

# Harmonic Dataset Distillation for Time Series Forecasting

Seungha Hong<sup>1</sup>, Sanghwan Jang<sup>1</sup>, Wonbin Kweon<sup>2</sup>, Suyeon Kim<sup>1</sup>, Gyuseok Lee<sup>2</sup>, Hwanjo Yu<sup>1\*</sup>

<sup>1</sup>Pohang University of Science and Technology (POSTECH), Republic of Korea

<sup>2</sup>University of Illinois Urban-Champaign (UIUC), United States

{shhong97, s.jang, kimsu, hwanjoju}@postech.ac.kr, {wonbin, gyuseok2}@illinois.edu

## Abstract

Time Series forecasting (TSF) in the modern era faces significant computational and storage cost challenges due to the massive scale of real-world data. Dataset Distillation (DD), a paradigm that synthesizes a small, compact dataset to achieve training performance comparable to that of the original dataset, has emerged as a promising solution. However, conventional DD methods are not tailored for time series and suffer from architectural overfitting and limited scalability. To address these issues, we propose **Harmonic Dataset Distillation for Time Series Forecasting (HDT)**. HDT decomposes the time series into its sinusoidal basis through the FFT and aligns the core periodic structure by *Harmonic Matching*. Since this process operates in the frequency domain, all updates during distillation are applied globally without disrupting temporal dependencies of time series. Extensive experiments demonstrate that HDT achieves strong cross-architecture generalization and scalability, validating its practicality for large-scale, real-world applications.

## 1 Introduction

Time Series Forecasting (TSF), a task that aims to predict future values based on historical data, is a critical task across numerous domains, including industrial manufacturing, healthcare, traffic, and meteorology (Miller et al. 2024). However, the practical application of TSF faces significant challenges related to data storage and computational cost. On the data side, sources such as industrial sensors and biomedical monitors collect data at high frequencies (often every minute or second), producing terabytes of sequential data daily and making it impractical to store the complete historical record (Cisco 2025; NOAA 2025; Lee et al. 2022). On the model side, the recent advent of large foundation models such as TimesFM (Das et al. 2024) and Moirai (Woo et al. 2024) has intensified these computational burdens, amplifying the need for efficient data handling.

A promising approach to address these challenges is *Dataset Distillation (DD)*, a paradigm that synthesizes a small, compact dataset whose training performance is comparable to that of the entire original dataset. Initially formulated as a bi-level optimization problem (Wang et al. 2018),

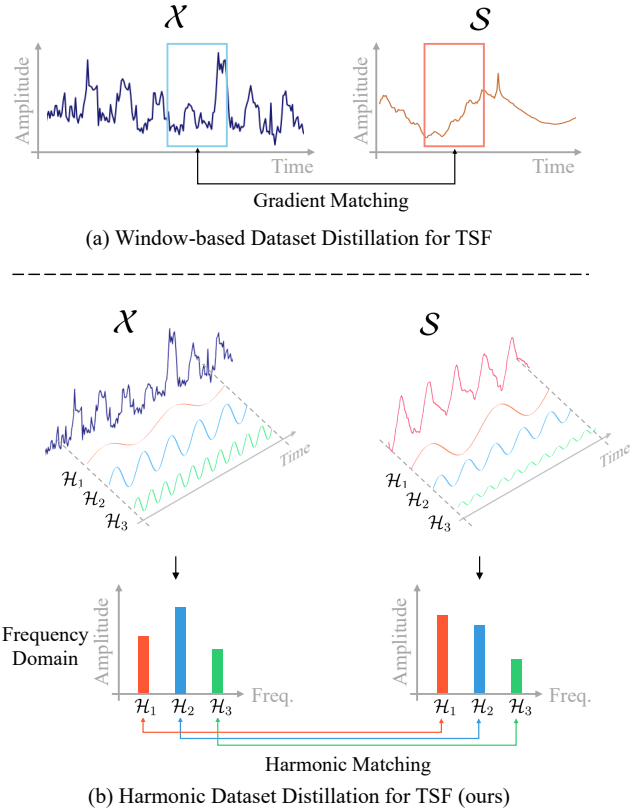


Figure 1: Illustrative examples of distilling an original dataset  $\mathcal{X}$  into a synthetic dataset  $\mathcal{S}$  with (a) Window-based and (b) Harmonic Dataset Distillation for TSF. In (a), windows randomly sampled from  $\mathcal{X}$  are distilled into arbitrary positions within  $\mathcal{S}$ . In (b), the sequence is decomposed into a sinusoidal basis, and selected harmonics ( $\mathcal{H}_i$ ) between  $\mathcal{X}$  and  $\mathcal{S}$  are aligned during distillation.

DD has demonstrated its effectiveness in image classification (Zhao and Bilen 2021; Nguyen, Chen, and Lee 2021; Cazenavette et al. 2022; Zhao and Bilen 2023). Recently, there have been attempts to extend it to other data modalities such as graphs (Jin et al. 2022; Liu, Zeng, and Zheng 2024), time series (Liu et al. 2024b; Ding et al. 2024), and natural language (Maekawa et al. 2023, 2024).

\*Corresponding Author

However, directly applying conventional DD methods to TSF fails to account for the global structure of time series data. In TSF, models use small, fixed-size windows of data as input and output, even when the overall time series is long. For example, a model might use 96 past steps to predict the next 96 steps while the total time series spans 20,000 to 70,000 data points. If we adopt existing DD methods directly to time series, each time window is treated as an independent data instance. Therefore, optimization is performed within these local windows, trying to match local windows from the synthetic data to those from the original data (Figure 1-(a)). This “local-to-local” approach, which we refer to as *Window-based Dataset Distillation for TSF*, completely disregards the characteristics of time series, such as long-range dependencies and periodicity. This failure to capture the global context leads to limited scalability with respect to the synthetic data size and degraded cross-architecture generalization.

To tackle these limitations, we propose *Harmonic Dataset Distillation for Time Series Forecasting (HDT)*. Instead of performing distillation on localized time domain window pairs, HDT shifts the optimization space to the frequency domain using the Fast Fourier Transform (FFT). Specifically, we apply the FFT to represent both the original and synthetic time series as a sum of sinusoidal basis functions. Among these, we define the dominant components as *harmonics*, which contain the core periodic information of the sequence. Our *Harmonic Matching* aligns the distributions of these harmonics between the original and synthetic data, thereby preserving the global structure of time series (Figure 1-(b)). Furthermore, since each basis has a global influence over the entire sequence, any updates in this optimization space result in a modification of the synthetic sequence as a whole. This ensures updates are applied without disrupting the temporal dependencies within the synthetic series, even with a conventional distillation loss.

To validate our proposed method, we provide a comprehensive evaluation from both theoretical and empirical standpoints. We first conduct a theoretical analysis to formally justify how *Harmonic Matching* preserves essential temporal structures. Complementing this, our extensive experiments on various TSF datasets confirm that HDT achieves state-of-the-art performance across diverse backbone architectures and synthetic dataset sizes.

In summary, our work makes the following contributions:

- We introduce HDT, a novel and effective dataset distillation method for time series forecasting.
- A theoretical analysis is provided, proving that the synthetic dataset distilled by HDT preserves the essential global structure of the original data.
- Extensive experiments on modern backbones demonstrate the method’s state-of-the-art performance and strong cross-architecture generalization.

## 2 Preliminaries

We first provide the problem formulation and key notations of Time Series Forecasting (TSF) and Dataset Distillation (DD). Then, we examine the scenario of directly applying

conventional DD methods to TSF and discuss their limitations.

### 2.1 Time Series Forecasting

A time series, denoted as  $\mathcal{X} \in \mathbb{R}^{N \times C}$ , is a sequence of data points measured at successive time intervals. Here,  $N$  is the total length of the time series, and  $C$  is the number of variables (or channels) measured at each time step. In TSF, given a historical lookback window  $\mathcal{X}_{s+1:s+l}$  of length  $l$  starting from an arbitrary time step  $s$ , the objective is to predict a future forecast horizon  $\mathcal{X}_{s+l+1:s+l+t}$  of length  $t$ .

For model training and evaluation, the entire series  $\mathcal{X}$  is divided chronologically into training set  $\mathcal{X}_{\text{train}}$  and test set  $\mathcal{X}_{\text{test}}$ . From each of these sets, a sliding window is applied to generate input-output pairs  $(x, y)$ , where  $x = \mathcal{X}_{s+1:s+l}$  and  $y = \mathcal{X}_{s+l+1:s+l+t}$  for all possible start times  $s$ . With the prepared training samples, the model is trained by minimizing the loss function  $\mathcal{L}$ :

$$\mathcal{L}(\theta, \mathcal{X}) := \frac{1}{|\mathcal{X}|} \sum_{(x,y) \in \mathcal{X}} d(f_{\theta}(x), y). \quad (1)$$

Here, the loss  $\mathcal{L}$  is the average of a distance metric  $d$  (e.g., L2 distance) between the model’s predictions  $f_{\theta}(x)$  and the true values  $y$  over all pairs in the dataset.

The training process finds the optimal parameters  $\theta^*$  by solving the following optimization problem on the training data:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \mathcal{L}(\theta, \mathcal{X}_{\text{train}}). \quad (2)$$

For convenience, we encapsulate this entire training process as a function  $\mathcal{T}$ . This function takes the initial parameter  $\theta$  and the training dataset  $\mathcal{X}_{\text{train}}$  as inputs, and returns the trained parameters:

$$\mathcal{T}(\theta, \mathcal{X}_{\text{train}}) \approx \underset{\theta}{\operatorname{argmin}} \mathcal{L}(\theta, \mathcal{X}_{\text{train}}). \quad (3)$$

### 2.2 Dataset Distillation

The goal of *Dataset Distillation* (Wang et al. 2018) is to create a small synthetic dataset  $\mathcal{S}$  that can train a model as effective as the original training dataset  $\mathcal{X}_{\text{train}}$ :

$$\mathcal{L}(\mathcal{T}(\theta, \mathcal{S}), \mathcal{X}_{\text{test}}) \approx \mathcal{L}(\mathcal{T}(\theta, \mathcal{X}_{\text{train}}), \mathcal{X}_{\text{test}}). \quad (4)$$

This is achieved by optimizing  $\mathcal{S}$  to minimize the evaluation loss of  $\mathcal{T}(\theta, \mathcal{S})$  over  $\mathcal{X}_{\text{test}}$ , which can be formulated as the following bi-level optimization problem:

$$\mathcal{S}^* = \underset{\mathcal{S}}{\operatorname{argmin}} \mathcal{L}(\mathcal{T}(\theta, \mathcal{S}), \mathcal{X}_{\text{test}}). \quad (5)$$

Here,  $\mathcal{S}$  is treated as learnable parameters and optimized using algorithms such as stochastic gradient descent (SGD). However, directly optimizing Equation 5 is computationally expensive due to the inclusion of the entire inner-loop training process  $\mathcal{T}(\theta, \mathcal{S})$ . To mitigate this, surrogate objectives that approximate the original optimization problem have been proposed (Zhao and Bilen 2021; Cazenavette et al. 2022; Zhao and Bilen 2023).

In dataset distillation, the effectiveness of a synthetic dataset  $\mathcal{S}$  hinges on its ability to effectively train arbitrary

models, not just the one used for its creation. This *cross-architecture generalization* ensures that  $\mathcal{S}$  captures the essential knowledge of the original dataset  $\mathcal{X}_{\text{train}}$  rather than merely memorizing a specific model’s training process.

### 2.3 Dataset Distillation for TSF

The objective is to obtain a synthetic sequence  $\mathcal{S} \in \mathbb{R}^{M \times C}$  that serves as a compressed substitute for the original time series  $\mathcal{X} \in \mathbb{R}^{N \times C}$ , where  $M \ll N$ .<sup>1</sup> A straightforward way to extend image dataset distillation methods to TSF datasets is to randomly sample windows from  $\mathcal{X}$  and  $\mathcal{S}$  and treat them as individual instances (Zhao and Bilén 2021; Cazenavette et al. 2022; Cui et al. 2023; Ding et al. 2024). Specifically, updating  $\mathcal{S}$  involves randomly sampling mini-batches from  $\mathcal{S}$  and  $\mathcal{X}$ , which are represented as follows:

$$\mathcal{B}_S \in \mathbb{R}^{|\mathcal{B}_S| \times (l+t) \times C}, \mathcal{B}_X \in \mathbb{R}^{|\mathcal{B}_X| \times (l+t) \times C}. \quad (6)$$

Here, each window has a length of  $l+t$  (input window + output window). The distillation loss (e.g., gradient matching (Zhao and Bilén 2021)) is computed for the sampled mini-batches, and the data points within the sampled windows  $\mathcal{B}_S$  are directly updated. We refer to this approach as *Window-based TSF Dataset Distillation*.

The “local-to-local” matching of window-based approaches presents critical limitations from two perspectives:

**L1. Limited Scalability** Ideally, increasing the size of the synthetic dataset ( $M$ ) should improve performance by capturing more diverse patterns within training dataset. However, with this approach, a larger  $M$  merely elongates the existing local patterns rather than capturing the broader, global structure of the time series. Consequently, this leads to diminishing returns, since the added data points provide little valuable information.

**L2. Architectural Overfitting** The window-based approaches update only the data points within a given window to minimize the specified distillation loss. This local optimization completely ignores the global dependencies that structure the entire time series. With the global context disregarded, the distillation leads the synthetic data to overfit to the specific inductive biases of the fixed backbone model, rather than learning the essential patterns inherent to the original data. Accordingly, this results in a degraded cross-architecture generalization.

## 3 Methodology

To address the limitations of window-based approaches, we propose *Harmonic Dataset Distillation for Time Series Forecasting (HDT)*. The core idea is to decompose the time series into its sinusoidal basis through FFT and perform distillation in the frequency domain. Instead of updating data points within local windows, our method directly updates the selected dominant basis functions, which we refer to as *harmonics*. The optimization is guided by two components: (1) *Harmonic Matching* to align the global harmonic

<sup>1</sup>For simplicity,  $\mathcal{X}$  will henceforth be considered to include only the training windows (i.e.,  $\mathcal{X} := \mathcal{X}_{\text{train}}$ ).

distributions, and (2) *Gradient Matching* as a distillation loss. This approach allows HDT to effectively capture global structures and overcome the limitations of window-based methods. An overview of our method is depicted in Figure 2.

### 3.1 Harmonic Matching

The primary objective of Harmonic Matching is to align the harmonic distributions of the synthetic data  $\mathcal{S}$  and the original data  $\mathcal{X}$ . To achieve this, we begin by decomposing the time series into a sum of sinusoidal basis functions with FFT.

A precise alignment of basis functions between  $\mathcal{X}$  and  $\mathcal{S}$  is required, but their different lengths prevent a direct one-to-one matching of components that represent the *same period*. Furthermore, a longer sequence produces a high-resolution spectrum in which numerous high-frequency components can act as noise, diluting the prominence of harmonics. To resolve these issues, we sample a subsequence  $\mathcal{X}_{\text{sub}}$  from  $\mathcal{X}$  with the same length  $M$  as  $\mathcal{S}$  and apply the FFT to both to obtain their frequency domain representations:

$$\mathcal{F}_X = \text{FFT}(\mathcal{X}_{\text{sub}}), \quad \mathcal{F}_S = \text{FFT}(\mathcal{S}). \quad (7)$$

Here, the FFT efficiently computes the Discrete Fourier Transform (DFT), which decomposes the sequences  $\mathcal{X}_{\text{sub}} = \{x_n\}_{n=0}^{M-1}$  and  $\mathcal{S} = \{s_n\}_{n=0}^{M-1}$  of length  $M$  into its frequency components  $\mathcal{F}_X$  and  $\mathcal{F}_S$  as follows:

$$\begin{aligned} \mathcal{F}_X[k] &= \sum_{n=0}^{M-1} x_n e^{-i2\pi kn/M}, \quad \text{and} \\ \mathcal{F}_S[k] &= \sum_{n=0}^{M-1} s_n e^{-i2\pi kn/M}, \quad k = 0, \dots, M-1. \end{aligned} \quad (8)$$

Using all these frequency components is often suboptimal because it incorporates unnecessary noise, which can obscure the true underlying patterns. Therefore, we selectively distill the most significant basis functions, which correspond to the top- $k$  frequencies with the largest amplitudes in  $\mathcal{F}_X$ . We refer to these top- $k$  components as *harmonics*  $\mathcal{H}$ , and obtain the following harmonics-based sequence representations:

$$\begin{aligned} \mathcal{H} &= \arg \text{top-}k_{i \in [0, \lfloor M/2 \rfloor]} (|\mathcal{F}_X[i]|) \\ \tilde{\mathcal{F}}_X[i] &:= \begin{cases} \mathcal{F}_X[i], & i \in \mathcal{H} \\ 0, & \text{otherwise} \end{cases} \quad \forall i \in [0, \lfloor M/2 \rfloor], \\ \tilde{\mathcal{F}}_S[i] &:= \begin{cases} \mathcal{F}_S[i], & i \in \mathcal{H} \\ 0, & \text{otherwise} \end{cases} \quad \forall i \in [0, \lfloor M/2 \rfloor]. \end{aligned} \quad (9)$$

During the distillation process, we directly update these selected harmonics instead of data points within local windows. Since each harmonic is a sinusoidal basis function with a global influence, every update modifies the synthetic series as a whole.

This global update mechanism is the key to resolving the scalability limitation (**L1**). The length of the synthetic series determines the range of periods that the sinusoidal basis can represent. Increasing  $M$  allows our method to capture the harmonics corresponding to *longer periods*, which contain a richer set of long-range global structures. This provides a

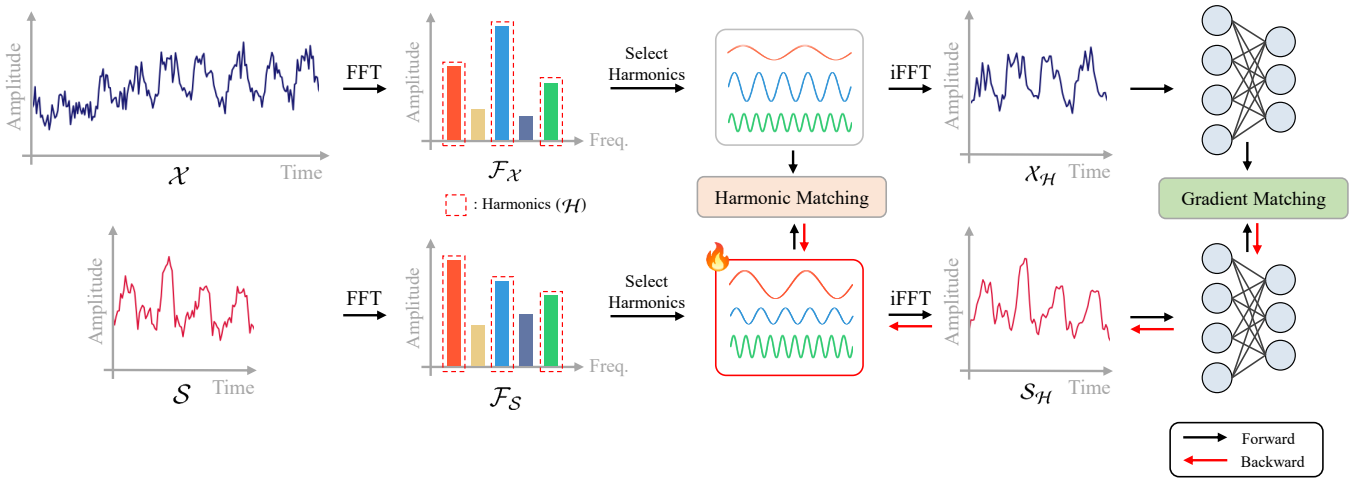


Figure 2: An overview of our method. Selected harmonics (red box) of synthetic data are updated through *Harmonic Matching* and *Gradient Matching*.

key advantage over window-based methods, which are restricted to learning additional local patterns without capturing a broader context. This ensures that the performance of our method scales meaningfully with the size of the synthetic dataset.

To enforce similarity between the selected harmonics, we introduce a harmonic loss  $\mathcal{L}_{\text{harm}}$ , which minimizes the  $L_p$ -norm distance between the amplitudes of these harmonic coefficients:

$$\mathcal{L}_{\text{harm}} = \|\tilde{\mathcal{F}}_{\mathcal{X}} - \tilde{\mathcal{F}}_{\mathcal{S}}\|_p. \quad (10)$$

Minimizing this loss acts as a regularizer, forcing the periodic structure of  $\mathcal{S}$  to align with that of  $\mathcal{X}$ . This resolves the architectural overfitting problem (L2), since the synthetic data learns to capture the harmonic distribution, which is an intrinsic model-agnostic property of the data, rather than overfitting to the specific biases of a single backbone. The formal justification for how minimizing  $\mathcal{L}_{\text{harm}}$  preserves the global structure of the original series is provided in Theorem 1.

**Theoretical Analysis for Harmonic Matching** We further provide a theoretical justification for how *Harmonic Matching* successfully preserves the global structure of time series data. Our analysis is grounded in the relationship between a series’s Power Spectral Density (PSD), which describes its power distribution across frequencies, and its Autocorrelation Function (ACF), which measures its temporal dependencies. The core principle is that by aligning the most significant components of the PSD, the *harmonics*, we ensure that the autocorrelation structures of the original and synthetic data are also aligned.

**Theorem 1.** *Let  $\mathcal{F}_{\mathcal{X}}$  and  $\mathcal{F}_{\mathcal{S}}$  are the DFTs of an  $M$ -point subset (segment) of  $\mathcal{X}$  and  $\mathcal{S}$ , and let  $r_{\mathcal{X}}(k)$  and  $r_{\mathcal{S}}(k)$  denote their respective ACFs at lag  $k$ . Suppose we choose  $\mathcal{S}$  so as to minimize  $\|\mathcal{F}_{\mathcal{X}} - \mathcal{F}_{\mathcal{S}}\|_p$ . Then, for a given maximum lag  $K$ , there exists a constant  $C > 0$  such that the difference*

*in their ACFs is bounded:*

$$\max_{|k| \leq K} |r_{\mathcal{S}}(k) - r_{\mathcal{X}}(k)| \leq C \varepsilon,$$

*where  $\varepsilon$  is a measure of how closely  $\mathcal{F}_{\mathcal{S}}$  approximates  $\mathcal{F}_{\mathcal{X}}$  in the frequency domain.*

Theorem 1 provides the theoretical foundation for *Harmonic Matching*. The inequality shows that the maximum error between the autocorrelation of the original and synthetic datasets is directly controlled by  $\varepsilon$ , the frequency domain approximation error. Since *Harmonic Matching* is explicitly designed to minimize this error, the theorem guarantees that it effectively preserves the underlying temporal dependencies of the original time series. For a detailed derivation and the necessary technical assumptions, refer to the full proof in Appendix.

### 3.2 Gradient Matching

For use in the subsequent gradient matching step, we reconstruct time-domain signals containing only these essential harmonic patterns via the inverse FFT (iFFT):

$$\mathcal{X}_{\mathcal{H}} = \text{iFFT}(\tilde{\mathcal{F}}_{\mathcal{X}}), \quad \mathcal{S}_{\mathcal{H}} = \text{iFFT}(\tilde{\mathcal{F}}_{\mathcal{S}}). \quad (11)$$

We employ the surrogate objective similar to previous works (Cazenavette et al. 2022; Cui et al. 2023), which matches the multi-step gradients with respect to the model parameters  $\theta$  when the model is trained on  $\mathcal{S}$  versus on  $\mathcal{X}$ . We train the model parameter  $\theta$  on  $\mathcal{X}_{\mathcal{H}}$  for  $i$  steps and on  $\mathcal{S}_{\mathcal{H}}$  for  $j$  steps, then use their normalized Euclidean distance as the distillation loss:

$$\mathcal{L}_{\text{grad}} = \frac{\|\mathcal{T}_j(\theta, \mathcal{S}_{\mathcal{H}}) - \mathcal{T}_i(\theta, \mathcal{X}_{\mathcal{H}})\|_2^2}{\|\theta - \mathcal{T}_i(\theta, \mathcal{X}_{\mathcal{H}})\|_2^2}. \quad (12)$$

Here,  $\mathcal{T}_i$  denotes the  $i$ -step training process.

By combining  $\mathcal{L}_{\text{grad}}$  with  $\mathcal{L}_{\text{harm}}$ , we arrive at the final objective function:

$$\underset{\mathcal{F}_{\mathcal{S}}}{\text{argmin}} \mathcal{L}_{\text{grad}} + \lambda \mathcal{L}_{\text{harm}}, \quad (13)$$

where  $\lambda$  serves as a hyperparameter to balance the contributions of the two losses. Once  $\mathcal{F}_S$  has converged, the final distilled dataset  $\mathcal{S}$  is recovered with the iFFT. The entire distillation process is described in Appendix.

## 4 Experiments

### 4.1 Setup

**Datasets** We evaluated our method on widely used TSF benchmarks: the ETT, Electricity, and Traffic datasets (Lai et al. 2018; Zhou et al. 2021). The details of each dataset are provided in Appendix

**Backbones** Previous state-of-the-art methods (Ding et al. 2024) have employed simple models (e.g., MLP, LSTM, CNN) as backbones. To evaluate the generality of the dataset distillation algorithm, we instead use recent state-of-the-art TSF models: DLinear (Zeng et al. 2023), iTransformer (Liu et al. 2024a), and xPatch (Stitsyuk and Choi 2025), which are the representatives of Linear, Transformer, and CNN architectures, respectively. The details of the models are in Appendix.

**Baselines** We compare our method against several baselines, starting with a naive ‘Random’ that trains on a randomly selected subsequence of the original data. We then include state-of-the-art dataset distillation methods: DC (Zhao and Bilen 2021), MTT (Cazenavette et al. 2022), TESLA (Cui et al. 2023), and CondTSF (Ding et al. 2024). The performance of ‘Full Data’ training is also reported as a lower bound of MSE. Detailed descriptions of each baseline are provided in the Appendix.

### 4.2 Overall Results

We conducted comprehensive experiments to validate the empirical performance of HDT, following the evaluation protocol described in Appendix. The evaluation covers both *fixed-architecture* settings, where the backbone and evaluation models are the same, and *cross-architecture* settings, where they differ. Table 1 presents the overall results for  $M = 384$ , using DLinear (L), iTransformer (T), and xPatch (C) as the backbone and the evaluation models.

A key finding is that while prior methods perform well in fixed-architecture settings, they suffer from unreliable performance in cross-architecture scenarios. Their performance often degrades significantly when the backbone and evaluation models differ, in some cases performing even worse than the ‘Random’ baseline (e.g., xPatch backbone evaluated on DLinear). This highlights their severe architectural overfitting (L2), as the distilled data is over-specialized to a single model.

In contrast, HDT not only achieves state-of-the-art performance across nearly all settings but also demonstrates remarkable robustness. Its performance remains strong and stable in cross-architecture settings, indicating that it distills a more universal, model-agnostic representation of the original data. The two rightmost plots of Figure 3 (derived from Table 1) clearly illustrate that only HDT maintains a minimal increase in MSE in cross-architecture settings.

Finally, to evaluate the scalability of our method (L1), we measured the performance as the synthetic data size  $M$  increases. The two leftmost plots in Figure 3 show that HDT’s performance consistently improves with a larger  $M$ , while others saturate beyond a certain size. This confirms HDT’s ability to effectively capture the long-range context of the time series.

To supplement our main findings, we provide additional results including a detailed hyperparameter analysis, experiments with a smaller synthetic dataset, and qualitative visualizations in Appendix.

### 4.3 Ablation Study

To validate the contribution of each key component in HDT, we conducted an ablation study using the DLinear backbone. We compared three distinct configurations, measuring their average performance across all evaluation models (L, T, and C): (1) Base, which uses only conventional gradient matching on time windows; (2) Base + Decomp., which performs gradient matching only with sinusoidal decomposition; and (3) HDT, the full proposed method.

The results in Table 2 show the effectiveness of our components. The improvement from ‘Base’ to ‘Base + Decomp.’ demonstrates the advantage of operating in the frequency domain, as each update has a global influence over the entire sequence. The final substantial improvement of ‘HDT’ is delivered by Harmonic Matching, which distills the essential periodic patterns by aligning the distribution of selected harmonics between the synthetic and original data.

### 4.4 Efficiency Analysis

**Distillation Runtime** We report the distillation runtime (in seconds) over 100 outer-loop steps in Table 3. Although HDT incorporates the FFT, its theoretical complexity of  $O(M \log M)$  is minor compared to the gradient computations of the backbone model. As Table 3 shows, the runtime gap between HDT and other methods narrows as model complexity increases (from DLinear to iTransformer), confirming that the overall overhead of HDT is marginal. These results, combined with the performance gains reported in Table 1, demonstrate that HDT achieves superior performance with only a slight increase in computational cost.

**Training Efficiency** We measured the actual training time of iTransformer on the full dataset versus our distilled dataset. As shown in Table 4, training on the distilled data achieves a dramatic speed-up, reducing training time from hours to mere seconds. This highlights the substantial computational benefits of the distilled dataset.

### 4.5 Large Scale Scenario

**Performance on Large-Scale Datasets** Evaluations in TSF often rely on lightweight datasets that may not represent the scale of real-world data. To address this gap, we evaluated HDT on the massive CA dataset from the LargeST (Liu et al. 2023) benchmark, which has a length of 201,363 and 8,600 features. For the experiment, we distilled this dataset to a size of  $M = 384$  using the DLinear backbone. As shown

Backbone: DLinear																		
Dataset	ETTh1			ETTh2			ETTm1			ETTm2			Electricity			Traffic		
	L	T	C	L	T	C	L	T	C	L	T	C	L	T	C	L	T	C
Random	0.945	0.757	0.664	1.860	0.406	0.359	0.919	0.732	0.666	1.504	0.256	0.234	0.400	0.327	0.351	1.112	0.938	0.911
DC	1.193	0.853	0.721	2.200	0.410	0.363	0.801	0.688	0.659	1.670	0.266	0.234	0.413	0.333	0.351	1.122	0.945	0.924
MTT	0.521	0.640	0.587	0.661	0.387	0.346	0.493	0.879	1.342	0.702	0.257	0.248	0.342	0.412	0.489	0.747	1.204	0.852
TESLA	0.576	0.535	0.530	0.724	0.381	0.348	0.443	0.475	0.438	0.428	0.231	0.220	0.388	0.413	0.394	1.087	0.956	0.953
CondTSF	<u>0.510</u>	<u>0.494</u>	<u>0.492</u>	<u>0.392</u>	<u>0.336</u>	<u>0.325</u>	<u>0.410</u>	<u>0.422</u>	<u>0.416</u>	<u>0.223</u>	<u>0.209</u>	<u>0.204</u>	<u>0.231</u>	<u>0.241</u>	<u>0.238</u>	<u>0.716</u>	<u>0.854</u>	<u>0.807</u>
HDT (ours)	<b>0.430</b>	<b>0.421</b>	<b>0.409</b>	<b>0.359</b>	<b>0.331</b>	<b>0.311</b>	<b>0.364</b>	<b>0.405</b>	<b>0.389</b>	<b>0.211</b>	<b>0.205</b>	<b>0.201</b>	<b>0.208</b>	<b>0.239</b>	<b>0.232</b>	<b>0.679</b>	<b>0.847</b>	<b>0.754</b>
Full Data	0.386	0.389	0.384	0.326	0.314	0.296	0.343	0.345	0.334	0.186	0.185	0.177	0.195	0.152	0.175	0.648	0.408	0.470

Backbone: iTransformer																		
Dataset	ETTh1			ETTh2			ETTm1			ETTm2			Electricity			Traffic		
	L	T	C	L	T	C	L	T	C	L	T	C	L	T	C	L	T	C
Random	0.945	0.757	0.664	1.860	0.406	0.359	0.919	0.732	0.666	1.504	0.256	0.234	0.400	0.327	0.351	1.112	0.938	0.911
DC	1.228	0.869	0.722	3.200	0.403	0.361	0.871	0.675	0.663	2.472	0.255	0.232	0.462	0.337	0.363	1.000	0.882	0.874
MTT	1.340	<u>0.452</u>	<u>0.441</u>	7.247	0.375	0.338	1.368	0.584	0.567	4.475	0.251	0.228	0.535	0.316	0.324	<u>0.881</u>	0.839	0.943
TESLA	1.079	0.480	0.503	4.914	<u>0.327</u>	<u>0.323</u>	1.299	0.452	0.428	5.028	0.204	0.205	0.422	0.305	0.351	1.027	0.875	0.861
CondTSF	<u>0.519</u>	0.496	0.492	<u>0.442</u>	<u>0.334</u>	<u>0.326</u>	<u>0.411</u>	<u>0.411</u>	<u>0.407</u>	<u>0.236</u>	<u>0.200</u>	<b>0.199</b>	<u>0.231</u>	<u>0.235</u>	<u>0.234</u>	0.965	<u>0.824</u>	<u>0.822</u>
HDT (ours)	<b>0.409</b>	<b>0.399</b>	<b>0.397</b>	<b>0.348</b>	<b>0.323</b>	<b>0.320</b>	<b>0.377</b>	<b>0.406</b>	<b>0.399</b>	<b>0.211</b>	<b>0.199</b>	<u>0.200</u>	<b>0.229</b>	<b>0.231</b>	<b>0.230</b>	<b>0.755</b>	<b>0.793</b>	<b>0.734</b>
Full Data	0.386	0.389	0.384	0.326	0.314	0.296	0.343	0.345	0.334	0.186	0.185	0.177	0.195	0.152	0.175	0.648	0.408	0.470

Backbone: xPatch																		
Dataset	ETTh1			ETTh2			ETTm1			ETTm2			Electricity			Traffic		
	L	T	C	L	T	C	L	T	C	L	T	C	L	T	C	L	T	C
Random	0.945	0.757	0.664	1.860	0.406	0.359	0.919	0.732	0.666	1.504	0.256	0.234	0.400	0.327	0.351	1.112	0.938	0.911
DC	1.290	0.829	0.712	3.928	0.408	0.363	0.751	0.657	0.654	2.477	0.250	0.230	0.372	0.322	0.343	1.048	0.882	0.868
MTT	1.740	<u>0.496</u>	<u>0.486</u>	6.968	<u>0.359</u>	<u>0.329</u>	0.904	0.589	0.538	5.021	0.240	0.229	0.420	0.334	0.348	1.175	1.025	0.828
TESLA	<u>1.112</u>	0.560	0.544	3.850	0.397	0.362	1.161	0.666	<u>0.421</u>	3.439	<u>0.207</u>	<u>0.203</u>	0.508	0.334	0.372	1.003	0.844	0.860
CondTSF	1.172	0.629	0.587	<u>2.897</u>	0.406	0.362	<u>0.574</u>	<u>0.550</u>	0.531	<u>0.349</u>	0.248	0.230	<u>0.237</u>	<u>0.261</u>	<u>0.246</u>	<u>0.904</u>	<u>0.827</u>	<u>0.808</u>
HDT (ours)	<b>0.412</b>	<b>0.416</b>	<b>0.394</b>	<b>0.350</b>	<b>0.349</b>	<b>0.319</b>	<b>0.370</b>	<b>0.523</b>	<b>0.396</b>	<b>0.251</b>	<b>0.204</b>	<b>0.198</b>	<b>0.237</b>	<b>0.253</b>	<b>0.240</b>	<b>0.716</b>	<b>0.787</b>	<b>0.739</b>
Full Data	0.386	0.389	0.384	0.326	0.314	0.296	0.343	0.345	0.334	0.186	0.185	0.177	0.195	0.152	0.175	0.648	0.408	0.470

Table 1: Overall dataset distillation performance in terms of MSE. The synthetic data size  $M$  is set to 384. **Bold** indicates the best result, while underlined denotes the second-best result.

Method	ETTh1	ETTh2	ETTm1	ETTm2	Electricity	Traffic
Base	0.583	0.465	0.905	0.402	0.414	0.934
Base + Decomp.	0.545	0.420	0.814	0.325	0.376	0.902
HDT	<b>0.420</b>	<b>0.334</b>	<b>0.386</b>	<b>0.206</b>	<b>0.226</b>	<b>0.760</b>

Table 2: Ablation study on the components of HDT. Base denotes conventional gradient matching, and Decomp. refers to the use of the sinusoidal decomposition.

in Table 5, HDT scales effectively to massive data, significantly outperforming previous distillation methods and approaching the performance of full-data training.

**Fine-tuning Large Foundation Models** We further investigated HDT’s utility in the context of large foundation models. We evaluated the pretrained Moirai-Large (311M parameters) (Woo et al. 2024) on the ETTh1 dataset under three scenarios: zero-shot, full fine-tuning, and few-shot fine-tuning with a dataset distilled with HDT. Table 6 shows that the distilled dataset enables substantial performance gains over zero-shot with a minimal fraction of the training cost of full fine-tuning, demonstrating a highly favorable cost-performance trade-off. Specifically, our approach is 80x

Backbone	Dataset	MTT	CondTSF	HDT
DLinear	ETTh1	30.06	32.13	36.37
	Electricity	35.61	38.04	42.76
iTransformer	ETTh1	197.69	199.95	209.46
	Electricity	253.90	262.10	264.90

Table 3: Comparison of running times over 100 outer-loops in seconds. The cost increase of HDT is marginal.

faster than full fine-tuning while incurring only a 2.5% degradation in MSE performance. Even including the distillation process (end-to-end), HDT still achieves an 8.51x speed-up.

## 5 Related Works

### 5.1 Time Series Forecasting

Recent research in Time Series Forecasting (TSF) has evolved along several distinct architectural paradigms. Transformer-based models are prominent for their ability to capture long-range dependencies (Liu et al. 2024a; Nie et al. 2023), simple Linear-based models have demonstrated

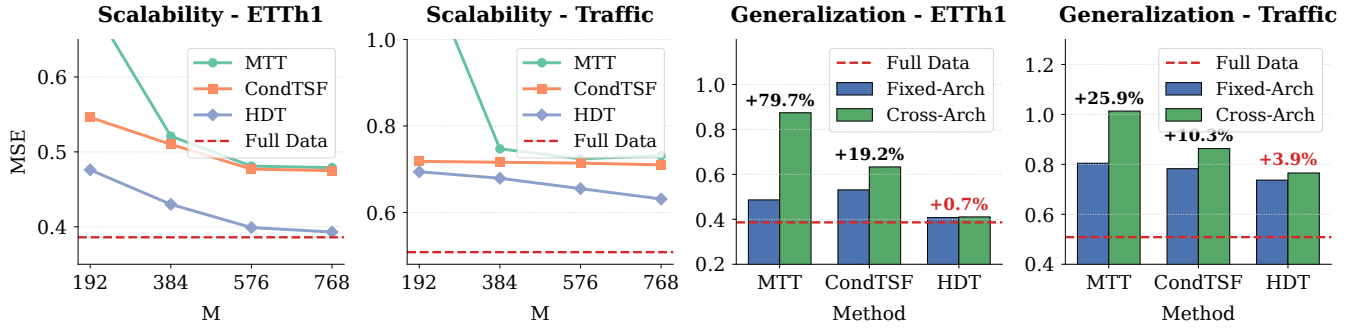


Figure 3: Scalability and cross-architecture generalization performance on the ETTh1 and Traffic datasets. The left two plots show scalability results for varying synthetic data sizes ( $M$ ). The right two plots compare the performance between fixed- and cross-architecture settings, highlighting that HDT maintains a significantly smaller increase in MSE compared to others.

Dataset	Full Training	Distilled	Speed-up
Electricity	1650.33	1.98	<b>834x</b>
Traffic	4266.45	2.32	<b>1839x</b>

Table 4: Training time comparison in seconds with the iTransformer backbone, demonstrating the significant efficiency gains achieved by training on the distilled dataset.

	Random	CondTSF	HDT	Full Data
MSE	358.51	197.95	<b>46.63</b>	44.25

Table 5: Distillation performance (MSE) on CA dataset with the DLinear backbone. HDT’s performance closely approaches that of full-data training.

surprisingly strong performance (Zeng et al. 2023; Ekambaram et al. 2023; Xu, Zeng, and Xu 2024; Wang et al. 2024), and CNN-based approaches have proven effective for capturing local temporal patterns (Wu et al. 2022; Luo and Xue 2024; Stitsyuk and Choi 2025). A more recent trend is the development of large pre-trained foundation models, which have shown zero-shot forecasting capabilities (Das et al. 2024; Woo et al. 2024). To demonstrate the general applicability of HDT across these distinct paradigms, we select representative models from each: iTransformer (Liu et al. 2024a), DLinear (Zeng et al. 2023), xPatch (Stitsyuk and Choi 2025), and Moirai (Woo et al. 2024).

## 5.2 Dataset Distillation

Dataset Distillation (DD) aims to create a compact synthetic dataset that yields training performance comparable to that of the original dataset. This concept was first introduced in (Wang et al. 2018), but it suffered from suboptimal performance due to the need for nested-loop optimization. To address this issue, three main approaches have been developed: (1) optimizing surrogate objective, (2) kernel-based methods, and (3) parameterization. Instead of directly minimizing the final performance, several works optimize surrogate objectives such as matching gradients (Zhao and Bilén 2021), training trajectories (Cazenavette et al. 2022; Cui et al. 2023;

Setting	MSE	Training Time
Zero-shot	1.972	-
Full Fine-tuning	1.383	3775.68s
Few-shot with HDT (end-to-end)	1.417	443.51s
Few-shot with HDT	1.417	47.23s
Comparison	<b>+2.5%</b>	<b>80x (8.51x)</b>

Table 6: Fine-tuning the Moirai-Large (311M) on ETTh1. The final row shows the MSE increase and training speed-up of our method relative to full fine-tuning.

Guo et al. 2024), or feature distributions (Zhao and Bilén 2023; Zhang et al. 2024) between the synthetic and original datasets. Kernel-based methods simplify the optimization process with kernel ridge regression to derive a closed-form solution for the inner loop (Nguyen, Chen, and Lee 2021; Nguyen et al. 2021; Zhou, Nezhadarya, and Ba 2022). Parameterization methods optimize alternative parameters instead of directly optimizing image pixels (Cazenavette et al. 2023; Liu and Wang 2023b,a; Wei et al. 2023; Shin, Shin, and Moon 2023). Our work combines a surrogate objective and a parameterization method, both tailored for the unique challenges of the TSF dataset distillation.

## 6 Conclusion

In this work, we introduced HDT, a novel dataset distillation paradigm for TSF. By shifting the distillation process to the frequency domain to align the *harmonics*, HDT effectively preserves the global structure of the original series. Our experiments demonstrate that this approach resolves the critical limitations of prior methods: limited scalability and architectural overfitting. Furthermore, we validated HDT’s practical utility in large-scale, real-world scenarios, including its performance on massive datasets and its effectiveness in fine-tuning foundation models. As the need for efficient data handling becomes increasingly critical in domains such as online learning and resource-constrained settings, the application of DD for TSF presents a promising avenue for future research.

## Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00217286, No. RS-2024-00335873), Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2019-II191906, Artificial Intelligence Graduate School Program (POSTECH)), and the Technology Innovation Program (No. RS-2025-02952974) funded by the Ministry of Trade, Industry & Energy (MOTIE, Korea).

## References

- Cazenavette, G.; Wang, T.; Torralba, A.; Efros, A. A.; and Zhu, J.-Y. 2022. Dataset Distillation by Matching Training Trajectories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Cazenavette, G.; Wang, T.; Torralba, A.; Efros, A. A.; and Zhu, J.-Y. 2023. Generalizing Dataset Distillation via Deep Generative Prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Cisco. 2025. Cisco for Oil and Gas. <https://www.cisco.com/site/us/en/solutions/industries/energy/oil-gas/index.html>. Accessed: July 21, 2025.
- Cui, J.; Wang, R.; Si, S.; and Hsieh, C.-J. 2023. Scaling Up Dataset Distillation to ImageNet-1K with Constant Memory. In *Proceedings of the International Conference on Machine Learning*.
- Dahlhaus, R. 1996. On the Kullback-Leibler information divergence of locally stationary processes. *Stochastic processes and their applications*, 62(1): 139–168.
- Das, A.; Kong, W.; Sen, R.; and Zhou, Y. 2024. A Decoder-only Foundation Model for Time-series Forecasting.
- Ding, J.; Liu, Z.; Zheng, G.; Jin, H.; and Kong, L. 2024. CondTSF: One-Line Plugin of Dataset Condensation for Time Series Forecasting. *Proceedings of the Advances in Neural Information Processing Systems*.
- Ekambaram, V.; Jati, A.; Nguyen, N.; Sinthong, P.; and Kalagnanam, J. 2023. TSMixer: Lightweight MLP-Mixer Model for Multivariate Time Series Forecasting. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.
- Guo, Z.; Wang, K.; Cazenavette, G.; Li, H.; Zhang, K.; and You, Y. 2024. Towards Lossless Dataset Distillation via Difficulty-Aligned Trajectory Matching. In *Proceedings of the International Conference on Learning Representations*.
- Jin, W.; Zhao, L.; Zhang, S.; Liu, Y.; Tang, J.; and Shah, N. 2022. Graph Condensation for Graph Neural Networks. In *Proceedings of the International Conference on Learning Representations*.
- Khintchine, A. 1934. Korrelationstheorie der stationären stochastischen Prozesse. *Mathematische Annalen*, 109(1): 604–615.
- Lai, G.; Chang, W.-C.; Yang, Y.; and Liu, H. 2018. Modeling Long- and Short-Term Temporal Patterns with Deep Neural Networks. In *Proceedings of the International ACM SIGIR Conference on Research & Development in Information Retrieval*, SIGIR '18. New York, NY, USA: Association for Computing Machinery. ISBN 9781450356572.
- Lee, H.; Jung, K.; Kang, S.; Yoo, S.; Park, R.; Park, W.; Lee, H.-C.; Jheon, S.; Kim, Y.; Lee, H.; et al. 2022. VitalDB, a high-fidelity multi-parameter vital signs database in surgical patients. *Scientific Data*, 9(1): 1–11.
- Liu, S.; and Wang, X. 2023a. Few-Shot Dataset Distillation via Translative Pre-Training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Liu, S.; and Wang, X. 2023b. MGDD: A Meta Generator for Fast Dataset Distillation. In *Proceedings of the Advances in Neural Information Processing Systems*.
- Liu, X.; Xia, Y.; Liang, Y.; Hu, J.; Wang, Y.; BAI, L.; Huang, C.; Liu, Z.; Hooi, B.; and Zimmermann, R. 2023. LargeST: A Benchmark Dataset for Large-Scale Traffic Forecasting. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Proceedings of the Advances in Neural Information Processing Systems*, volume 36, 75354–75371. Curran Associates, Inc.
- Liu, Y.; Hu, T.; Zhang, H.; Wu, H.; Wang, S.; Ma, L.; and Long, M. 2024a. iTransformer: Inverted Transformers Are Effective for Time Series Forecasting. In *Proceedings of the International Conference on Learning Representations*.
- Liu, Z.; Hao, K.; Zheng, G.; and Yu, Y. 2024b. Dataset Condensation for Time Series Classification via Dual Domain Matching. *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.
- Liu, Z.; Zeng, C.; and Zheng, G. 2024. Graph Data Condensation via Self-Expressive Graph Structure Reconstruction. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.
- Luo, D.; and Xue, W. 2024. ModernTCN: A Modern Pure Convolution Structure for General Time Series Analysis. In *Proceedings of the International Conference on Learning Representations*.
- Maekawa, A.; Kobayashi, N.; Funakoshi, K.; and Okumura, M. 2023. Dataset Distillation with Attention Labels for Fine-Tuning BERT. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Maekawa, A.; Kosugi, S.; Funakoshi, K.; and Okumura, M. 2024. DiLM: Distilling Dataset into Language Model for Text-Level Dataset Distillation. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics*.
- Miller, J. A.; Aldosari, M.; Saeed, F.; Barna, N. H.; Rana, S.; Arpinar, I. B.; and Liu, N. 2024. A Survey of Deep Learning and Foundation Models for Time Series Forecasting. arXiv:2401.13912.
- Nguyen, T.; Chen, Z.; and Lee, J. 2021. Dataset Meta-Learning from Kernel Ridge-Regression. In *Proceedings of the International Conference on Learning Representations*.
- Nguyen, T.; Novak, R.; Xiao, L.; and Lee, J. 2021. Dataset Distillation with Infinitely Wide Convolutional Networks. *Proceedings of the Advances in Neural Information Processing Systems*.

- Nie, Y.; Nguyen, N. H.; Sinthong, P.; and Kalagnanam, J. 2023. A Time Series Is Worth 64 Words: Long-Term Forecasting with Transformers. *Proceedings of the International Conference on Learning Representations*.
- NOAA. 2025. Open Data Dissemination (NODD). <https://www.noaa.gov/information-technology/open-data-dissemination>. Accessed: July 21, 2025.
- Priestley, M. B. 1965. Evolutionary spectra and non-stationary processes. *Journal of the Royal Statistical Society: Series B (Methodological)*, 27(2): 204–229.
- Shin, D.; Shin, S.; and Moon, I.-C. 2023. Frequency Domain-based Dataset Distillation. In *Proceedings of the Advances in Neural Information Processing Systems*.
- Stitsyuk, A.; and Choi, J. 2025. xPatch: Dual-Stream Time Series Forecasting with Exponential Seasonal-Trend Decomposition. arXiv:2412.17323.
- Wang, S.; Wu, H.; Shi, X.; Hu, T.; Luo, H.; Ma, L.; Zhang, J. Y.; and ZHOU, J. 2024. TimeMixer: Decomposable Multiscale Mixing for Time Series Forecasting. In *The Twelfth International Conference on Learning Representations*.
- Wang, T.; Zhu, J.-Y.; Torralba, A.; and Efros, A. A. 2018. Dataset Distillation. *arXiv Preprint arXiv:1811.10959*.
- Wei, X.; Cao, A.; Yang, F.; and Ma, Z. 2023. Sparse Parameterization for Epitomic Dataset Distillation. In *Proceedings of the Advances in Neural Information Processing Systems*.
- Wiener, N. 1930. Generalized harmonic analysis. *Acta mathematica*, 55(1): 117–258.
- Woo, G.; Liu, C.; Kumar, A.; Xiong, C.; Savarese, S.; and Sahoo, D. 2024. Unified Training of Universal Time Series Forecasting Transformers.
- Wu, H.; Hu, T.; Liu, Y.; Zhou, H.; Wang, J.; and Long, M. 2022. TimesNet: Temporal 2D-Variation Modeling for General Time Series Analysis. *arXiv Preprint arXiv:2210.02186*.
- Xu, Z.; Zeng, A.; and Xu, Q. 2024. FITS: Modeling Time Series with 10k Parameters. In *The Twelfth International Conference on Learning Representations*.
- 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*.
- Zhang, H.; Li, S.; Wang, P.; Zeng, D. G.; and Shiming. 2024. M3D: Dataset Condensation by Minimizing Maximum Mean Discrepancy. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Zhao, B.; and Bilal, H. 2021. Dataset Condensation with Gradient Matching. In *Proceedings of the International Conference on Learning Representations*.
- Zhao, B.; and Bilal, H. 2023. Dataset Condensation with Distribution Matching. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*.
- 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. *Proceedings of the AAAI Conference on Artificial Intelligence*, 11106–11115.
- Zhou, Y.; Nezhadarya, E.; and Ba, J. 2022. Dataset Distillation Using Neural Feature Regression. In *Proceedings of the Advances in Neural Information Processing Systems*.