $\bar{x}_{t+1:t+H}^1$. Then the prediction fusion module performs adaptive integration of these multi-scale predictions to produce the final system output $\bar{x}_{t+1:t+H}$.

$$\begin{aligned} \theta'_0, \theta'_1 &= \text{softmax}(\theta_0, \theta_1) \\ \bar{x}_{t+1:t+H} &= \bar{x}_{t+1:t+H}^0 \odot \theta'_0 + \bar{x}_{t+1:t+H}^1 \odot \theta'_1 \end{aligned} \quad (10)$$

where θ'_0, θ'_1 corresponds to the learnable parameters for the prediction fusion module.

This fusion strategy ensures an optimal balance between different periodic features while maintaining computational efficiency and prediction accuracy.

4 Experiments

Experimental Setup

Datasets To validate the model’s performance and generalization capability, we conduct experiments on seven datasets, including publicly benchmarks like ETT dataset series (Zhou et al. 2021), Electricity (ECL), Weather and Traffic datasets (Wu et al. 2021). These datasets span diverse domains with varying scales and sampling frequencies, providing a comprehensive testbed to evaluate the model’s adaptability and scalability across different forecasting tasks. Detailed dataset information is summarized in Appendix B.

Technical Implementation Details All experiments were implemented in PyTorch 2.6.0 and conducted on an NVIDIA RTX 4090 GPU with 24GB of memory.

Baselines In this work, we benchmark our model against the representative and state-of-the-art approaches. We are evaluating and comparing with Linear and MLP-based models, including DLinear (Zeng et al. 2023), CycleNet (Lin et al. 2024a), and TimeMixer (Wang et al. 2024a). We further evaluate against frequency learning methods, including FreTS (Yi et al. 2023b), FilterNet (Yi et al. 2024), FITS (Xu, Zeng, and Xu 2024), and Amplifier (Fei et al. 2025). To demonstrate the effectiveness of our model, we conduct comprehensive comparisons with Transformer-based models, including iTransformer (Liu et al. 2023) and PatchTST (Nie et al. 2023).

Evaluation Metrics Following established practices in prior work, we employ two primary metrics to assess forecasting performance: Mean Squared Error (MSE) and Mean Absolute Error (MAE).

Main Results

Table 1 presents the training results of the FreqCycle model and nine baselines on seven open-source datasets. Overall, FreqCycle secured first place in 10 out of 14 metrics and second place in 2 metrics. Compared to other models, our model demonstrates significant performance improvements, which we attribute to the FECF’s learning of periodic patterns and the SFPL’s enhancement of key frequencies.

Datasets	L	FreqCycle		CycleNet		FITS	
		MSE	MAE	MSE	MAE	MSE	MAE
ETTh2	96	0.371	0.399	0.383	0.404	0.383	0.408
ETTh2	96	0.263	0.311	0.267	0.314	0.286	0.328
ETTm2	168	0.241	0.340	0.388	0.415	0.333	0.382
ETTm2	672	0.167	0.279	0.262	0.325	0.250	0.312

Table 2: Performance comparison with different lookback window lengths. The prediction length is $H \in \{96, 192, 336, 720\}$. The best results are highlighted in bold. Results are averaged across all prediction lengths. Full results are available in Table 7 of Appendix D.

To evaluate the effectiveness of the MFreqCycle model, we design Table 2 with lookback windows of 96 and a whole week length data 672(168), comparing the model against CycleNet and FITS. As shown in Table 2, simply increasing the lookback window may lead to minor performance improvements or even slight degradation in some cases. The reason could be that the inherent limitations of the models prevent them from handling the complex temporal dependencies or noise interference introduced by longer lookback windows. However, the MFreqCycle model, specifically designed for longer lookback windows, achieves remarkable performance improvements on the ETT datasets, demonstrating its pioneering capability in learning both long-period patterns and their aperiodic components.

Model Analysis

Ablation Study To validate the effectiveness of the FECF and SFPL modules in FreqCycle, we designed the ablation experiments shown in Table 3. We compared the performance of: 1) the complete FreqCycle framework, 2) the variant without the FECF module, and 3) the variant replacing SFPL with MLP. The results show that both modules significantly contribute to improving prediction performance. Specifically, the FECF module achieves the most substantial improvements on Traffic and Weather datasets with prominent cycles, showing a 21.30% MAE improvement on the Traffic dataset. The SFPL module exhibits notable performance gains on ETTh datasets, which exhibit stronger non-stationary characteristics and richer transient frequency components, achieving a 6.91% MSE improvement. SFPL and FECF demonstrate complementary performance characteristics, and their synergistic integration enables the FreqCycle model to achieve superior predictive performance.

Effectiveness of SFPL In recent years, researchers have introduced signal processing techniques, like low-pass filters (LPFs), into time series forecasting and developed learnable filter schemes, aiming to extract more critical frequency components. Unlike conventional LPFs with rigid cutoff frequencies or purely gradient-based learnable filters, our approach draws inspiration from STFT to enhance the extraction of crucial mid to high frequency information.

We conduct a systematic comparison between our SFPL

Datasets	ETTh1		ETTh2		Traffic		Weather	
	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
FreqCycle	0.369	0.390	0.282	0.336	0.438	0.262	0.159	0.203
No FECF	0.375	0.397	0.292	0.344	0.501	0.318	0.179	0.218
Boost	1.65%	1.82%	3.26%	2.56%	14.54%	21.30%	12.59%	7.29%
2MLP	0.379	0.401	0.302	0.348	0.444	0.262	0.161	0.206
Boost	2.66%	2.93%	6.91%	3.85%	1.44%	0.04%	1.51%	1.28%

Table 3: Ablation study of FreqCycle module, using parameters L=96 and H=96. 2MLP means the variant replacing SFPL with MLP.

Datasets	ETTh2		ECL		Traffic	
	MSE	MAE	MSE	MAE	MSE	MAE
FreqCycle	0.282	0.336	0.139	0.233	0.438	0.262
SFPL2LPF	0.294	0.345	0.143	0.240	0.449	0.268
Boost	4.00%	2.68%	3.39%	3.09%	2.63%	2.29%
SFPL2PaiFilter	0.293	0.344	0.143	0.239	0.445	0.262
Boost	3.68%	2.59%	2.96%	2.58%	1.81%	0.08%

Table 4: Comparison of frequency processing methods, using parameters L=96 and H=96.

module and existing frequency processing methods within the FreqCycle framework. We evaluate: 1) replacing SFPL’s frequency processing with PaiFilter (Yi et al. 2024), and 2) substituting SFPL with the LPF scheme from FITS (Xu, Zeng, and Xu 2024). As shown in Table 4, these results indicate that existing frequency-domain filtering approaches fail to adequately capture critical frequency components, whereas our SFPL module effectively extracts frequency features.

Effectiveness of FECF Our FECF technique fundamentally operates as a Seasonal-Trend Decomposition (STD) method. To validate its effectiveness, we conducted experiments comparing it against two established STD approaches: the moving average decomposition (MOV) used in DLinear (Zeng et al. 2023) and the local decomposition (LD) (Yu et al. 2024). As demonstrated in Table 5, when replacing FECF with either MOV+Linear (DFreq) or LD+Linear (LDFreq) variants in our FreqCycle framework, the forecasting performance consistently degraded across all seven benchmark datasets. These empirical results clearly establish that our FECF-based decomposition achieves superior temporal pattern extraction compared to conventional STD methods.

Efficiency Analysis The theoretical complexity of FreqCycle is $O(L \log L)$. We conducted a comprehensive comparison with representative models across three dimensions, Forecasting performance, Peak memory consumption and Training speed. Figure 4 demonstrates that FreqCycle achieves the optimal balance between model performance and computational efficiency. This efficiency advantage stems from our time and frequency processing paradigm, which avoids the quadratic complexity of atten-

Datasets	FreqCycle		DFreq		LDFreq	
	MSE	MAE	MSE	MAE	MSE	MAE
ETTm2	0.263	0.311	0.312	0.365	0.329	0.368
Weather	0.243	0.270	0.259	0.309	0.248	0.291
ECL	0.168	0.259	0.204	0.293	0.201	0.288
Traffic	0.448	0.261	0.529	0.316	0.527	0.307

Table 5: Performance comparison of different STD methods. The configuration look-back window length of 96 and The reported results are averaged across all prediction horizons of $H \in \{96, 192, 336, 720\}$, with full results available in Table 5 of Appendix D.

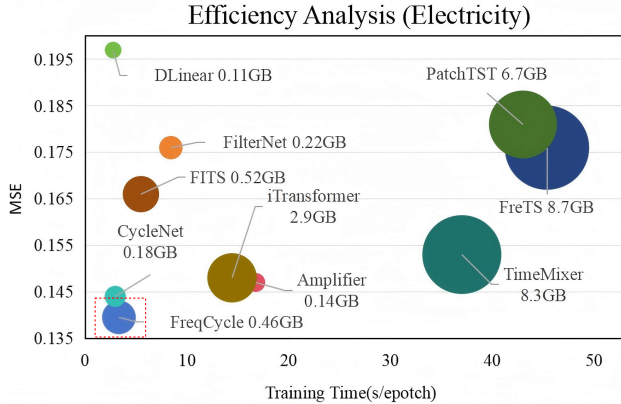


Figure 4: Model efficiency comparison in terms of MSE, Peak memory consumption and Training speed. Using an input length $L = 96$ and prediction horizon $H = 96$ on the Electricity dataset.

tion mechanisms while preserving temporal pattern learning capabilities.

Visualization

Visualization of SFPL’s Key Frequency Enhancement

After training the model on benchmark datasets, we analyzed the frequency amplitude spectra before and after SFPL processing. As demonstrated in Figure 5, SFPL effectively Amplifies the energy proportion of mid to high frequency bands as well as Preserves the structural integrity of dominant low-frequency components.

Visualization of Periodic Patterns Extracted by FECF

Figure 6 shows the learned periodic patterns across multiple datasets after hyperparameter W optimization. Our analysis reveals that for the Electricity and Traffic datasets, FreqCycle successfully captures weekly periodicity patterns through hyperparameter tuning. For the remaining datasets, the optimized FreqCycle configuration learns daily cycle patterns.

Visualization of Forecasting Results The prediction results on benchmark datasets are available in Appendix E.




Figure 5: Frequency spectrum of Electricity dataset before and after SFPL enhancement. The other datasets’ results are available in Appendix E.

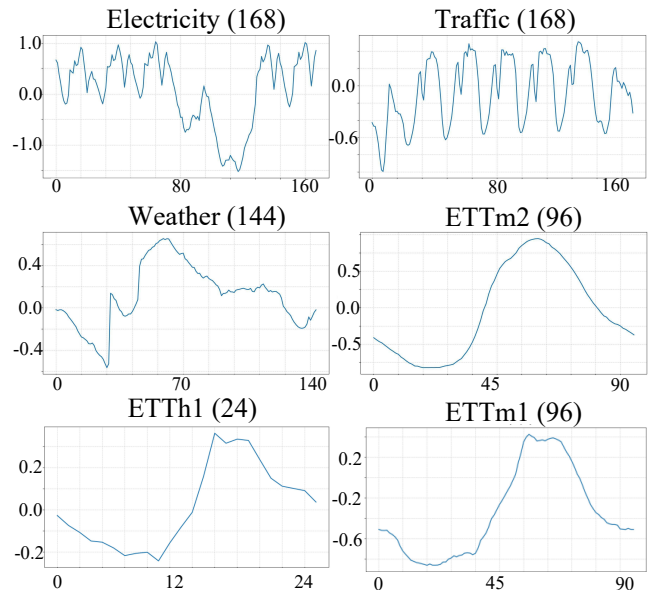


Figure 6: Learned periodic patterns by conducting hyperparameter W tuning on each dataset.

5 Conclusion

This paper presents FreqCycle, a frequency aware framework for TSF that addresses the under explored challenge of modeling mid to high frequency components. FreqCycle integrates two key innovations, FECF for periodic pattern extraction, and SFPL to amplify mid to high frequency energy via learnable filters and adaptive weighting. For coupled multi-scale periodicity, we further propose MFreqCycle, which achieves remarkable performance improvements on several datasets while maintaining efficiency. Comprehensive experiments across seven real world datasets demonstrate both the superior accuracy and computational efficiency of our proposed approach. The FreqCycle framework establishes a new paradigm for time series forecasting that achieves state-of-the-art performance while maintaining lightweight architecture.

Acknowledgments

We would like to express our sincere gratitude to our advisor, Associate Professor Xing He, for his continuous guidance and invaluable support throughout this research. We also thank all members of the Intelligent Information Processing and Control Laboratory (IIC) at Shanghai Jiao Tong University for their insightful discussions and kind assistance during the development of this work.

References

- Ariyo, A. A.; Adewumi, A. O.; and Ayo, C. K. 2014. Stock Price Prediction Using the ARIMA Model. In *2014 UKSim-AMSS 16th International Conference on Computer Modelling and Simulation*, 106–112.
- Bai, S.; Kolter, J. Z.; and Koltun, V. 2018. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. [arXiv:1803.01271](https://arxiv.org/abs/1803.01271).
- Fan, W.; Wang, P.; Wang, D.; Wang, D.; Zhou, Y.; and Fu, Y. 2023. Dish-TS: A General Paradigm for Alleviating Distribution Shift in Time Series Forecasting. [arXiv:2302.14829](https://arxiv.org/abs/2302.14829).
- Fei, J.; Yi, K.; Fan, W.; Zhang, Q.; and Niu, Z. 2025. Amplifier: Bringing Attention to Neglected Low-Energy Components in Time Series Forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 11645–11653.
- Jerome, A.; Ishii, T.; and Chen, H. 2018. Forecasting and Anomaly Detection on Application Metrics using LSTM. In Abe, N.; Liu, H.; Pu, C.; Hu, X.; Ahmed, N.; Qiao, M.; Song, Y.; Kossmann, D.; Liu, B.; Lee, K.; Tang, J.; He, J.; and Saltz, J., eds., *2018 IEEE INTERNATIONAL CONFERENCE ON BIG DATA (BIG DATA)*, IEEE International Conference on Big Data, 2221–2227. IEEE; IEEE Comp Soc; Expedia Grp; Baidu; Squirrel AI Learning; Ankura; Springer. ISBN 978-1-5386-5035-6. IEEE International Conference on Big Data (Big Data), Seattle, WA, DEC 10-13, 2018.
- Jia, Y.; Lin, Y.; Hao, X.; Lin, Y.; Guo, S.; and Wan, H. 2023. WITRAN: Water-wave Information Transmission and Recurrent Acceleration Network for Long-range Time Series Forecasting. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Lin, S.; Lin, W.; Hu, X.; Wu, W.; Mo, R.; and Zhong, H. 2024a. CycleNet: enhancing time series forecasting through modeling periodic patterns. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, 106315–106345.
- Lin, S.; Lin, W.; Wu, W.; Chen, H.; and Yang, J. 2024b. SparseTSF: modeling long-term time series forecasting with 1k parameters. In *Proceedings of the 41st International Conference on Machine Learning*, 30211–30226.
- Lin, S.; Lin, W.; Wu, W.; Zhao, F.; Mo, R.; and Zhang, H. 2023. SegRNN: Segment Recurrent Neural Network for Long-Term Time Series Forecasting. [arXiv:2308.11200](https://arxiv.org/abs/2308.11200).
- Liu, M.; Zeng, A.; Chen, M.; Xu, Z.; Lai, Q.; Ma, L.; and Xu, Q. 2022. SCINet: Time Series Modeling and Forecasting with Sample Convolution and Interaction. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS 35 (NEURIPS 2022)*, Advances in Neural Information Processing Systems. ISBN 978-1-7138-7108-8. 36th Conference on Neural Information Processing Systems (NeurIPS), ELECTR NETWORK, NOV 28-DEC 09, 2022.
- Liu, Y.; Hu, T.; Zhang, H.; Wu, H.; Wang, S.; Ma, L.; and Long, M. 2023. iTransformer: Inverted Transformers Are Effective for Time Series Forecasting. *CoRR*.
- Nawab, S.; Quatieri, T.; and Lim, J. 1983. Signal reconstruction from short-time Fourier transform magnitude. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 31(4): 986–998.
- Nie, Y.; Nguyen, N. H.; Sinthong, P.; and Kalagnanam, J. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *The Eleventh International Conference on Learning Representations*.
- Oreshkin, B. N.; Carpov, D.; Chapados, N.; and Bengio, Y. 2019. N-BEATS: Neural basis expansion analysis for interpretable time series forecasting. *arXiv preprint arXiv:1905.10437*.
- Wang, S.; Wu, H.; Shi, X.; Hu, T.; Luo, H.; Ma, L.; Zhang, J. Y.; and ZHOU, J. 2024a. TimeMixer: Decomposable Multiscale Mixing for Time Series Forecasting. In *The Twelfth International Conference on Learning Representations*.
- Wang, Y.; Zou, R.; Liu, F.; Zhang, L.; and Liu, Q. 2021. A review of wind speed and wind power forecasting with deep neural networks. *APPLIED ENERGY*, 304.
- Wu, H.; Xu, J.; Wang, J.; and Long, M. 2021. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems*, 34: 22419–22430.
- Wu, H.; Zhou, H.; Long, M.; and Wang, J. 2023. Interpretable weather forecasting for worldwide stations with a unified deep model. *NATURE MACHINE INTELLIGENCE*, 5(6): 602+.
- Xu, Z.; Zeng, A.; and Xu, Q. 2024. FITS: Modeling Time Series with 10k Parameters. In *The Twelfth International Conference on Learning Representations*.
- Yi, K.; Fei, J.; Zhang, Q.; He, H.; Hao, S.; Lian, D.; and Fan, W. 2024. Filtnet: Harnessing frequency filters for time series forecasting. *Advances in Neural Information Processing Systems*, 37: 55115–55140.
- Yi, K.; Zhang, Q.; Cao, L.; Wang, S.; Long, G.; Hu, L.; He, H.; Niu, Z.; Fan, W.; and Xiong, H. 2023a. A survey on deep learning based time series analysis with frequency transformation. *arXiv preprint arXiv:2302.02173*.
- Yi, K.; Zhang, Q.; Fan, W.; Wang, S.; Wang, P.; He, H.; An, N.; Lian, D.; Cao, L.; and Niu, Z. 2023b. Frequency-domain MLPs are more effective learners in time series forecasting. *Advances in Neural Information Processing Systems*, 36: 76656–76679.
- Yu, G.; Zou, J.; Hu, X.; Aviles-Rivero, A. I.; Qin, J.; and Wang, S. 2024. Revitalizing Multivariate Time Series Forecasting: Learnable Decomposition with Inter-Series Depen-

dencies and Intra-Series Variations Modeling. In *Forty-first International Conference on Machine Learning*.

Zeng, A.; Chen, M.; Zhang, L.; and Xu, Q. 2023. Are Transformers Effective for Time Series Forecasting? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11121–11128. Association for the Advancement of Artificial Intelligence (AAAI).

Zhang, T.; Zhang, Y.; Cao, W.; Bian, J.; Yi, X.; Zheng, S.; and Li, J. 2022. Less is more: Fast multivariate time series forecasting with light sampling-oriented mlp structures. *arXiv preprint arXiv:2207.01186*.

Zhao, Y.; Li, T.; Chen, J.; Ye, S.; Jiang, F.; Zhang, T.; and Gao, X. 2024. Disentangled Interpretable Representation for Efficient Long-term Time Series Forecasting. *arXiv preprint arXiv:2411.17257*.

Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; Xiong, H.; and Zhang, W. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 11106–11115.

Zhou, T.; Ma, Z.; Wen, Q.; Sun, L.; Yao, T.; Yin, W.; Jin, R.; et al. 2022a. Film: Frequency improved legendre memory model for long-term time series forecasting. *Advances in neural information processing systems*, 35: 12677–12690.

Zhou, T.; Ma, Z.; Wen, Q.; Wang, X.; Sun, L.; and Jin, R. 2022b. FEDformer: Frequency Enhanced Decomposed Transformer for Long-term Series Forecasting. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 27268–27286. PMLR.