# LTSM-Bundle: A Toolbox and Benchmark on Large Language Models for Time Series Forecasting

Anonymous authors
Paper under double-blind review

### **Abstract**

Time Series Forecasting (TSF) has long been a challenge in time series analysis. Inspired by the success of Large Language Models (LLMs), researchers are now developing Large Time Series Models (LTSMs), universal transformer-based models that use autoregressive prediction to improve TSF. However, training LTSMs on heterogeneous time series data poses unique challenges, including diverse frequencies, dimensions, scalability, and patterns across datasets. Recent efforts have studied and evaluated various design choices aimed at enhancing LTSM training and generalization capabilities. Despite progress in both paradigms, there is no unified framework for systematically evaluating models and design choices across them. However, these design choices are typically studied and evaluated in isolation and are not compared collectively. In this work, we introduce LTSM-Bundle, a comprehensive toolbox and benchmark for training LTSMs, spanning pre-processing techniques, model configurations, and dataset configurations. Modularized and benchmarked LTSMs from multiple dimensions, encompassing prompting strategies, tokenization approaches, training paradigms, base model selection, data quantity, and dataset diversity. By combining the most effective design choices, the combination achieves state-of-the-art zero-shot and few-shot performance while providing a reproducible foundation for evaluating both traditional LSF models and emerging LTSMs. Our source code is available at https://anonymous.4open.science/r/LTSM-bundle-5B70/

### 1 Introduction

Time series forecasting (TSF) is a long-standing task in time series analysis, aiming to predict future values based on historical data points. Over the decades, TSF has transitioned from traditional statistical methods (Ariyo et al., 2014) to machine learning (Friedman, 2001), and more recently, to deep learning approaches (Elsworth & Güttel, 2020; Livieris et al., 2020). In particular, transformers (Vaswani et al., 2017), which are often regarded the most powerful architecture for sequential modeling, have demonstrated superior performance in TSF, especially for long-term forecasting (Wu et al., 2021; Zhou et al., 2021; Nie et al., 2022; Woo et al., 2022; Kitaev et al., 2020). Moving forward, inspired by the remarkable capabilities of Large Language Models (LLMs), many researchers have begun to explore Large Time Series Models (LTSMs) as the natural next phase, seeking to train universal transformer-based models for TSF (Woo et al., 2024; Garza & Mergenthaler-Canseco, 2023; Dooley et al., 2024; Rasul et al., 2024; Das et al., 2023; Gruver et al., 2024; Chang et al., 2023; Zhou et al., 2023; Jin et al., 2023).

Unlike textual data, where tokens typically hold semantic meanings transferable across documents, time series data exhibits high heterogeneity, presenting unique challenges for LTSM training. Across different datasets, time series often have diverse frequencies (such as hourly and daily), dimensions (in terms of varying numbers of variables) and patterns (where, for example, traffic time series may differ significantly from electricity data). This diversity not only poses difficulties in training an LTSM to fit all the datasets but also impedes the model's generalization to unseen data.

Recent endeavors have proposed various innovative designs to enhance the training and generalization capability of LTSMs. To name a few, (i) in terms of pre-processing, prompting strategies have been proposed

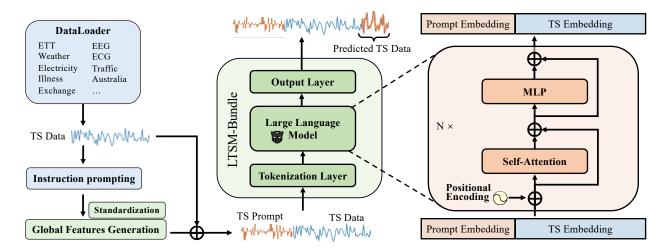


Figure 1: An overview of the key design choices supported by the LTSM-Bundle framework. The library is a general-purpose benchmarking toolbox that standardizes data preprocessing, tokenization strategies, training paradigms, and evaluation pipelines across traditional LSF models, LLM-based models, and pretrained LTSMs, rather than being limited to LLM-centric designs.

to generate dataset-specific prompts (Jin et al., 2023), while various tokenization strategies have been studied for converting time series into tokens to be inputted into transformer layers (Zhou et al., 2023; Ansari et al., 2024); (ii) for the model configurations, prior research has involved reusing weights from pre-trained language models and adapting them to downstream tasks (Zhou et al., 2023); (iii) regarding dataset configurations, different datasets have been utilized for training purposes (Garza & Mergenthaler-Canseco, 2023; Zhou et al., 2023; Chang et al., 2023). Despite these advances, evaluating models across both paradigms remains challenging. Existing studies often adopt inconsistent pre-processing pipelines, heterogeneous tokenization strategies, and non-uniform evaluation metrics, making it difficult to draw fair comparisons or identify best practices, where these designs are typically studied and evaluated in isolation. There is no existing package that integrates these components or benchmarks them collectively. This makes it difficult to understand, select, and combine these components to effectively train LTSMs in practice.

To address the gaps, we introduce LTSM-Bundle, a comprehensive toolbox and benchmark for training LLMs for time series forecasting, spanning pre-processing techniques, model configurations, and dataset configurations, as depicted in Figure 1. The goal of LTSM-Bundle is designed to serve as a foundation for benchmarking both current and emerging LTSMs and TSFMs, while maintaining full transparency about its current model and dataset coverage. We modularize and benchmark LTSMs under the same settings and from multiple dimensions, including prompting strategies, tokenization approaches, training paradigms, base model selection, data quantity, and dataset diversity. In addition to existing components, we introduce time series prompt, a statistical prompting strategy tailored for time series data, as one of our benchmarking components. It generates prompts by extracting global features from the training dataset, providing robust statistical descriptions for heterogeneous data.

Through extensive benchmarking with LTSM-Bundle, we combine the most effective design choices identified in our study for training LTSMs. Our empirical results suggest that the identified combination produces superior zero-shot and few-shot (with 5% training data) performances compared to the state-of-the-art LTSMs in benchmark data sets. Additionally, even with just 5% of the data, our result is comparable to the strong baselines trained on the full training data, demonstrating the practical value of LTSM-Bundle in developing and training LTSMs. In summary, we have made the following contributions:

• We present LTSM-Bundle, a comprehensive toolbox and benchmark for LTSMs. LTSM-Bundle not only includes various components with easy-to-use interfaces for training LTSMs but is also integrated with TDengine, a state-of-the-art time series database, to build up an end-to-end training pipeline from time series data storage to report visualization and generation.

- We perform systematic analysis with LTSM-Bundle. Our analysis yields numerous insightful observations, paving the path for future research endeavors.
- Tokenization design critically affects cross-architecture transferability. Our results show that different tokenization strategies have a significant impact on forecasting performance across diverse model architectures.
- Model size is not always correlated with accuracy. We find that smaller and medium-sized models can
  achieve competitive or even superior performance compared to larger models, particularly on long-horizon
  forecasting tasks.

# 2 Notations and Problem Formulation

We denote a multi-variate time series as  $\mathbf{Z} = \{\mathbf{z}_1, \mathbf{z}_2, ..., \mathbf{z}_T\}$ , where  $\mathbf{z}_t \in \mathbb{R}^d$  is a vector of multivariate variable with dimension d, and T is the total number of timestamps. We typically partition  $\mathbf{Z}$  chronologically to create training, validation, and testing sets, denoted as  $\mathbf{Z}^{\text{train}} = \{\mathbf{z}_1, \mathbf{z}_2, ..., \mathbf{z}_{T^{\text{train}}}\}$ ,  $\mathbf{Z}^{\text{val}} = \{\mathbf{z}_{T^{\text{train}}+1}, \mathbf{z}_{T^{\text{train}}+2}, ..., \mathbf{z}_{T^{\text{train}}}\}$ , and  $\mathbf{Z}^{\text{test}} = \{\mathbf{z}_{T^{\text{train}}+T^{\text{val}}+1}, \mathbf{z}_{T^{\text{train}}+T^{\text{val}}+2}, ..., \mathbf{z}_{T}\}$ , where  $T^{\text{train}}$  and  $T^{\text{test}}$  denote the number of timestamps for training and validation, respectively. In traditional TSF, we aim to train a model using  $\mathbf{Z}^{\text{train}}$  such that, on  $\mathbf{Z}^{\text{test}}$ , given the observations from the historical P timestamps  $\mathbf{X} = \{\mathbf{z}_{t_1}, \mathbf{z}_{t_2}, ..., \mathbf{z}_{t_P}\}$ , the model can accurately predict the values of future Q timestamps  $\mathbf{Y} = \{\mathbf{z}_{t_{P+1}}, \mathbf{z}_{t_{P+2}}, ..., \mathbf{z}_{t_{P+Q}}\}$ , where  $\mathbf{X}$  and  $\mathbf{Y}$  are sub-sequences of  $\mathbf{Z}^{\text{test}}$ . In our work, the LTSMs are trained by minimizing mean square error loss  $\mathcal{L}(\text{LTSM}(\mathbf{X}), \mathbf{Y})$  between the given sub-sequences.

We focus on training LTSMs, where the objective is to develop a model that performs well across various test sets, denoted as  $\mathcal{Z}^{\text{test}} = \{\mathbf{Z}_1^{\text{test}}, \mathbf{Z}_2^{\text{test}}, ..., \mathbf{Z}_N^{\text{test}}\}$ , where N represents the number of datasets for testing. Each  $\mathbf{Z}_i^{\text{test}}$  may originate from a distinct domain, with different lengths, dimensions, and frequencies. The training sets for an LTSM may comprise training data associated with  $\mathcal{Z}^{\text{test}}$  or data from other sources, provided they are not included in  $\mathcal{Z}^{\text{test}}$ . Training LTSMs presents a notable challenge compared to traditional TSF models due to the inherent difficulty in accommodating diverse patterns across datasets, often necessitating specialized designs. Nevertheless, it also offers opportunities to transfer knowledge from existing time series to new scenarios.

# 3 LTSM-Bundle Library

### 3.1 Package Design

We showcase the system overview of LTSM-Bundle in Figure 2. LTSM-Bundle is a modular toolkit that supports the complete life-cycle of large time-series models (LTSMs), from raw data ingestion to deployment-ready evaluation. The framework is organized around four interoperable subsystems, all exposed through a unified API that eliminates boilerplate engineering and accelerates experimentation. First, TS Tokenizing converts multivariate time series into token sequences via linear and dynamic schemes that preserve both global trends and local temporal dynamics, making the data directly consumable by Transformer-style backbones. Second, TS Prompting embeds task instructions and statistical context—through hard, soft, and statistics-aware prompts—into the token stream, enabling zero- and few-shot adaptation to forecasting, anomaly detection, and classification tasks. Third, the **Data Processing** layer supplies scalable loaders, windowing utilities, and feature-engineering pipelines that abstract away dataset idiosyncrasies and seamlessly handle large benchmarks. Fourth, LTSM Training offers a uniform optimization interface for fine-tuning or training from scratch a range of backbone architectures and parameter scales, with built-in support for curriculum learning, transfer learning, and large-scale hyperparameter sweeps. These subsystems are orchestrated by a reproducible workflow engine that chains tokenizing, prompting, processing, and training steps into end-to-end pipelines. The toolkit further provides a library of loss functions, data-augmentation routines, and evaluation metrics, alongside visualization and reporting utilities that generate publication-ready artifacts. Collectively, LTSM-Bundle delivers a comprehensive and extensible platform for constructing, analyzing, and deploying large time-series models, thereby lowering the barrier to reproducible research and accelerating industrial adoption.

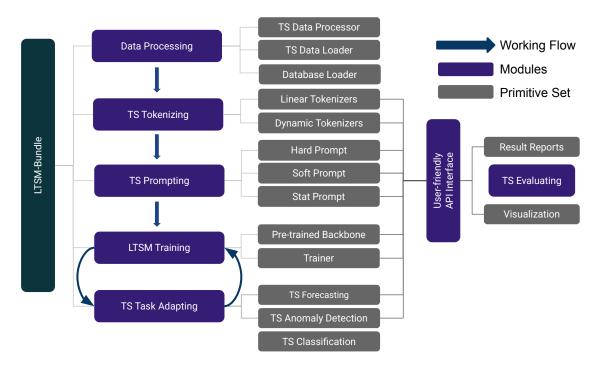


Figure 2: System Overview of LTSM-Bundle library. LTSM-Bundle provides an end-to-end training and evaluation pipeline constructed from data preprocessing to visualization. It also provides a database reader to better incorporate with large-scale datasets.

### 3.2 Package Interface

The LTSM-Bundle package is implemented with a scikit-learn-like interface to easily train a customized LTSM. To increase the flexibility of LTSM-Bundle in utilizing a broader spectrum of backbone models and training paradigms, we have integrated it with the Huggingface Transformers package<sup>1</sup>. This integration allows for the incorporation of diverse pre-trained weights and supports various training approaches, enhancing the overall adaptability of our framework. Researchers and practitioners only need to provide their own time series data and specify the chosen training configurations; then, every training pipeline can be created using the LTSM-Bundle package. Moreover, LTSM-Bundle supports linking with time series vector database, TDengine<sup>2</sup>, a state-of-the-art time series database, to build an end-to-end training pipeline from time series data storage to report visualization and generation. Without struggling to build up extra efforts on clear report generation and visualization, LTSM-Bundle automatically generates all results with an easy-to-read template at the top of TDengine.

# 4 Benchmarking LTSM Training

We benchmark existing LTSM components by involving and coordinating them into four fundamental components of LTSM-Bundle package. We aim to answer the following research questions: 1) How do to-kenization methods and prompting techniques impact model convergence? 2) How do base model selection and the training paradigms impact the time series forecasting performance? and 3) How do different dataset configurations impact the model generalization?

To answer the questions, we evaluate the impact of each component individually by keeping the rest of the components fixed. First, we fix the base model and dataset configurations with smaller model sizes, limited data quantities, and reduced data diversity, excluding prompt tokenization, to identify the optimal prompting

https://github.com/huggingface/transformers

<sup>&</sup>lt;sup>2</sup>https://github.com/taosdata/TDengine

strategy. Next, we incorporate the best prompting strategy with the same base model selection and dataset configuration to assess the tokenization methods. Afterward, we keep all the components constant except the base model to study the impact of different model initialization and training strategies. Finally, with the best tokenization and prompting methods, we select a list of candidate base models following the guidelines from the previous step. We also control the quantity and diversity of the training data to assess their impacts on model generalizability and prediction performances.

We follow the experimental settings outlined in Timesnet (Wu et al., 2022) and Time-LLM (Jin et al., 2023), employing the unified evaluation framework. The input time series length  $\mathcal{E}$  is set to 336, with four different prediction lengths in  $\{96, 192, 336, 720\}$ . Evaluation metrics include mean square error (MSE) and mean absolute error (MAE). We also calculate the average scores among all prediction horizons. The results highlighted in **red** represent the **best** performance and highlighted in **blue** represent the **second best** performance. The details of hyperparameter settings of our banchmarking experiments are in Appendix B, respectively.

Our evaluations use widely adopted benchmarks: the ETT series (ETTh1, ETTh2, ETTm1, ETTm2) (Zhou et al., 2021), Traffic, Electricity, Weather, and Exchange-Rate datasets (Lai et al., 2018; Wu et al., 2022). ETT comprises four subsets—two with hourly data (ETTh) and two with 15-minute data (ETTm)—each containing seven features from July 2016 to July 2018. The Traffic dataset provides hourly road occupancy rates from San Francisco freeways (2015–2016); the Electricity dataset records hourly consumption for 321 clients (2012–2014); the Weather dataset offers 21 indicators every 10 minutes in Germany during 2020; and the Exchange-Rate dataset includes daily rates for eight countries (1990–2016). For more details, refer to Appendix A.

We first train our framework on the diverse time series data collection using LTSM-Bundle package and then assess the best combination identified on joint learning and zero-shot transfer learning to different domains of time series knowledge. For clarity, we use the term "LTSM-Bundle" to represent the best practice of the combination under the described experiment settings in each of the following sections.

# 4.1 Pre-processing: Instruction Prompts

The pre-processing step plays a crucial role in enabling LLM-based models to better adapt to time series datasets. In this section, we present a detailed analysis aimed at recommending the most effective pre-processing prompting strategy to compose LTSM-bundle. Instruction prompts enhance the effectiveness of LTSM training by providing auxiliary information. This prompt helps the model adjust its internal state and focus more on relevant features in different domains of the dataset, thereby improving learning accuracy. With the aid of prompts, LTSM aims to optimize forecasting ability across diverse dataset domains. We explore two types of prompts: the *Text Prompts* (Jin et al., 2023) written in task-specific information, and the *time series prompts* developed by global features of time series data. This comparison determines the most effective prompt type for LTSM training.

Time Series Prompts Time series prompts aim to capture the comprehensive properties of time series data. Unlike instruct prompts, they are derived from a diverse set of global features extracted from the entire training dataset. This approach ensures a robust representation of the underlying dynamics, in addition to enhancing model performance. The time series prompts are generated by extracting global features from each variate of the time series training data. The extracted global features are specified in Appendix D. After extracting the global features, we proceed to standardize their values across all varieties and instances within the dataset. This standardization is crucial to prevent the overflow issue during both training and inference stages. Let  $\mathbf{P} = \{p_1, \dots, p_M\}$  denote the global features of  $\mathbf{Z}$  after the standardization, where  $\mathbf{p}_t \in \mathbb{R}^d$ . Subsequently,  $\mathbf{P}$  serves as prompts, being concatenated with each timestamp  $\mathbf{X}$  derived from the time series data. Consequently, the large time series models take the integrated vector  $\tilde{\mathbf{X}} = \mathbf{P} \cup \mathbf{X} = \{\mathbf{p}_1, \dots, \mathbf{p}_M, \mathbf{z}_{t_1}, \mathbf{z}_{t_2}, \dots, \mathbf{z}_{t_P}\}$  as input data throughout both training and inference phases, as illustrated in Figure 3. The time series prompts are generated separately for the training and testing datasets, without leaking the information from testing data. We leverage the package<sup>3</sup> to generate the time series prompts.

<sup>3</sup>https://github.com/thuml/Time-Series-Library

Table 1: Performance of different prompting strategies

Metric	Input	ETTh1	ETTh2	ETTm1	ETTm2	Traffic	Weather	Electricity	Avg.
MSE	No prompt	0.308	0.237	0.367	0.157	0.306	0.177	0.148	0.243
	TS prompt	0.307	0.234	0.285	0.155	0.305	0.172	0.145	0.229
	Text prompt	0.319	0.241	0.490	0.190	0.345	0.212	0.185	0.283
MAE	No prompt	0.375	0.325	0.411	0.258	0.272	0.232	0.246	0.303
	TS prompt	0.377	0.326	0.369	0.266	0.279	0.242	0.247	0.301
	Text prompt	0.386	0.329	0.476	0.289	0.326	0.269	0.299	0.339

Table 2: Performance of linear and time series tokenization

Metric	Tokenizer	ETTh1	ETTh2	ETTm1	ETTm2	Traffic	Weather	Exchange	Electricity	Avg.
MSE	Linear tokenizer Time series tokenizer	0.301 1.798	0.228 0.855	$0.261 \\ 1.671$	$0.149 \\ 0.625$	0.300 2.199	0.163 0.983	$\frac{0.058}{3.729}$	0.140 2.206	0.214 1.663
MAE	Linear tokenizer Time series tokenizer	0.372 1.057	0.319 0.606	0.346 0.991	0.265 0.488	0.268 1.083	0.230 0.619	$\frac{0.173}{1.495}$	0.241 1.108	0.281 0.895

Experimental Results We begin by evaluating the effectiveness of instruction prompts. Specifically, we assess two distinct types of instruction prompts, both initialized by the same pre-trained GP2-Medium weights within the context of commonly used linear tokenization. The experimental results are shown in Table 1. Our observations suggest that ① statistical prompts outperform traditional text prompts in enhancing the training of LTSM models with up to 8% lower MAE scores. Additionally, ② it is observed that the use of statistical prompts results in superior performance compared to scenarios where no prompts are employed, yielding up to 3% lower MSE scores. The superiority of statistical prompt is evident in the more effective leveraging of LTSM capabilities, leading to improved learning outcomes across various datasets. Based on the above observations, we select time series prompts as the focus in the following analysis and incorporate them into LTSM-bundle.

#### 4.2 Pre-processing: Tokenizations

In addition to employing instructional prompts to enhance generalization in LTSM training, this section provides a detailed analysis aimed at identifying the most effective tokenization strategy for LTSMs. We explore two distinct tokenization approaches – linear tokenization (Zhou et al., 2023) and time series tokenization (Ansari et al., 2024) – to determine the superior method for training LTSM models.

**Details of Tokenization** To harness the power of LLMs, a prevalent strategy involves mapping time series values to tokens (Zhou et al., 2023; Jin et al., 2023). However, converting time series data to natural language formats for LLMs is not trivial, as LLMs are pre-trained with predetermined tokenizers designed for NLP datasets. However, this implies that time series data cannot be directly fed into LLMs for training on forecasting purposes; it requires a specialized transformation of the time series data into specific indices suitable for processing by the LLMs. In this manner, we utilize two advanced types of tokenizations, linear tokenization and time series tokenization, to better evaluate their effectiveness in transferring data for training LTSMs. Specifically, the linear tokenization (Zhou et al., 2023) leverages one trainable linear layer  $f: \mathbb{R}^{\mathcal{E}} \to \mathbb{R}^K$  to transfer time series numbers to specific tokens, where  $\mathcal{E}$  denotes time series length, and K refers to input size of pre-trained LLM backbone. The trainable time series tokenization (Ansari et al., 2024) aims to covert continuous time series data into discrete tokens by scaling and quantizing their values to the specific number of token bins with a given Dirichlet function.

**Experimental Results** We investigate the impact of two tokenization methods on training LTSMs. By comparing different tokenization strategies, we aim to identify which approach best complements the LTSM architecture, enhancing its ability to process and learn from complex and multi-domain datasets. Specifically, we conduct experiments comparing linear tokenization and time series tokenization, utilizing pre-trained GPT-2-medium models along with time series prompts. The experimental results shown in Table 2 demonstrates

	Metric		M	SE		MAE			
	Predict length	96	192	336	720	96	192	336	720
TS prompt	From scratch	0.325	0.296	0.323	0.355	0.355	0.375	0.374	0.409
	LoRA fine-tuning	0.343	0.381	0.399	0.466	0.374	0.403	0.426	0.478
	Fully fine-tuning	0.229	0.260	0.297	0.351	0.301	0.324	0.354	0.397
Text prompt	From scratch	0.494	0.434	0.597	0.485	0.463	0.438	0.512	0.475
	LoRA fine-tuning	0.347	0.379	0.406	0.473	0.373	0.404	0.431	0.484
	Fully fine-tuning	0.294	0.286	<mark>0.353</mark>	0.358	0.330	0.338	0.378	0.429

Table 3: Performance of learning from scratch, LoRA fine-tuning, and fully fine-tuning

strate that linear tokenization more effectively facilitates the training process of LTSM compared to time series tokenization. In our study, we focus on smaller and more accessible training data. Under these circumstances, we observe that a linear tokenization is a more suitable choice than a time series tokenization, as the pre-trained time series tokenization may not have the transferability toward different model architecture. Unlike textual tokenization, the time series tokenization is determined by a pre-trained time series dataset with a designated LTSM architecture In summary, (3) linear tokenization is flexible and adaptive for different settings of LTSM training compared to time series tokenization under a smaller amount of training data.

### 4.3 Model Configuration: Training Paradigm

Different training paradigms exhibit unique characteristics that influence how well LLMs fit a specific training dataset. In this section, we explore three distinct training paradigms, fully fine-tuning, training from scratch, and LoRA (Hu et al., 2021), to identify the most effective approaches for training the LTSM framework.

**Training Paradigm.** In the full fine-tuning paradigm, we utilize the pre-trained weights of each base LLM, which finetune all parameters using the given time series dataset. Conversely, in the training-from-scratch paradigm, we only preserve the original model architecture but initialize all parameters anew before training with the time series dataset. In the LoRA paradigm, we employ low-rank adapters on base LLMs.

Experimental Results. We assess the effectiveness of the training paradigm under the settings of time series prompt and text prompt usage. Table 3 presents the results of various training strategies using GPT-2-Medium as the backbone. In general, the experimental results indicate that full fine-tuning is the most effective strategy for training the LTSM framework whether leveraging time series prompts or text prompts. Based on the results, we summarize the observations as follows. 4 Although training-from-scratch achieves competitive performance compared to full fine-tuning, the large number of trainable model parameters may lead to overfitting, ultimately degrading performance. 5 Fully fine-tuning paradigm leads to the best performance with up to 11% of improvement on MSE and up to 17% of improvement on MAE under the length of {96, 192, 336}, and performance competitive under the length of 720. Training LTSM-bundle under the full fine-tuning paradigm is recommended, as it converges twice as fast as training from scratch, ensuring efficient and effective forecasting.

# 4.4 Model Configuration: Base Model Selection

Base Model Candidates As for the base models of our framework, we leverage four different pre-trained models, including GPT-2-small, GPT-2-medium, GPT-2-large (Radford et al., 2019), and Phi-2 (Javaheripi & Bubeck, 2023). GPT-2 employs a transformer architecture with up to 48 layers, and it is trained on a diverse corpus of internet text, resulting in a model size of 124M (small), 355M (medium), and 774M (large) parameters. Phi-2 also uses a transformer-based architecture but emphasizes high-quality ("textbook-quality") data, comprising 2.7 billion parameters. Despite its smaller size compared to the largest contemporary models, Phi-2 incorporates innovative scaling techniques to optimize performance. Different from the absolute positional encoding used by GPT-2, Phi-2 employs relative positional encoding, which considers the pairwise distance between each token pair for encoding position information of tokens. Following the settings

Table 4: Performance of different backbones

Metric   MSI			SE	MAE				
GPT-2	96	192	336	720	96	192	336	720
Small	0.252	0.306	0.316	0.352	0.313	0.363	0.367	0.400
Medium	0.229	0.260	0.297	0.351	0.301	0.324	0.354	0.397
Large	0.224	0.257	0.301	0.358	0.292	0.322	0.356	0.399

Table 5: Performance of different down-sampling rates

Metric		$\mathbf{M}$	SE	MA			AE	Æ		
DS rate	96	192	336	720	96	192	336	720		
2.5%	0.227	0.268	0.308	0.369	0.294	0.335	0.365	0.415		
5%	0.229	0.261	0.297	0.350	0.301	0.324	0.354	0.397		
10%	0.233	0.260	0.291	0.350	0.289	0.323	0.348	0.396		

in (Zhou et al., 2023), we utilize the top three self-attention layers of every pre-trained model as our backbone structure in LTSM-bundle framework.

Experimental Results We explore the impact of using different pre-trained LLM weights as backbones in LTSM models, with the goal of identifying the most suitable pre-trained LLM weights for processing time series data. The findings are detailed in Table 4. We assess the performance of different backbones with time series prompts under the fully fine-tuning paradigm. We summarize our observations as follows: 6 GPT-2-Small demonstrates a performance improvement of up to 2% in relatively long-term forecasting (i.e., 336 and 720 hours) compared to the GPT-2-Large model. 7 GPT-2-Medium outperforms GPT-2-Large in relatively short-term forecasting (i.e., 96 and 192 hours), as larger models may be prone to overfitting during training, degrading forecasting performance.

While our benchmarking results in Table 4 indicate that variations in the number of parameters within the same architecture have minimal impact on performance, we did not limit our analysis to a single model size alone. To ensure the rigor of our evaluation, we compared Phi-2—an alternative model architecture—with GPT-2 models of varying sizes (large, medium, and small) using different prompting methods, as detailed in Table 18 and Table 19 of Appendix F. The results show that GPT-2 Small and Medium obtains higher performance than Phi-2 on both time series prompt and textual instruction prompt settings. Based on the above findings, we recommend incorporating GPT-2-Medium or GPT-2-Small as the backbone of LTSM-bundle.

### 4.5 Dataset configuration: Quantity

The quantity of datasets is often the key to the success of LLMs due to the consistent semantic meaning of tokens. Nevertheless, time series tokens are less informative and semantically meaningful compared to natural language tokens. In this section, we investigate the impact of data quantity to determine whether the principle that more training data leads to better LTSMs.

Quantity Configuration We conduct the down-sampling strategies to study the impact of data quantity on prediction performance. Specifically, each time series in the training data are periodically down-sampled along the timestamps to reduce the granularity of the entire time series while maintaining the general pattern. Each dataset is split into training, validation, and testing sets, and then down-sampling is applied to the training set for model training. We compare the models trained with 10%, 5%, and 2.5% of the full-size time series in the training set. In the following experiment section, we annotate partial training data usage as few-shot training.

**Experimental Results** Table 5 lists the results of models trained with different data quantities under GPT-2 Medium as the model backbone. The model trained with 5% and 10% down-sampled data leads to the best result compared to 2.5%. (8) We observe that increasing the amount of data does not positively correlate with improved model performance. The rationale is that increasing data points enhances the granularity of

Table 6: Average performance with different downsampling under different domain LTSM-Bundle

(MAE/MSE)	1 dataset	2 datasets	4 datasets	8 datasets
2.5%			0.396/0.357	
5%		0.415/0.436	0.383/0.341	
10%	0.414/0.375	0.415/0.440	0.394/0.355	0.348/0.288

Table 7: Performance of LTSM trained on different numbers of datasets

		1 dataset	2 datasets	3 datasets	4 datasets	5 datasets	6 datasets	7 datasets	8 datasets
	96	0.333	0.366	0.269	0.276	0.232	0.229	0.227	0.215
MSE	192	0.394	0.440	0.309	0.318	0.271	0.267	0.269	0.254
MSE	336	0.403	0.427	0.356	0.351	0.302	0.325	0.308	0.302
	720	0.478	0.511	0.436	0.419	0.373	0.364	0.369	0.361
	96	0.351	0.351	0.327	0.329	0.298	0.294	0.292	0.282
MAE	192	0.407	0.395	0.363	0.368	0.333	0.332	0.334	0.323
MAL	336	0.420	0.420	0.401	0.392	0.361	0.384	0.371	0.362
	720	0.464	0.494	0.466	0.445	0.423	0.411	0.419	0.411

time series but may reduce the model's generalization ability, while excessive down-sampling loses critical information, hindering pattern learning. Thus, optimizing model performance requires carefully balancing the amount of training data with its diversity. To explore this balance, we conducted experiments using different downsampling rates across various time series datasets. After comparing with the different downsampling rates (i.e., 2.5%, 10%, and 25% in Table 5 and Table 6). We use 5% of the dataset to benchmark the LTSM-Bundle family across diverse time series data, as increasing the dataset size beyond this point results in only marginal performance improvements while significantly increasing training time (i.e., 10% increase doubles the computational cost but yields only a slight enhancement in forecasting performance).

#### 4.6 Dataset Configuration: Diversity

Impact of Dataset Diversity Recall that we utilized eight datasets for training purposes, encompassing ETT variant, Traffic, Electricity, Weather, and Exchange. Here, we focus on evaluating the performance of LTSM models when trained with subsets of these datasets. Specifically, we employ the first M datasets from the aforementioned list for training, where  $M \in \{1, 2, ..., 8\}$ . For instance, when  $M = \{1, 3\}$  solely ETTh1 is utilized for training; when  $M = \{1, 3\}$  ETTh1, ETTh2, ETTm1, ETTm2, and Weather are utilized. Subsequently, we evaluate the trained model's performance across all datasets to understand the impact of dataset diversity.

Experimental Results Table 7 summarizes the results. (9) Augmenting dataset diversity generally leads to improved performance. This is expected because more diverse data has the potential to enhance the generalization capabilities of LTSMs across various patterns. We conclude two reasons: (1) Although originating from different domains with distinct characteristics, datasets may share underlying knowledge that can be transferred, enhancing model generalization and performance. (2) Prompting strategies, particularly time series prompts, can further facilitate knowledge transfer by guiding the model to implicitly learn which information to retain and which to discard.

### 5 Comparison with Baselines

Based on the observations in Section 4, we identify a strong combination using LTSM-Bundle with the settings as follows: (1) Base model backbone: GPT-2-Medium, (2) Instruction prompts: the time series prompts, (3) Tokenization: linear tokenization, and (4) Training paradigm: fully fine-tuning. We compare this combination against SoTA TSF models on zero-shot and few-shot settings.

	LTSM-	Bundle	TIME	E-LLM	GPT	Γ4TS	LLM	Time	DLi	near	Patcl	hTST	Time	$_{ m esNet}$
Metric	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
$\rm ETTh1 \rightarrow ETTh2$	0.319	0.402	0.353	0.387	0.406	0.422	0.992	0.708	0.493	0.488	0.380	0.405	0.421	0.431
$\text{ETTh1} \rightarrow \text{ETTm2}$	0.312	0.406	0.273	0.340	0.325	0.363	1.867	0.869	0.415	0.452	0.314	0.360	0.327	0.361
$\text{ETTm1} \rightarrow \text{ETTh2}$	0.306	0.391	0.381	0.412	0.433	0.439	0.992	0.708	0.464	0.475	0.439	0.438	0.457	0.454
$\mathrm{ETTm1}  ightarrow \mathrm{ETTm2}$	0.217	0.319	0.268	0.320	0.313	0.348	1.867	0.869	0.335	0.389	0.296	0.334	0.322	0.354
$\mathrm{ETTm2}  ightarrow \mathrm{ETTh2}$	0.314	0.393	0.354	0.400	0.435	0.443	1.867	0.869	0.455	0.471	0.409	0.425	0.435	0.443
$ETTm2 \rightarrow ETTm1$	0.403	0.430	0.414	0.438	0.769	0.567	1.933	0.984	0.649	0.537	0.568	0.492	0.769	0.567

Table 8: Zero-shot performance. "LTSM-Bundle" denotes the best combination of LTSM training components

### 5.1 Experimental Settings

We follow the same settings as in Time-LLM (Jin et al., 2023). Specifically, for zero-shot experiments, we test the model's cross-domain adaptation under the long-term forecasting scenario and evaluate it on various cross-domain scenarios utilizing the ETT datasets. The hyperparameter settings of training LTSM-Bundle are in Appendix B. For the few-shot setting, we train our LTSM-Bundle on 5% of the data and compare it with other baselines under the 5% as well. We cite the performance of other models when applicable (Zhou et al., 2023). Furthermore, we compare LTSM-Bundle trained on 5% training data against baselines trained on the full training set. Our findings in Appendix F indicate that LTSM-Bundle achieves comparable results, further underscoring its superiority.

The baseline methods consist of various Transformer-based methods, including PatchTST (Nie et al., 2022), ETSformer (Woo et al., 2022), Non-Stationary Transformer (Liu et al., 2022), FEDformer (Zhou et al., 2022), Autoformer (Chen et al., 2021), Informer (Zhou et al., 2021), and Reformer (Kitaev et al., 2020). Additionally, we evaluate our model against recent competitive models like Time-LLM (Jin et al., 2023), TEST (Sun et al., 2023), LLM4TS (Chang et al., 2023), GPT4TS (Zhou et al., 2023), DLinear (Zeng et al., 2023), TimesNet (Wu et al., 2022), and LightTS (Campos et al., 2023). More details of the baseline methods can be found in Section 6.

#### 5.2 Zero-shot and Few-shot Results

Zero-shot Performance In the zero-shot learning experiments shown in Table 8 shows that the best component combination from benchmarking LTSM-Bundle consistently delivers superior performance across various cross-domain scenarios using the ETT datasets. For example, in the ETTh1 to ETTh2 dataset transfer task, LTSM-Bundle achieves an MSE of 0.319 and an MAE of 0.402, outperforming all other methods, including TIME-LLM, GPT4TS, and DLinear. Similarly, in the ETTm1 to ETTm2 dataset transfer scenario, LTSM-Bundle records the lowest MSE and MAE scores of 0.217 and 0.319, showing its strong generalization capability across different domains. The consistent improvements across transfer tasks of LTSM-Bundle in zero-shot learning.

Few-shot Performance Table 9 presents the performance of the best component combination from benchmarking compared to the baseline models in the few-shot setting, utilizing 5% of the training data. Notably, LTSM-Bundle exhibits a significant advantage over both traditional baselines and existing LTSMs. Across the 7 datasets, LTSM-Bundle outperforms all baselines regarding MSE in 5 datasets and regarding MAE in 4 datasets. Moreover, LTSM-Bundle achieves the top rank 40 times among the reported results. These findings underscore the effectiveness of our model in few-shot scenarios, where it demonstrates high accuracy even with limited training data. Its capability to excel with minimal data not only highlights its adaptability but also its potential for practical applications, particularly in contexts where data availability is constrained. All further and full versions of results on the full datasets are provided in Appendix F.

### 6 Related Works

Table 9: Performance comparison in the few-shot setting with 5% training data. The full results are provided in Appendix F. "LTSM-Bundle" denotes the best combination of LTSM training components

		LTSM-	Bundle	TIME	C-LLM	LLN	14TS	GP	Γ4TS	DLi	near	Patcl	hTST	Time	esNet	FEDf	ormer
Me	tric	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTh1	96 192 336 720	0.307 0.329 0.346 0.370 0.338	0.377 0.391 0.405 0.441 0.403	0.483 0.629 0.768 - 0.627	0.464 0.540 0.626 - 0.543	0.509 0.717 0.728 - 0.651	0.484 0.581 0.589 - 0.551	0.543 0.748 0.754 - 0.681	0.506 0.580 0.595 - 0.560	0.547 0.720 0.984 - 0.750	0.503 0.604 0.727 - 0.611	0.557 0.711 0.816 - 0.694	0.519 0.570 0.619 - 0.569	0.892 0.940 0.945 - 0.925	0.625 0.665 0.653 - 0.647	0.593 0.652 0.731 - 0.658	0.529 0.563 0.594 - 0.562
ETTh2	96 192 336 720 Avg	0.335 0.235 0.283 0.320 0.378 0.303	0.326 0.365 0.401 0.456 0.387	0.336 0.406 0.405 - 0.382	0.397 0.425 0.432 - 0.418	0.314 0.365 0.398 - 0.359	0.375 0.408 0.432 - 0.405	0.376 0.418 0.408 - 0.400	0.421 0.441 0.439 - 0.433	0.730 0.442 0.617 1.424 - 0.694	0.456 0.542 0.849 - 0.577	0.401 0.452 0.464 - 0.439	0.421 0.455 0.469 - 0.448	0.409 0.483 0.499 - 0.439	0.420 0.464 0.479 - 0.448	0.390 0.457 0.477 - 0.463	0.424 0.465 0.483 - 0.454
ETTm1	96 192 336 720 Avg	0.285 0.319 0.378 0.464 0.362	0.369 0.393 0.425 0.477 0.416	0.316 0.450 0.450 0.483 0.425	0.377 $0.464$ $0.424$ $0.471$ $0.434$	$ \begin{vmatrix} 0.349 \\ 0.374 \\ 0.411 \\ 0.516 \\ 0.412 \end{vmatrix} $	0.379 0.394 0.417 0.479 0.417	0.386 0.440 0.485 0.577 0.472	0.405 0.438 0.459 0.499 0.450	$\begin{array}{c} 0.332 \\ 0.358 \\ 0.402 \\ 0.511 \\ 0.400 \end{array}$	0.374 0.390 0.416 0.489 0.417	0.399 0.441 0.499 0.767 0.526	0.414 0.436 0.467 0.587 0.476	0.606 0.681 0.786 0.796 0.717	0.518 0.539 0.597 0.593 0.561	0.628 0.666 0.807 0.822 0.730	0.544 0.566 0.628 0.633 0.592
ETTm2	96 192 336 720 Avg	0.156 0.203 0.255 0.342 0.239	0.266 0.307 0.349 0.417 0.335	0.174 0.215 0.273 0.433 0.274	0.261 0.287 0.330 0.412 0.323	0.192 0.249 0.301 0.402 0.286	0.273 0.309 0.342 0.405 0.332	0.199 0.256 0.318 0.460 0.308	0.280 0.316 0.353 0.436 0.346	0.236 0.306 0.380 0.674 0.399	0.326 0.373 0.423 0.583 0.426	0.206 0.264 0.334 0.454 0.314	0.288 0.324 0.367 0.432 0.352	0.220 0.311 0.338 0.509 0.344	0.299 0.361 0.366 0.465 0.372	0.229 0.394 0.378 0.523 0.381	0.320 0.361 0.427 0.510 0.404
Weather	96 192 336 720 Avg	0.172 0.218 0.276 0.339 0.251	0.242 0.278 0.329 0.373 0.305	0.172 0.224 0.282 0.366 0.260	0.263 0.271 0.321 0.381 0.309	0.173 0.218 0.276 0.355 0.251	0.227 0.265 0.310 0.366 0.292	0.175 0.227 0.286 0.366 0.263	0.230 0.276 0.322 0.379 0.301	0.184 0.228 0.279 0.364 0.263	0.242 0.283 0.322 0.388 0.308	0.171 0.230 0.294 0.384 0.269	0.224 0.277 0.326 0.387 0.303	0.207 0.272 0.313 0.400 0.298	0.253 0.307 0.328 0.385 0.318	0.229 0.265 0.353 0.391 0.309	0.309 0.317 0.392 0.394 0.353
Electricity	96 192 336 720 Avg	0.145 0.159 0.180 0.215 0.175	0.247 0.259 0.284 0.317 0.276	0.147 0.158 0.178 0.224 0.179	0.242 0.241 0.277 0.312 0.268	0.139 0.155 0.174 0.222 0.173	0.235 0.249 0.269 0.310 0.266	0.143 0.159 0.179 0.233 0.178	0.241 0.255 0.274 0.323 0.273	0.150 0.163 0.175 0.219 0.176	0.251 0.263 0.278 0.311 0.275	0.145 0.163 0.183 0.233 0.181	0.244 0.260 0.281 0.323 0.277	0.315 0.318 0.340 0.635 0.402	0.389 0.396 0.415 0.613 0.453	0.235 0.247 0.267 0.318 0.266	0.322 0.341 0.356 0.394 0.353
	96 192 336 720 Avg	0.305 0.313 0.326 0.346 0.323	0.279 0.274 0.287 0.301 0.285	0.414 0.419 0.437 - 0.423	0.291 0.291 0.314 - 0.298	0.401 0.418 0.436 - 0.418	0.285 0.293 0.308 - 0.295	0.419 0.434 0.449 - 0.434	0.298 0.305 0.313 - 0.305	0.427 0.447 0.478 - 0.450	0.304 0.315 0.333 - 0.317	0.404 0.412 0.439 - 0.418	0.286 0.294 0.310 - 0.296	0.854 0.894 0.853 - 0.867	0.492 0.517 0.471 - 0.493	0.670 0.653 0.707 - 0.676	0.421 0.405 0.445 - 0.423
$1^{st}$ (	Count	4	7		6	1	16		0		2	:	2		0	(	0

In this work, we focus on benchmarking the training paradigms of LTSMs on top of decoder-only single models. To provide a comprehensive comparison, we categorize existing LLM-based time series forecasting approaches into three representative groups: 1. Pretrained LLM Adaptation. These methods leverage general-purpose large language models (LLMs), such as GPT, LLaMA, and Phi, for time series forecasting through fine-tuning, in-context learning, or lightweight adapter modules. Representative works include GPT4TS (Zhou et al., 2023), which adapts the Frozen Pretrained Transformer (FPT) to predict future sequences, and LLM4TS (Chang et al., 2023), which extends general-purpose LLMs for temporal modeling. 2. Time Series-Specific LLMs. These models are explicitly designed and trained on large-scale temporal datasets to better capture temporal dependencies and improve forecasting robustness. Examples include Time-LLM (Jin et al., 2023), which treats time series data as sequential events for enhanced modeling; TEST (Sun et al., 2023), which introduces transformer enhancements for complex temporal dependencies; and other recent works such as Chronos and Moirai, which aim to establish foundation-style LLMs for time series. 3. Hybrid Architectures. Hybrid approaches combine the strengths of LLMs with specialized temporal modeling components to enhance efficiency, interpretability, and robustness. For instance, DLinear (Zeng et al., 2023) integrates linear trend modeling with lightweight forecasting modules, while TimesNet (Wu et al., 2022) captures multiscale temporal patterns using neural operators. Similarly, LightTS (Campos et al., 2023) emphasizes efficient tokenization and fast inference for real-time applications. In addition, we benchmark widely adopted transformer-based models, including PatchTST (Nie et al., 2022), ETSformer (Woo et al., 2022), Non-Stationary Transformer (Liu et al., 2022), FEDformer (Zhou et al., 2022),

Autoformer (Chen et al., 2021), Informer (Zhou et al., 2021), and Reformer (Kitaev et al., 2020), which serve as essential baselines in our evaluation.

To the best of our knowledge, no prior art has provided a comprehensive benchmark to analyze the effectiveness of each component in the training of LTSMs. Some works (Li et al., 2024; Liu et al., 2025) maintain a fair platform to compare different time series forecasting methods. The others (Fons et al., 2024; Qiu et al., 2024) analyze the forecasting performance from the perspectives of time series patterns. This benchmark provides an accessible and modular pipeline for evaluating a diverse set of training components in LTSM development, leveraging a time series database and user-friendly visualization. With the debates on whether LLMs can benefit from time series forecasting tasks, our toolbox offers a scikit-learn-like API interface to efficiently explore each component's effectiveness in training LTSMs.

# 7 Conclusion and Future Perspectives

In this study, we present the first comprehensive toolbox and benchmark for understanding different design choices in training LLMs for time series forecasting. Our benchmark covers various aspects, including data preprocessing, model configuration, and dataset configuration. We delve into detailed design choices such as prompting, tokenization, training paradigms, base model selection, data quantity, and dataset diversity. Through this analysis, we derive 9 observations and identify the best combination for training LTSMs. We demonstrate that this combination achieves strong zero-shot and few-shot performance compared to state-of-the-art LTSMs, and it requires only 5% of the data to achieve comparable performance to state-of-the-art baselines on benchmark datasets. We hope that our findings will inspire future research in this direction, and this combination based on LTSM-bundle could serve as a simple yet strong baseline for future comparison.

An ongoing debate exists about whether using pre-trained weights from large language models (LLMs) can enhance time series forecasting performance (Tan et al., 2024). In response, the LTSM-Bundle package provides a robust platform for further investigation. Our current findings indicate that incorporating pre-trained weights can indeed improve forecasting performance. However, as LTSM training techniques evolve, future enhancements may offer additional insights and even different perspectives on the benefits of these pre-trained weights. Notably, using pre-trained weights as a starting point may reduce training time and help the models converge faster. Drawing from our analysis, we suggest two potential directions for enhancing the LTSM training components, as outlined below:

Advancing Prompting Strategies. The heterogeneity of time series datasets presents a significant challenge in training a universal model that can effectively fit all datasets while generalizing to unseen ones. Our analysis highlights the potential of prompting as a solution by enriching datasets with additional contextual information. Specifically, we demonstrate the effectiveness of the time series prompt, which extracts statistical insights to improve model performance. Looking ahead, we anticipate the development of more sophisticated prompting strategies to further enhance generalization. For instance, incorporating variate-specific prompts in multivariate time series data could provide richer context and improve predictive accuracy. We believe this direction holds substantial promise for future research advancements.

Constructing Synthetic Training Data. Our analysis highlights the significance of dataset diversity in training transferable LTSMs. Specifically, increasing the number of datasets can enhance performance significantly (observation ③). This insight suggests that LTSMs could achieve better transferability when exposed to more patterns during training. Thus, there is a potential for enhancing LTSMs through synthetic datasets that simulate various patterns. However, this requires more research endeavor, as synthetic data may introduce artifacts into the training dataset. This differs from previous pre-trained LTSM work like Chronos (Ansari et al., 2024), which demonstrates strong performance due to the extremely large and synthetic dataset they prepared. One of our future research directions is to augment the training data with more precise but feasible amounts of synthetic data for training LTSM for time series forecasting tasks.

# References

- Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen, Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, et al. Chronos: Learning the language of time series. arXiv preprint arXiv:2403.07815, 2024.
- Adebiyi A Ariyo, Adewumi O Adewumi, and Charles K Ayo. Stock price prediction using the arima model. In 2014 UKSim-AMSS 16th international conference on computer modelling and simulation, pp. 106–112. IEEE, 2014.
- Bishnu S Atal. Effectiveness of linear prediction characteristics of the speech wave for automatic speaker identification and verification. the Journal of the Acoustical Society of America, 55(6):1304–1312, 1974.
- Marília Barandas, Duarte Folgado, Letícia Fernandes, Sara Santos, Mariana Abreu, Patrícia Bota, Hui Liu, Tanja Schultz, and Hugo Gamboa. Tsfel: Time series feature extraction library. *SoftwareX*, 11:100456, 2020.
- David Campos, Miao Zhang, Bin Yang, Tung Kieu, Chenjuan Guo, and Christian S Jensen. Lightts: Lightweight time series classification with adaptive ensemble distillation. *Proceedings of the ACM on Management of Data*, 1(2):1–27, 2023.
- Ching Chang, Wen-Chih Peng, and Tien-Fu Chen. Llm4ts: Two-stage fine-tuning for time-series forecasting with pre-trained llms. arXiv preprint arXiv:2308.08469, 2023.
- Minghao Chen, Houwen Peng, Jianlong Fu, and Haibin Ling. Autoformer: Searching transformers for visual recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 12270–12280, 2021.
- Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. A decoder-only foundation model for time-series forecasting. arXiv preprint arXiv:2310.10688, 2023.
- Steven Davis and Paul Mermelstein. Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences. *IEEE transactions on acoustics, speech, and signal processing*, 28 (4):357–366, 1980.
- Samuel Dooley, Gurnoor Singh Khurana, Chirag Mohapatra, Siddartha V Naidu, and Colin White. Forecastpfn: Synthetically-trained zero-shot forecasting. *Advances in Neural Information Processing Systems*, 36, 2024.
- Steven Elsworth and Stefan Güttel. Time series forecasting using lstm networks: A symbolic approach. arXiv preprint arXiv:2003.05672, 2020.
- Letícia Fernandes, Marília Barandas, and Hugo Gamboa. Learning human behaviour patterns by trajectory and activity recognition. In *BIOSIGNALS*, pp. 220–227, 2020.
- Carina Figueira, Ricardo Matias, and Hugo Gamboa. Body location independent activity monitoring. In *International Conference on Bio-inspired Systems and Signal Processing*, volume 5, pp. 190–197. SciTePress, 2016.
- Elizabeth Fons, Rachneet Kaur, Soham Palande, Zhen Zeng, Tucker Balch, Manuela Veloso, and Svitlana Vyetrenko. Evaluating large language models on time series feature understanding: A comprehensive taxonomy and benchmark. arXiv preprint arXiv:2404.16563, 2024.
- Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pp. 1189–1232, 2001.
- Azul Garza and Max Mergenthaler-Canseco. Timegpt-1. arXiv preprint arXiv:2310.03589, 2023.
- Nate Gruver, Marc Finzi, Shikai Qiu, and Andrew G Wilson. Large language models are zero-shot time series forecasters. Advances in Neural Information Processing Systems, 36, 2024.

- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. arXiv preprint arXiv:2106.09685, 2021.
- Mojan Javaheripi and Sébastien Bubeck. Phi-2: The surprising power of small language models, December 2023.
- Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, et al. Time-llm: Time series forecasting by reprogramming large language models. arXiv preprint arXiv:2310.01728, 2023.
- Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. arXiv preprint arXiv:2001.04451, 2020.
- Çağri Kocaman and Muammer Özdemir. Comparison of statistical methods and wavelet energy coefficients for determining two common pq disturbances: Sag and swell. In 2009 International Conference on Electrical and Electronics Engineering-ELECO 2009, pp. I–80. IEEE, 2009.
- Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. Modeling long-and short-term temporal patterns with deep neural networks. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pp. 95–104, 2018.
- 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.
- Haoxin Liu, Shangqing Xu, Zhiyuan Zhao, Lingkai Kong, Harshavardhan Prabhakar Kamarthi, Aditya Sasanur, Megha Sharma, Jiaming Cui, Qingsong Wen, Chao Zhang, et al. Time-mmd: Multi-domain multimodal dataset for time series analysis. *Advances in Neural Information Processing Systems*, 37: 77888–77933, 2025.
- 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, 2022.
- Ioannis E Livieris, Emmanuel Pintelas, and Panagiotis Pintelas. A cnn-lstm model for gold price time-series forecasting. *Neural computing and applications*, 32:17351–17360, 2020.
- 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.
- Travis E Oliphant et al. Guide to numpy, volume 1. Trelgol Publishing USA, 2006.
- YN Pan, Jin Chen, and XL Li. Spectral entropy: A complementary index for rolling element bearing performance degradation assessment. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 223(5):1223–1231, 2009.
- Geoffroy Peeters, Bruno L Giordano, Patrick Susini, Nicolas Misdariis, and Stephen McAdams. The timbre toolbox: Extracting audio descriptors from musical signals. The Journal of the Acoustical Society of America, 130(5):2902–2916, 2011.
- Xiangfei Qiu, Xiuwen Li, Ruiyang Pang, Zhicheng Pan, Xingjian Wu, Liu Yang, Jilin Hu, Yang Shu, Xuesong Lu, Chengcheng Yang, et al. Easytime: Time series forecasting made easy. arXiv preprint arXiv:2412.17603, 2024.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Sebastian Raschka. Mlxtend: Providing machine learning and data science utilities and extensions to python's scientific computing stack. *Journal of open source software*, 3(24):638, 2018.

- Kashif Rasul, Arjun Ashok, Andrew Robert Williams, Hena Ghonia, Rishika Bhagwatkar, Arian Khorasani, Mohammad Javad Darvishi Bayazi, George Adamopoulos, Roland Riachi, Nadhir Hassen, et al. Lag-llama: Towards foundation models for probabilistic time series forecasting. *Preprint*, 2024.
- Chenxi Sun, Yaliang Li, Hongyan Li, and Shenda Hong. Test: Text prototype aligned embedding to activate llm's ability for time series. arXiv preprint arXiv:2308.08241, 2023.
- Mingtian Tan, Mike A Merrill, Vinayak Gupta, Tim Althoff, and Thomas Hartvigsen. Are language models actually useful for time series forecasting? arXiv preprint arXiv:2406.16964, 2024.
- Henriksson U. Power spectrum and bandwidth. ., 2003. URL https://www.commsys.isy.liu.se/TSDT45/Material/UlfsSpectrum2003.pdf.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Peter Welch. The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on audio and electroacoustics*, 15 (2):70–73, 1967.
- Gerald Woo, Chenghao Liu, Doyen Sahoo, Akshat Kumar, and Steven Hoi. Etsformer: Exponential smoothing transformers for time-series forecasting. arXiv preprint arXiv:2202.01381, 2022.
- Gerald Woo, Chenghao Liu, Akshat Kumar, Caiming Xiong, Silvio Savarese, and Doyen Sahoo. Unified training of universal time series forecasting transformers. arXiv preprint arXiv:2402.02592, 2024.
- 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. In *The eleventh international conference on learning representations*, 2022.
- BF Yan, Ayaho Miyamoto, and Eugen Brühwiler. Wavelet transform-based modal parameter identification considering uncertainty. *Journal of Sound and Vibration*, 291(1-2):285–301, 2006.
- 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.
- 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, 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, 2022.
- Tian Zhou, Peisong Niu, Liang Sun, Rong Jin, et al. One fits all: Power general time series analysis by pretrained lm. Advances in neural information processing systems, 36:43322–43355, 2023.
- Daniel Zwillinger and Stephen Kokoska. CRC standard probability and statistics tables and formulae. Crc Press, 1999.

# **Appendix**

### A Details of Datasets

In this paper, the training datasets include ETT (Electricity Transformer Temperature) (Zhou et al., 2021)<sup>4</sup>, Traffic<sup>5</sup>, Electricity<sup>6</sup>, Weather<sup>7</sup>, and Exchange-Rate(Lai et al., 2018). ETT<sup>8</sup> (Zhou et al., 2021) comprises four subsets: two with hourly-level data (ETTh) and two with 15-minute-level data (ETTm). Each subset includes seven features related to oil and load metrics of electricity transformers, covering the period from July 2016 to July 2018. The traffic dataset includes hourly road occupancy rates from sensors on San Francisco freeways, covering the period from 2015 to 2016. The electricity dataset contains hourly electricity consumption data for 321 clients, spanning from 2012 to 2014. The weather data set comprises 21 weather indicators, such as air temperature and humidity, recorded every 10 minutes throughout 2020 in Germany. Exchange-Rate(Lai et al., 2018) contains daily exchange rates for eight countries, spanning from 1990 to 2016. We first train our framework on the diverse time series data collection and then assess the abilities of LTSM-Bundle on jointly learning and zero-shot transfer learning to different domains of time series knowledge.

# **B** Hyper-parameter Settings of Experiments

The hyper-parameter settings of LTSM-Bundle training for all experiments are shown in Table 10. Other training hyper-parameters follow the default values in the TrainingArguments class<sup>9</sup> of the huggingface transformers package.

Hyperparameter name	Value
Number of Transformer layers $N$	3
Training / evaluation / testing split	0.7 / 0.1 / 0.2
Gradient accumulation steps	64
Learning rate	0.001
Optimizer	Adam
LR scheduler	CosineAnnealingLR
Number of epochs	10
Number of time steps per token	16
Stride of time steps per token	8
Dimensions of TS prompt	133
Transformer architectures	GPT-2-{small, medium, large}, Phi-2
Length of prediction	96, 192, 336, 720
Length of input TS data	336
Data type	torch.bfloat16
Downsampling rate of training data	20

Table 10: Hyperparameter settings of LTSM-Bundle training

### **C** Computation Infrastructure

All experiments described in this paper are conducted using a well-defined physical computing infrastructure, the specifics of which are outlined in Table 11. This infrastructure is essential for ensuring the reproducibility and reliability of our results, as it details the exact hardware environments used during the testing phases.

 $<sup>^4</sup>$ https://github.com/zhouhaoyi/ETDataset

<sup>5</sup>http://pems.dot.ca.gov

 $<sup>^6 \</sup>texttt{https://archive.ics.uci.edu/dataset/321/electricityloaddiagrams20112014}$ 

<sup>&</sup>lt;sup>7</sup>https://www.bgc-jena.mpg.de/wetter/

<sup>8</sup>https://github.com/laiguokun/multivariate-time-series-data

 $<sup>^9\</sup>mathrm{https://github.com/huggingface/transformers/blob/main/src/transformers/training\_args.py}$ 

Table 11:	Computing	infrastructure	for the	experiments

Device attribute	Value
Computing infrastructure	GPU
GPU model	Nvidia A5000 / Nvidia A100
GPU number	$8 \times A5000 / 4 \times A100$
GPU memory	$8 \times 24 \text{GB} / 4 \times 80 \text{GB}$

# **D** Global Features for Prompts

Time series prompts are developed to encapsulate the comprehensive characteristics of time series data. Specifically, let s be a certain variate of the time-series data z. We identified in Table 12 the partial global features that we leverage to craft these prompts, each selected for its ability to convey critical information about the data's temporal structure and variability. For the inter-quartile and histogram in Table 12,  $Q_3(s)$  and  $Q_1(s)$  represent the first and third quartile of the Time series data, respectively; and  $m_i$  represents the histogram in which n is the total number of observations and k the total number of bins.

Beyond those shown in Table 12, we also consider the global features according to the following references: Fast Fourier Transform, Wavelet transform, Zero crossing rate, Maximum peaks, Minimum peaks, ECDF percentile count, Slope, ECDF slope, Spectral distance, Fundamental frequency, Maximum frequency, Median frequency, Spectral maximum peaks (Barandas et al., 2020); Maximum Power Spectrum (Welch, 1967), Spectral Centroid (Peeters et al., 2011), Decrease (Peeters et al., 2011), Kurtosis (Peeters et al., 2011), Skewness (Peeters et al., 2011), Spread (Peeters et al., 2011), Slope (Peeters et al., 2011), Variation (Peeters et al., 2011), Spectral Roll-off (Figueira et al., 2016), Roll-on (Figueira et al., 2016), Human Range Energy (Fernandes et al., 2020), MFCC (Davis & Mermelstein, 1980), LPCC (Atal, 1974), Power Bandwidth (U., 2003), Spectral Entropy (Pan et al., 2009), Wavelet Entropy (Yan et al., 2006) and Wavelet Energy (Kocaman & Özdemir, 2009), Kurtosis (Zwillinger & Kokoska, 1999), Skewness (Zwillinger & Kokoska, 1999), Maximum (Oliphant et al., 2006), Minimum (Oliphant et al., 2006), Median (Oliphant et al., 2006) and ECDF (Raschka, 2018), ECDF Percentile (Raschka, 2018). For the implementation, we leverage the TSFEL library<sup>10</sup> (Barandas et al., 2020) to estimate the global features. The features are extracted separately for each variate in the time-series data.

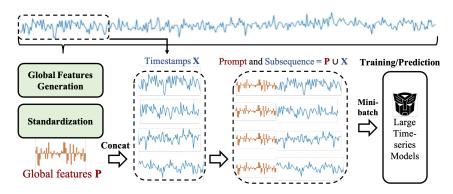


Figure 3: Time series prompt generation process based on the given training time series data.

 $<sup>^{10} \</sup>mathtt{https://tsfel.readthedocs.io/en/latest/descriptions/feature\_list.html}$ 

Table 12: Partial global features in time series prompts

Feature	Formula	Feature	Formula
Autocorrelation	$\sum_{i\in\mathbb{Z}} s_i s_{i-l}$	Centroid	$ \frac{\sum_{i=0}^{T} t_i \cdot s_i^2 / \sum_{i=0}^{T} s_i^2}{} $
Max differences	$\max_{i}(s_{i+1}-s_i)$	Mean differences	$  \operatorname{mean}_i(s_{i+1} - s_i)$
Median differences	$median_i(s_{i+1}-s_i)$	Max absolute differences	$\left  \max_{i} \left  s_{i+1} - s_i \right  \right $
Mean absolute differences	$  \operatorname{mean}_i   s_{i+1} - s_i  $	Median absolute differences	$  \operatorname{median}_i   s_{i+1} - s_i  $
Distance	$ \sum_{i=0}^{T-1} \sqrt{1 + (s_{i+1} - s_i)^2} $	Sum of absolute differences	$ \sum_{i=0}^{T-1}  s_{i+1} - s_i  $
Total energy	$\sum_{i=0}^T s_i^2 \cdot (t_T - t_0)$	Entropy	$-\sum_{x \in s} P(x) \log_2 P(x)$
Peak-to-peak distance	$ \max(s) - \min(s) $	Area under curve	
Absolute energy	$\sum_{i=0}^{T} s_i^2$	Histogram	
Inter-quartile range	$Q_3(\boldsymbol{s}) - Q_1(\boldsymbol{s})$	Mean absolute deviation	$\frac{1}{T}\sum_{i=1}^{T} s_i^2-\operatorname{mean}(s) $
Median absolute deviation	$   \operatorname{median}_i( s_i - \operatorname{median}(\boldsymbol{s}) )   $	Root mean square	$\sqrt{\frac{1}{T}\sum_{i=1}^{T}s_i^2}$
Standard deviation (STD)	$\sqrt{\frac{1}{T}\sum_{i=1}^{T}(s_i - \text{mean}(\boldsymbol{s}))^2}$	Variance (VAR)	$\frac{1}{T}\sum_{i=1}^{T}(s_i - \operatorname{mean}(s))^2$
Wavelet absolute mean	$ \operatorname{mean}(\operatorname{wavelet}(s)) $	Wavelet standard deviation	$\big    \mathrm{std}(\mathrm{wavelet}(\boldsymbol{s}))  $
Wavelet variance	$\big  \operatorname{var}(\operatorname{wavelet}(\boldsymbol{s}))  $	Skewness	$\frac{1}{T(STD)^3} \sum_{i=0}^{T} (s_i - \text{mean}(\boldsymbol{s}))^3$

Table 13: Feature comparison with other LTSM open source packages

	LTSM-Bundle	OpenLTM	Time-LLM	LLM-Time
Support for multiple model architectures and prompting strategies	Yes	Yes	No	No
Integration with database	Yes	No	No	No
Data preprocessing and pipeline integration	Yes	No	No	No
Zero-shot	Yes	Yes	Yes	Yes
Visualization	Yes	No	No	Yes

# E Comparison with Other Packages

In this section, we highlight the difference and advantages of LTSM-Bundle comparing to other existing open source LTSM packages, including OpenLTM<sup>11</sup>, Time-LLM <sup>12</sup>, and LLM-Time<sup>13</sup>. Our package involve more industrial-oriented and user-friendly features, such as database integration and report visualization.

# F Additional Experimental Results on LTSM-Bundle

In this section, we show additional results regarding comparing LTSM-Bundle with other baselines in Tables 16 and 17, results of zero-shot transfer learning in Table 14, results of different training paradigms in Table 18, results of different backbones in Table 19, results of different downsampling ratios in Table 20.

 $<sup>^{11} {\</sup>tt https://github.com/thuml/OpenLTM}$ 

 $<sup>^{12}\</sup>mathrm{https://github.com/KimMeen/Time-LLM}$ 

 $<sup>^{13} \</sup>mathtt{https://github.com/ngruver/llmtime}$ 

### F.1 Performance Comparison with Additional Baselines

Extending the analysis presented in Section 5.2, we introduce full performance comparison with new baselines. We evaluate the proposed LTSM-Bundle in zero-shot and few-shot settings to highlight its efficacy and robustness in Table 16 and 17.

# F.2 Zero-shot Transfer Learning Comparisons

In addition to the results in Section 5.2, this section introduces the full zero-shot transfer learning comparisons. We evaluate the proposed LTSM-Bundle in the zero-shot transfer scenarios, detailed in shown in Table 14.

### F.3 Training Paradigm Comparisons

Expanding upon the results in Section 4.3, this section presents the full experimental results for the training paradigms analysis, including different backbones and prompting strategies. The analytic results are detailed in Table 18.

### F.4 Backbone Architecture Comparisons

We provide all the numbers of analytics on different backbone architectures, continuing from Section 4.4. Results are in different language model backbones, including GPT-2-Small, GPT-2-Medium, GPT-2-Large, and Phi-2, shown in Table 19.

### F.5 Down-sampling Ratio Comparisons

We here present the full version of our experimental results on the different down-sampling ratios in Section 4.6. We test LTSM-Bundle with GPT-Medium as backbones with the proposed TS prompt under a fully tuning paradigm. The results in the ratio of  $\{40, 20, 10\}$  (i.e., downsample rate in  $\{2.5\%, 5\%, 10\%\}$ ) are all demonstrated in Table 20.

#### F.6 Different Numbers of Layer Adaptation Comparisons

We compared the average performance among all datasets of a 3-layer model and a full 24-layer model using GPT-medium as the backbone. Our results (in Table 15) show that the 24-layer model performs worse when trained with the same number of iterations. We believe this suggests that the 3-layer configuration is a reasonable strategy for benchmarking at this stage.

Table 14: Results of zero-shot transfer learning. A time-series model is trained on a source dataset and transferred to the target dataset without adaptation.

Methods		LTSM-	Bundle	TIME	E-LLM	LLM	Time	GP	T4TS	DLi	near	Patc	hTST	Time	esNet	Autof	former
Metric		MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
-	96	0.229	0.326	0.279	0.337	0.510	0.576	0.335	0.374	0.347	0.400	0.304	0.350	0.358	0.387	0.469	0.486
	192	0.310	0.395	0.351	0.374	0.523	0.586	0.412	0.417	0.417	0.460	0.386	0.400	0.427	0.429	0.634	0.567
$ETTh1 \rightarrow ETTh2$	336	0.336	0.414	0.388	0.415	0.640	0.637	0.441	0.444	0.515	0.505	0.414	0.428	0.449	0.451	0.655	0.588
	720	0.401	0.474	0.391	0.420	2.296	1.034	0.438	0.452	0.665	0.589	0.419	0.443	0.448	0.458	0.570	0.549
	Avg	0.319	0.402	0.353	0.387	0.992	0.708	0.406	0.422	0.493	0.488	0.380	0.405	0.421	0.431	0.582	0.548
	96 192	0.197 $0.314$	0.318 $0.420$	0.189 $0.237$	$0.293 \\ 0.312$	0.646 0.934	0.563 $0.654$	0.236 0.287	0.315 $0.342$	0.255 0.338	0.357 $0.413$	0.215 $0.275$	0.304 $0.339$	0.239 $0.291$	0.313 $0.342$	0.352	0.432 $0.460$
$\mathrm{ETTh1}  ightarrow \mathrm{ETTm2}$	336	0.314	0.420	0.237	0.312 $0.365$	1.157	0.034 $0.728$	0.287	0.344	0.338	0.415 $0.465$	0.275	0.373	0.291	0.342 $0.371$	0.415	0.489
E11111 7 E111112	720	0.313	0.483	0.231	0.390	4.730	1.531	0.435	0.374	0.425	0.403	0.334	0.373	0.434	0.419	0.403	0.469
	Avg	0.312	0.406	0.273	0.340	1.867	0.869	0.325	0.363	0.415	0.452	0.314	0.360	0.327	0.361	0.457	0.483
	96	0.390	0.439	0.450	0.452	1.130	0.777	0.732	0.577	0.689	0.555	0.485	0.465	0.848	0.601	0.693	0.569
	192	0.417	0.460	0.465	0.461	1.242	0.820	0.758	0.559	0.707	0.568	0.565	0.509	0.860	0.610	0.760	0.601
$ETTh2 \rightarrow ETTh1$	336	0.462	0.501	0.501	0.482	1.382	0.864	0.759	0.578	0.710	0.577	0.581	0.515	0.867	0.626	0.781	0.619
	720	0.568	0.588	0.501	0.502	4.145	1.461	0.781	0.597	0.704	0.596	0.628	0.561	0.887	0.648	0.796	0.644
	Avg	0.459	0.497	0.479	0.474	1.961	0.981	0.757	0.578	0.703	0.574	0.565	0.513	0.865	0.621	0.757	0.608
	96	0.200	0.316	0.174	0.276	0.646	0.563	0.253	0.329	0.240	0.336	0.226	0.309	0.248	0.324	0.263	0.352
	192	0.250	0.359	0.233	0.315	0.934	0.654	0.293	0.346	0.295	0.369	0.289	0.345	0.296	0.352	0.326	0.389
$\text{ETTh2} \rightarrow \text{ETTm2}$	336	0.327	0.416	0.291	0.337	1.157	0.728	0.347	0.376	0.345	0.397	0.348	0.379	0.353	0.383	0.387	0.426
	720	0.573	0.563	0.392	0.417	4.730	1.531	0.446	0.429	0.432	0.442	0.439	0.427	0.471	0.446	0.487	0.478
	Avg	0.337	0.413	0.272	0.341	1.867	0.869	0.335	0.370	0.328	0.386	0.325	0.365	0.342	0.376	0.366	0.411
	96	0.246	0.342	0.321	0.369	0.510	0.576	0.353	0.392	0.365	0.415	0.354	0.385	0.377	0.407	0.435	0.470
$\text{ETTm1} \rightarrow \text{ETTh2}$	192 336	$0.290 \\ 0.326$	0.374 $0.406$	0.389 $0.408$	0.410 $0.433$	0.523 0.640	$0.586 \\ 0.637$	0.443	0.437 $0.461$	0.454	0.462 $0.464$	0.447 0.481	0.434 $0.463$	0.471 $0.472$	0.453 $0.484$	0.495	0.489 $0.472$
E111111 7 E11112	720	0.363	0.440	0.406	0.436	2.296	1.034	0.466	0.461	0.490	0.529	0.431	0.403	0.472	0.484	0.470	0.472
	Avg	0.306	0.391	0.381	0.412	0.992	0.708	0.433	0.439	0.464	0.475	0.439	0.438	0.457	0.454	0.470	0.479
	96	0.144	0.257	0.169	0.257	0.646	0.563	0.217	0.294	0.221	0.314	0.195	0.271	0.222	0.295	0.385	0.457
	192	0.193	0.302	0.227	0.318	0.934	0.654	0.277	0.327	0.286	0.359	0.258	0.311	0.288	0.337	0.433	0.469
$\mathrm{ETTm1}  ightarrow \mathrm{ETTm2}$	336	0.240	0.342	0.290	0.338	1.157	0.728	0.331	0.360	0.357	0.406	0.317	0.348	0.341	0.367	0.476	0.477
	720	0.292	0.379	0.375	0.367	4.730	1.531	0.429	0.413	0.476	0.476	0.416	0.404	0.436	0.418	0.582	0.535
	Avg	0.217	0.320	0.268	0.320	1.867	0.869	0.313	0.348	0.335	0.389	0.296	0.334	0.322	0.354	0.469	0.484
	96	0.257	0.346	0.298	0.356	0.510	0.576	0.360	0.401	0.333	0.391	0.327	0.367	0.360	0.401	0.353	0.393
	192	0.309	0.382	0.359	0.397	0.523	0.586	0.434	0.437	0.441	0.456	0.411	0.418	0.434	0.437	0.432	0.437
$\text{ETTm2} \rightarrow \text{ETTh2}$	336	0.341	0.413	0.367	0.412	0.640	0.637	0.460	0.459	0.505	0.503	0.439	0.447	0.460	0.459	0.452	0.459
	720	0.350	0.432	0.393	0.434	2.296	1.034	0.485	0.477	0.543	0.534	0.459	0.470	0.485	0.477	0.453	0.467
	Avg	0.314	0.393	0.354	0.400	0.992	0.708	0.435	0.443	0.455	0.471	0.409	0.425	0.435	0.443	0.423	0.439
	96	0.364	0.410	0.359	0.397	1.179	0.781	0.747	0.558	0.570	0.490	0.491	0.437	0.747	0.558	0.735	0.576
ETT ETT 1	192	0.405	0.432	0.390	0.420	1.327	0.846	0.781	0.560	0.590	0.506	0.530	0.470	0.781	0.560	0.753	0.586
$\text{ETTm2} \rightarrow \text{ETTm1}$	336	0.413 $0.432$	0.433	0.421	0.445	1.478 3.749	0.902 $1.408$	0.778	0.578	0.706	0.567	0.565	0.497	0.778	0.578	0.750	0.593 $0.609$
	720 Avg	0.432 $0.403$	$0.446 \\ 0.430$	0.487 $0.414$	0.488 $0.438$	1.933	0.984	0.769 0.769	0.573 $0.567$	0.731	0.584 $0.537$	0.686 0.568	0.565 $0.492$	0.769 $0.769$	0.573 $0.667$	0.782	0.609
	Avg	0.400	0.400	0.414	0.400	1.500	0.304	0.109	5.501	0.043	0.001	0.000	0.432	0.103	0.001	1 0.100	0.001

Table 15: Performance comparison between 3-layer and 24-layer of LTSM-Bundle.

	3-layer	24-layer
MAE MSE	0.2003 $0.2770$	0.2439 0.3162

Table 16: Performance comparison with additional baselines (Full data)

Reformer	E   MSE MAE	3   0.837   0.728 2   0.923   0.766 9   1.097   0.835 5   1.257   0.889 5   1.029   0.805	2 2.626 1.317 1 11.12 2.979 5 9.323 2.769 5 3.874 1.697 6 6.736 2.191	9 0.658 0.528 0.658 0.592 1 0.898 0.721 3 1.102 0.841 4 0.799 0.671	3 0.658 0.619 1.078 0.827 7 1.549 0.972 8 2.631 1.242 0 1.479 0.915	4         0.689         0.596           4         0.752         0.638           6.39         0.596           1         1.130         0.792           8         0.803         0.656	0.312 0.402 0.348 0.433 0.350 0.433 0.340 0.433
Informer	MSE MAE	0.865 0.713 1.008 0.792 1.107 0.809 1.181 0.865 1.040 0.795	3.755 1.525 5.602 1.931 4.721 1.835 3.647 1.625 4.431 1.729	0.672 0.571 0.795 0.669 1.212 0.871 1.166 0.823 0.961 0.734	0.365 0.453 0.533 0.563 1.363 0.887 3.379 1.338 1.410 0.810	0.300 0.384 0.598 0.544 0.578 0.523 1.059 0.741 0.634 0.548	0.274 0.368 0.296 0.386 0.300 0.394 0.373 0.439
LightTS	MSE MAE	0.424 0.432 0.475 0.462 0.518 0.488 0.547 0.533 0.491 0.479	0.397 0.437 0.520 0.504 0.626 0.559 0.863 0.672 0.602 0.543	0.374 0.400 0.400 0.407 0.438 0.438 0.527 0.502 0.435 0.437	0.209 0.308 0.311 0.382 0.442 0.466 0.675 0.587 0.409 0.436	0.182 0.242 0.227 0.287 0.282 0.334 0.352 0.386 0.261 0.312	0.207 0.307 0.213 0.316 0.230 0.333 0.265 0.360
ETSformer	MAE	0.479 0.504 0.521 0.535 0.510	0.391 0.439 0.479 0.497 0.452	0.398   0.410   0.428   0.462   0.425	0.280 0.319 0.357 0.413 0.342	0.281 0.312 0.353 0.288 0.334	0.304 0.315 0.329 0.345
l—	MAE   MSE	0.491 0.494 0.504 0.538 0.535 0.574 0.616 0.562 0.537 0.542	0.458 0.340 0.493 0.430 0.551 0.485 0.560 0.500 0.516 0.439	0.398 0.375 0.444 0.408 0.464 0.435 0.516 0.429 0.456 0.429	0.274 0.189 0.339 0.253 0.361 0.314 0.413 0.413 0.347 0.293	0.223         0.197           0.285         0.237           0.338         0.298           0.410         0.352           0.314         0.271	0.273 0.187 0.286 0.199 0.304 0.212 0.321 0.233 0.296 0.208
r   Non-Stationary	MSE	9 0.513 0.534 6 0.588 2 0.643 7 0.570	8 0.476 2 0.512 6 0.552 1 0.562 9 0.526	5 0.386 6 0.459 7 0.495 1 0.585 7 0.481	9 0.192 0 0.280 2 0.334 2 0.417 1 0.306	6 0.173 7 0.245 5 0.321 8 0.414 2 0.288	7 0.169 4 0.182 8 0.200 1 0.222 8 0.193
Autoformer	MSE MAE	0.449 0.459 0.500 0.482 0.521 0.496 0.514 0.512 0.496 0.487	0.346 0.388 0.456 0.452 0.482 0.486 0.515 0.511 0.450 0.459	0.505 0.475 0.553 0.496 0.621 0.537 0.671 0.561 0.588 0.517	0.255 0.339 0.281 0.340 0.339 0.372 0.433 0.432 0.327 0.371	0.266 0.336 0.307 0.367 0.359 0.395 0.419 0.428 0.338 0.382	0.201 0.317 0.222 0.334 0.231 0.338 0.254 0.361 0.227 0.338
FED former	MSE MAE	0.376 0.419 0.420 0.448 0.459 0.465 0.506 0.507 0.440 0.460	0.358 0.397 0.429 0.439 0.496 0.487 0.463 0.474 0.443 0.444	0.379 0.419 0.426 0.441 0.445 0.459 0.543 0.490 0.448 0.452	0.203 0.287 0.269 0.328 0.325 0.366 0.421 0.415 0.305 0.349	7.217 0.296 7.276 0.336 7.339 0.380 7.403 0.428 7.309 0.360	0.193 0.308 0.201 0.315 0.214 0.329 0.246 0.355 0.214 0.327
TimesNet   I	MAE	0.402   0.429   0.469   0.500   0.450	0.374   0 0.414   0 0.452   0 0.468   0	0.375 0.387 0.411 0.450	0.267 0.309 0.351 0.403 0.333	0.220 0.261 0.306 0.359 0.287	0.272 0.289 0.300 0.320 0.295
-	MAE   MSE	0.399     0.384       0.421     0.436       0.436     0.491       0.466     0.521       0.430     0.491	0.336 0.340 0.379 0.402 0.380 0.452 0.422 0.462 0.379 0.414	0.342 0.338 0.369 0.374 0.392 0.410 0.420 0.478 0.380 0.400	0.255 0.187 0.292 0.249 0.329 0.321 0.385 0.408 0.315 0.291	0.198   0.172 0.241   0.219 0.282   0.280 0.334   0.365 0.264   0.259	0.222 0.168 0.240 0.184 0.259 0.198 0.290 0.220 0.252 0.192
r   PatchTST	MAE   MSE	0.399   0.370 0.416   0.413 0.443   0.422 0.490   0.447 0.437   0.413	0.353   0.274 0.418   0.339 0.465   0.329 0.551   0.379 0.446   0.330	0.343   0.290 0.365   0.332 0.386   0.366 0.421   0.416 0.378   0.351	0.269 0.165 0.303 0.220 0.342 0.274 0.421 0.362 0.333 0.255	0.237   0.149 0.282   0.194 0.319   0.245 0.362   0.314 0.300   0.225	0.237   0.129 0.249   0.157 0.267   0.163 0.301   0.197 0.263   0.161
DLinear	MSE	0.375 0.3 0.405 0.4 0.439 0.4 0.472 0.4 0.422 0.4	0.289 0.3 0.383 0.4 0.448 0.4 0.605 0.7 0.431 0.4	0.299 0.3 0.335 0.3 0.369 0.3 0.425 0.4 0.357 0.3	0.167 0.2 0.224 0.3 0.281 0.3 0.397 0.4 0.267 0.3	0.176 0.2 0.220 0.2 0.265 0.3 0.333 0.3 0.248 0.3	0.140 0.2 0.153 0.2 0.169 0.2 0.203 0.3 0.166 0.2
GPT4TS	MSE MAE	0.376 0.397 0.416 0.418 0.442 0.433 0.477 0.456 0.465 0.455	0.285 0.342 0.354 0.389 0.373 0.407 0.406 0.441 0.381 0.412	0.292 0.346 0.332 0.372 0.366 0.394 0.417 0.421 0.388 0.403	0.173 0.262 0.229 0.301 0.286 0.341 0.378 0.401 0.284 0.339	0.162 0.212 0.204 0.248 0.254 0.286 0.326 0.337 0.237 0.270	0.139 0.238 0.153 0.251 0.169 0.266 0.206 0.297 0.167 0.263
LLM4TS	MAE	0.394 0.412 0.422 0.444 0.418	0.332 0.377 0.369 0.425 0.376	0.343 0.366 0.385 0.419 0.378	0.254 0.292 0.326 0.380 0.313	0.196 0.238 0.277 0.329 0.260	0.223 0.240 0.258 0.292 0.293
TEST   L	MAE   MSE	0.400   0.371 0.422   0.403 0.437   0.420 0.467   0.422 0.431   0.404	0.338   0.269 0.379   0.328 0.381   0.353 0.423   0.383 0.380   0.333	0.346   0.285 0.369   0.324 0.392   0.353 0.420   0.408 0.382   0.343	- 0.165 - 0.220 - 0.268 - 0.350 - 0.350	0.202   0.147 0.246   0.191 0.286   0.241 0.342   0.313 0.271   0.223	0.223   0.128 0.241   0.146 0.260   0.163 0.291   0.200 0.253   0.159
_	MAE   MSE	0.392         0.372           0.418         0.414           0.427         0.422           0.457         0.447           0.423         0.414	0.328 0.275 0.375 0.340 0.409 0.329 0.420 0.381 0.383 0.331	0.334         0.293           0.358         0.332           0.384         0.368           0.411         0.418           0.372         0.353	0.253 - 0.293 - 0.329 - 0.379 - 0.313 - 0.313	0.201 0.150 0.234 0.198 0.279 0.245 0.316 0.324 0.257 0.229	0.224 0.132 0.241 0.158 0.248 0.163 0.298 0.199 0.252 0.162
dle   TIME-LLM	MAE   MSE	0.377   0.362 0.391   0.398 0.405   0.430 0.441   0.442 0.403   0.408	0.326         0.268           0.365         0.329           0.401         0.368           0.456         0.372           0.387         0.334	0.369 0.272 0.393 0.310 0.425 0.352 0.477 0.383 0.416 0.329	0.266 0.161 0.307 0.219 0.349 0.271 0.417 0.352 0.335 0.251	0.242 0.147 0.278 0.189 0.329 0.262 0.373 0.304 0.305 0.225	0.247 0.131 0.259 0.152 0.284 0.160 0.317 0.192 0.276 0.158
Is   LTSM-Bundle	MSE	0.307 0.329 0.346 0.370 0.338	0.235 0.283 0.320 0.378 0.304	0.285 0.319 0.378 0.464 0.362	0.156 0.203 0.255 0.342 0.239	0.172 0.218 0.276 0.339 0.251	0.145 0.159 0.180 0.215 0.175
Methods	Metric	ЕТТЪ1 336 Avg	ETTh2	ETTm1 96 720 Avg	ETTm2 Avg Avg	Weather 720 Avg	Electricity 8 22 8 29 96

Table 17: Performance comparison with additional baselines (5% Few Shot data)

Reformer	MAE	0.795	0.857	0.835	1.518 1.473 1.427 -	0.798 0.839 0.842 0.828 0.828	1.545 1.484 1.460 1.460 1.487	0.435 0.450 0.459 0.468 0.468	0.855 0.919 0.921 0.921 0.904	0.841 0.844 0.868 - 0.851
$\mathbf{Refo}$	MSE	1.198	1.254	1.241	3.753 3.516 3.312 - 3.527	1.234 1.287 1.288 1.247 1.247	3.883 3.553 3.446 3.445 3.581	0.406 0.446 0.465 0.471 0.447	1.414 1.240 1.253 1.249 1.289	1.586 1.602 1.668 -
mer	MAE	0.812	0.811	0.817	1.508 1.933 1.520 -	0.775 0.788 0.809 0.794 0.791	1.478 1.473 1.533 1.489	0.497 0.545 0.547 0.522 0.527	0.919 0.939 0.942 0.919 0.929	0.821 0.834 0.841 - 0.832
Informer	MSE	1.225	1.202	1.225	3.837 3.975 3.956 - 3.922	1.130 1.150 1.198 1.175 1.163	3.599 3.578 3.561 3.896 3.658	0.497 0.620 0.649 0.570 0.584	1.265 1.298 1.302 1.259 1.281	1.557 1.596 1.621 - 1.591
_ st	MAE	0.91	0.87	0.903	1.006 1.348 1.451 - 1.268	0.733 0.756 0.775 0.799 0.765	0.772 0.850 0.879 0.984 0.984	0.285   0.323   0.355   0.418   0.345	0.609 0.678 0.745 0.871 0.725	0.636 0.848 0.903 -
LightTS	MSE	1.483	1.347	1.451	2.022 3.534 4.063 - 3.206	1.048 1.097 1.147 1.200 1.123	1.108 1.317 1.415 1.822 1.415	0.230 0.274 0.318 0.401 0.305	0.639 0.772 0.901 1.200 0.878	1.157 1.688 1.826 -
mer	MAE	0.832	0.832	0.839	0.619 0.697 0.727 -	0.747 0.766 0.787 0.831 0.782	0.485 0.521 0.555 0.627 0.547	0.295 0.331 0.398 0.461 0.371	0.638   0.648   0.667   0.788	0.855 0.928 0.999 -
ETSformer	MSE	1.169	1.179	1.189	0.678 0.845 0.905 -	1.031 1.087 1.138 1.245 1.125	0.404 0.479 0.552 0.701 0.534	0.218 0.294 0.359 0.461 0.333	0.697 0.718 0.758 1.028 0.800	1.643 1.856 2.080 - 1.859
onary	MAE	0.650	0.644	0.646	0.423 0.468 0.481 - 0.489	0.587 0.591 0.603 0.611 0.598	0.316 0.349 0.380 0.445 0.372	0.252 0.307 0.348 0.407 0.328	0.518   0.531   0.578   0.786   0.603	0.821 0.838 0.860 - 0.839
Non-Stationary	MSE		0.935	0.943	0.408 0.497 0.507 -	0.823 0.844 0.870 0.893 0.857	0.238 0.298 0.353 0.475 0.341	0.215 0.290 0.353 0.452 0.327	0.484 0.501 0.574 0.952	1.468 1.509 1.602 - 1.526
_	$MAE \mid M$		0.624 0	0.598 0	0.468   0 0.504   0 0.496   0 -	0.578 0 0.591 0 0.659 0 0.655 0	0.322   0 0.357   0 0.517   0 0.538   0 0.433   0	0.299 0 0.333 0 0.393 0 0.389 0	0.367   0 0.375   0 0.411   0 0.466   0	0.481 1 0.503 1 0.523 1 - 0.502 1
Autoformer	MSE N		0.761 0	0.722 0	0.428 0 0.496 0 0.486 0 -	0.726 0 0.750 0 0.851 0 0.857 0	0.232 0 0.291 0 0.478 0 0.553 0	0.227 0 0.278 0 0.351 0 0.387 0	0.297 0 0.308 0 0.354 0 0.426 0	0.795 0.837 0.867 0.833
—	$MAE \mid N$		0.594 0	0.562 0	0.424 0 0.465 0 0.483 0 -	0.544 0 0.566 0 0.628 0 0.633 0	0.320 0.361 0.3427 0.510 0.5404 0.5404	0.309 0 0.317 0 0.392 0 0.394 0 0.353 0	0.322 0 0.341 0 0.356 0 0.394 0	0.421   0 0.405   0 0.445   0 -
FEDformer	MSE	eo	0.731 (	0.658 (	0.390 ( 0.457 ( 0.477 ( -	0.628 ( 0.666 ( 0.807 ( 0.822 ( 0.730 (	0.229 (0.394 (0.378 (0.523 (0.381 (0.381	0.229 (0.265 (0.353 (0.391 (0.309 (0.	0.235 ( 0.247 ( 0.267 ( 0.318 ( 0.266 (	0.670 0.653 0.707 0.707 0.676
—	MAE   1		0.653	0.647	0.420 0.464 0.0479 0.448 0.448	0.518 0 0.539 0 0.597 0 0.593 0	0.299 0.00 0.361 0.00 0.366 0.00 0.465 0.00 0.372 0.00	0.253 (0.2307 (0.328 (0.3385 (	0.389 (0.396 (0.	0.492   0 0.517   0 0.471   0
TimesNet	MSE N	_	0.945 0	0.925 0	0.409 0 0.483 0 0.499 0 -	0.606 0 0.681 0 0.786 0 0.717 0	0.220 0 0.311 0 0.338 0 0.509 0	0.207 0 0.272 0 0.313 0 0.400 0	0.315 0 0.318 0 0.340 0 0.635 0	0.854 0 0.894 0 0.853 0 -
_ LS	$MAE \mid N$		0.619 0	0.569 0	0.421 0 0.455 0 0.469 0 -	0.414   0 0.436   0 0.467   0 0.587   0 0.476   0	0.288   0 0.324   0 0.367   0 0.432   0 0.352   0	0.224   0 0.277   0 0.326   0 0.387   0	0.244   0 0.260   0 0.281   0 0.323   0 0.277   0	0.286   0 0.294   0 0.310   0 -
PatchTST	MSE N		0.816 0	0.694 0	0.401 0 0.452 0 0.464 0 -	0.399 0 0.441 0 0.499 0 0.767 0	0.206 0 0.264 0 0.334 0 0.454 0	0.171 0 0.230 0 0.294 0 0.384 0	0.145 0 0.163 0 0.183 0 0.233 0	0.404 0 0.412 0 0.439 0 -
ar –	$MAE \mid N$		0.727 0	0   1197	0.456   0 0.542   0 0.849   0 -	0.374   0 0.390   0 0.416   0 0.489   0	0.326 0 0.373 0 0.423 0 0.583 0 0.426 0	0.242 0 0.283 0 0.322 0 0.388 0 0.308 0	0.251   0 0.263   0 0.278   0 0.311   0 0.275   0	0.304   0 0.315   0 0.333   0 -
DLinear	MSE N		0.984 (	0.750 0	0.442 0 0.617 0 1.424 0	0.332 0 0.358 0 0.402 0 0.511 0	0.236 C 0.306 C 0.380 C 0.674 C	0.184 (0.228 (0.279 (0.364 (0.263 (0.263	0.150 ( 0.163 ( 0.175 ( 0.219 ( 0.176 (	0.427 0 0.447 0 0.478 0
 SI	MAE   1	0.506 0.580 0		0.560	0.421   0 0.441   0 0.439   1 -	0.405   0.438   0.459   0.450	0.280   C 0.316   C 0.353   C 0.436   C	0.230 0.276 0.322 0.379 0.301	0.241   0 0.255   0 0.274   0 0.323   0	0.298   0.305   0.313   0.305
GPT4TS	MSE	0.543		0.681	0.376 0.418 0.408 -	0.386 0.440 0.485 0.577 0.472	0.199 0.256 0.318 0.460 0.308	0.175 0.227 0.286 0.366 0.263	0.143 0.159 0.179 0.233 0.178	0.419 0.434 0.449 -
 SI	MAE ]	0.484		0.551	0.375   0.408   0.432   0.405	0.379   0.394   0.417   0.417   0.417	0.273   0.309   0.342   0.405   0.332	0.227   0.265   0.310   0.366   0.292   0.292	0.235   0.249   0.269   0.310   0.266	0.285   0.293   0.308   0.295   0.255   0.255   0.255   0.255   0.255   0.255   0.255   0.255   0.255   0.255   0.255   0.255
LLM4TS	MSE	0.509		0.651	0.314 0.365 0.398 -	0.349 0.374 0.411 0.516 0.412	0.192 0.249 0.301 0.402 0.286	0.173 0.218 0.276 0.355 0.251	0.139 0.155 0.174 0.222 0.173	0.401 0.418 0.436 -
CLM	MAE	0.464		0.543	0.397 0.425 0.432 -	0.377 0.464 0.424 0.471 0.434	0.261   0.287   0.330   0.412   0.323	0.263   0.271   0.321   0.309	0.242   0.241   0.277   0.312   0.268	0.291 0.291 0.314 -
TIME-LLM	MSE	0.483		0.627	0.336 0.406 0.405 - 0.382	0.316 0.450 0.450 0.483 0.425	0.174 0.215 0.273 0.433 0.274	0.172 0.224 0.282 0.366 0.260	0.147 0.158 0.178 0.224 0.179	0.414 0.419 0.437 - 0.423
	MAE	0.377		_	0.326 0.365 0.401 0.456	0.369 0.393 0.425 0.477 0.416	0.266 0.307 0.349 0.417	0.242 0.278 0.329 0.373	0.247 0.259 0.284 0.317 0.276	0.279 0.274 0.287 0.301 0.285
LTSM-E	MSE	0.307	0.346	0.338	0.235 0.283 0.320 0.378 0.304	0.285 0.319 0.378 0.464 0.362	0.156 0.203 0.255 0.342 0.239	0.172 0.218 0.276 0.339 0.251	0.145 0.159 0.180 0.215 0.175	0.305 0.313 0.326 0.346 0.323
Methods   LTSM-Bundle	Metric	96	336	Avg	96 192 336 720 Avg	96 192 336 720 Avg	96 192 336 720 Avg	96 192 336 720 Avg	96 192 336 720 Avg	96 192 336 720 Avg
2	-	14	ELL		ELLP2	ETTm1	ELLIMS	Weather	Electricity	этвяТ

Table 18: Results of different backbones, training paradigms, and prompting strategies.

Datasets		ET	Th1	ETTh2		ET	Tm1	ET	Tm2	Tra	affic	Wea	ther	Excl	ange	E	CL
Metric		MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
	96	0.354	0.415	0.259	0.350	0.551	0.507	0.215	0.319	0.393	0.377	0.222	0.292	0.108	0.246	0.250	0.342
From scratch	192	0.364	0.421	0.537	0.505	0.231	0.331	0.235	0.335	0.373	0.356	0.246	0.309	0.143	0.288	0.231	0.331
+ GPT-Medium + TS prompt	336 720	0.359 0.357	0.420 $0.430$	0.321 0.372	0.402 $0.449$	0.423	0.454 $0.444$	0.267	0.360 $0.434$	0.357 0.347	0.329 $0.311$	0.283 $0.342$	0.335 $0.388$	0.207 $0.358$	0.344 $0.461$	0.217 $0.211$	0.323 $0.317$
15 prompt	Avg	0.358	0.421	0.372	0.427	0.401	0.434	0.269	0.362	0.367	0.343	0.273	0.331	0.204	0.335	0.227	0.328
	96	0.453	0.483	0.350	0.422	0.757	0.613	0.338	0.420	0.659	0.552	0.343	0.398	0.223	0.353	0.605	0.507
From scratch	192	0.422	0.470	0.348	0.423	0.708	0.601	0.326	0.415	0.509	0.475	0.323	0.388	0.212	0.352	0.352	0.403
+ GPT-Medium + Text prompt	$\frac{336}{720}$	0.481 0.437	0.502 $0.482$	0.449 0.396	0.487 $0.461$	0.938 0.634	$0.701 \\ 0.563$	0.483	$0.506 \\ 0.459$	0.562 0.536	0.496 $0.463$	0.457 $0.415$	0.474 $0.434$	0.430 $0.457$	$0.502 \\ 0.517$	0.729 0.460	0.456 $0.438$
	Avg	0.448	0.484	0.385	0.448	0.759	0.619	0.389	0.450	0.566	0.496	0.384	0.423	0.330	0.431	0.537	0.451
	96	0.323	0.392	0.243	0.341	0.394	0.437	0.185	0.301	0.334	0.321	0.200	0.279	0.098	0.236	0.197	0.291
From scratch + GPT-Small	$\frac{192}{336}$	0.332	0.399 $0.405$	0.275	$0.362 \\ 0.394$	0.369 0.352	0.426 $0.415$	0.204 0.263	0.313 $0.364$	0.333 0.324	$0.306 \\ 0.288$	0.219 0.265	0.286 $0.323$	0.127 $0.179$	0.268 $0.321$	0.195 $0.181$	0.290 $0.282$
+ TS prompt	720	0.362	0.432	0.364	0.447	0.389	0.439	0.324	0.402	0.341	0.300	0.339	0.377	0.333	0.457	0.207	0.309
	Avg	0.340	0.407	0.300	0.386	0.376	0.429	0.244	0.345	0.333	0.304	0.256	0.316	0.184	0.320	0.195	0.293
T 11 4	96	0.317	0.383	0.240	0.334	0.431	0.443	0.178	0.285	0.315	0.293	0.188	0.257	0.083	0.208	0.161	0.265
Fully tune + GPT-Small	192 336	0.355 0.357	0.413 $0.414$	0.285	$0.370 \\ 0.388$	0.479	0.482 $0.486$	0.221 0.252	0.323 $0.346$	0.352	0.336 $0.310$	0.238 $0.275$	0.304 $0.326$	0.132 0.196	0.277 $0.336$	0.213 0.203	0.311 $0.302$
+ TS prompt	720	0.363	0.434	0.361	0.442	0.479	0.483	0.345	0.420	0.350	0.313	0.345	0.384	0.426	0.496	0.224	0.326
	Avg	0.348	0.411	0.297	0.384	0.469	0.473	0.249	0.343	0.339	0.313	0.261	0.317	0.209	0.329	0.200	0.301
Fully tune	96 192	0.305 0.335	0.377 $0.397$	0.226	0.320 $0.359$	0.276	0.360 $0.389$	0.143	0.253 $0.295$	0.305 0.315	0.279 $0.283$	0.162 $0.212$	0.227 $0.275$	0.060 0.118	0.178 $0.253$	0.144 0.161	0.246 $0.261$
+ GPT-Small	336	0.348	0.406	0.278	0.399	0.314	0.369	0.191	0.293	0.313	0.285	0.212	0.273	0.118	0.233	0.101	0.201 $0.277$
+ Text prompt	720	0.371	0.454	0.364	0.446	0.404	0.452	0.352	0.405	0.344	0.305	0.332	0.366	0.379	0.470	0.208	0.309
	Avg	0.340	0.409	0.294	0.379	0.334	0.403	0.231	0.323	0.322	0.288	0.243	0.296	0.189	0.308	0.172	0.273
Fully tune	96 192	0.301 0.332	$0.372 \\ 0.397$	0.229	$0.320 \\ 0.368$	0.261 0.288	0.346 $0.370$	0.149 0.204	0.266 $0.303$	0.300 0.316	0.268 $0.282$	0.163 0.215	0.230 $0.282$	0.058 $0.133$	0.173 $0.277$	0.141 0.158	0.241 $0.258$
+ GPT-Medium	336	0.352	0.337	0.230	0.392	0.233	0.370	0.204	0.376	0.310	0.282	0.213	0.332	0.133	0.369	0.135	0.276
+ Text prompt	720	0.368	0.436	0.378	0.452	0.371	0.431	0.492	0.494	0.344	0.303	0.350	0.385	0.321	0.442	0.207	0.308
	Avg	0.338	0.404	0.303	0.383	0.316	0.390	0.285	0.360	0.322	0.287	0.252	0.307	0.184	0.315	0.170	0.271
Fully tune	96 192	0.320 0.342	0.387 $0.403$	0.242 0.270	$0.330 \\ 0.352$	0.490 0.376	0.477 $0.423$	0.191	0.290 $0.287$	0.346 0.355	0.326 $0.327$	0.212 0.236	0.270 $0.286$	0.134 $0.173$	0.269 $0.305$	0.185 0.204	0.300 $0.316$
+ GPT-Medium	336	0.348	0.409	0.284	0.367	0.530	0.501	0.253	0.335	0.379	0.345	0.298	0.331	0.311	0.421	0.224	0.334
+ Text prompt	720	0.368	0.433	0.424	0.479	0.375	0.429	0.361	0.423	OOM	OOM	OOM	OOM	0.333	0.457	0.206	0.307
	Avg	0.344	0.408	0.305	0.382	0.443	0.458	0.250	0.334	0.360	0.333	0.249	0.295	0.238	0.363	0.205	0.314
Fully tune	96 192	0.296 0.318	$0.371 \\ 0.386$	0.234 0.273	$0.328 \\ 0.355$	0.309	0.381 $0.381$	0.150 0.190	0.263 $0.293$	0.299 0.311	0.278 $0.278$	$0.175 \\ 0.212$	0.248 $0.279$	0.073 $0.129$	0.204 $0.271$	0.145 $0.164$	0.249 $0.266$
+ Phi-2	336	0.337	0.402	0.311	0.389	0.346	0.419	0.283	0.381	0.323	0.290	0.282	0.345	0.233	0.374	0.179	0.281
+ TS prompt	720	0.372	0.445 $0.401$	0.317	0.407 $0.370$	0.404	0.461 $0.411$	0.439 0.265	0.484 $0.355$	0.347	0.305 $0.288$	0.354 0.256	0.382 $0.313$	0.404 0.210	0.501 $0.337$	0.218 0.176	0.319 $0.279$
	Avg 96																0.249
Fully tune	192	0.296	0.371 $0.385$	0.234 0.269	$0.328 \\ 0.355$	0.309	0.381 $0.383$	0.150	0.263 $0.295$	0.299 0.307	0.278 $0.275$	0.175 $0.212$	0.248 $0.283$	0.073 $0.134$	0.204 $0.281$	0.145 0.161	0.249
+ Pi-2	336	0.337	0.402	0.311	0.389	0.346	0.419	0.283	0.381	0.323	0.290	0.282	0.345	0.233	0.374	0.179	0.281
+ Text prompt	720 Avg	0.356 $0.327$	$0.430 \\ 0.397$	0.359 0.293	$0.442 \\ 0.378$	0.392	0.454 $0.409$	0.383	0.451 $0.347$	0.345 0.318	$0.302 \\ 0.286$	0.345 $0.254$	0.377 $0.313$	0.561 $0.250$	$0.606 \\ 0.366$	0.212 $0.174$	0.315 $0.277$
	96	0.362	0.419	0.233	0.363	0.589	0.533	0.231	0.332	0.318	0.396	0.224	0.293	0.129	0.274	0.227	0.333
LoRA-dim-16	192	0.394	0.419	0.213	0.397	0.582	0.531	0.259	0.361	0.502	0.437	0.339	0.280	0.129	0.345	0.257	0.358
+ GPT-Medium	336	0.403	0.457	0.321	0.413	0.560	0.532	0.293	0.392	0.547	0.457	0.320	0.369	0.266	0.409	0.291	0.386
+ TS prompt	720 Avg	0.444	0.499 $0.455$	0.366	0.451 $0.406$	0.576	0.547 $0.536$	0.355	0.436 $0.380$	0.660	0.519 $0.452$	0.369 0.313	$0.406 \\ 0.337$	0.457 $0.263$	0.532 $0.390$	0.406 $0.295$	0.479 $0.389$
	96	0.365	0.422	0.270	0.361	0.596	0.593	0.222	0.329	0.438	0.408	0.223	0.294	0.117	0.259	0.233	0.341
LoRA-dim-32	192	0.401	0.449	0.314	0.398	0.594	0.537	0.261	0.363	0.503	0.443	0.223	0.340	0.204	0.346	0.259	0.361
+ GPT-Medium	336	0.403	0.457	0.321	0.413	0.563	0.533	0.294	0.393	0.547	0.459	0.321	0.370	0.267	0.410	0.294	0.390
+ TS prompt	720 Avg	0.444	0.498 $0.457$	0.367	0.452 $0.406$	0.572 0.581	0.545 $0.552$	0.357 0.283	0.437 $0.380$	0.647	0.513 $0.456$	0.369 0.298	0.406 $0.352$	0.454 $0.260$	$0.530 \\ 0.386$	0.399 0.296	0.473 $0.391$
	96	0.377	0.431	0.284	0.376	0.603	0.538	0.239	0.348	0.462	0.423	0.244	0.313	0.154	0.302	0.242	0.348
LoRA-dim-16	192	0.394	0.445	0.313	0.400	0.578	0.530	0.263	0.367	0.511	0.441	0.284	0.344	0.203	0.351	0.262	0.363
+ GPT-Medium	336	0.412	0.465	0.325	0.417	0.571	0.538	0.299	0.397	0.567	0.471	0.323	0.373	0.267	0.413	0.308	0.402
+ Word prompt	720 Avg	0.448	$0.501 \\ 0.461$	0.368	0.452 $0.411$	0.582	0.550 $0.539$	0.357	0.437 $0.387$	0.672 0.553	0.526 $0.465$	0.370 0.305	0.407 $0.359$	0.452 0.269	0.529 $0.399$	0.416 $0.307$	0.486 $0.400$
-	96	0.365	0.423	0.276	0.367	0.590	0.533	0.230	0.337	0.449	0.410	0.234	0.305	0.133	0.277	0.237	0.343
LoRA-dim-32	192	0.400	0.449	0.311	0.397	0.572	0.527	0.261	0.364	0.515	0.447	0.284	0.345	0.207	0.352	0.265	0.366
+ GPT-Medium	336	0.410	0.463	0.324	0.416	0.570	0.538	0.299	0.398	0.563	0.469	0.323	0.373	0.268	0.414	0.305	0.399
+ Word prompt	720 Avg	0.447	0.500 $0.459$	0.368	0.453 $0.408$	0.589	0.553 $0.538$	0.359 0.287	0.459 $0.389$	0.664	0.522 $0.462$	0.370 0.303	0.407 $0.357$	0.453 $0.265$	$0.530 \\ 0.393$	0.414 0.305	0.486 $0.398$
	8			1		1											

Table 19: Results of different backbones.

Datasets		ET	Th1	ET	Th2	ET	Γm1	ET	Γm2	Tra	affic	Wea	ther	Exchange		ECL	
Metric		MSE	MAE	MSE	MAE	MSE	MAE										
	96	0.297	0.368	0.224	0.319	0.277	0.360	0.147	0.259	0.304	0.280	0.168	0.240	0.069	0.192	0.148	0.251
Fully tune	192	0.328	0.391	0.279	0.362	0.300	0.377	0.207	0.311	0.315	0.279	0.212	0.277	0.121	0.264	0.159	0.258
+ GPT-Large	336	0.347	0.407	0.365	0.420	0.322	0.397	0.284	0.364	0.324	0.287	0.291	0.341	0.356	0.461	0.174	0.276
+ TS prompt	720	0.358	0.427	0.430	0.478	0.358	0.421	0.436	0.467	0.348	0.303	0.370	0.388	0.424	0.525	0.207	0.308
	Avg	0.333	0.398	0.325	0.395	0.314	0.389	0.269	0.350	0.323	0.287	0.260	0.311	0.242	0.360	0.172	0.273
	96	0.307	0.377	0.235	0.326	0.285	0.369	0.156	0.266	0.305	0.278	0.172	0.242	0.065	0.186	0.145	0.247
Fully tune	192	0.329	0.391	0.283	0.365	0.319	0.393	0.203	0.307	0.313	0.274	0.218	0.278	0.115	0.248	0.159	0.259
+ GPT-Medium	336	0.346	0.405	0.320	0.401	0.378	0.425	0.255	0.349	0.326	0.287	0.276	0.329	0.206	0.339	0.180	0.284
+ TS prompt	720	0.370	0.441	0.378	0.456	0.464	0.477	0.342	0.417	0.346	0.301	0.339	0.373	0.409	0.487	0.215	0.317
	Avg	0.338	0.403	0.304	0.387	0.362	0.416	0.239	0.335	0.323	0.285	0.251	0.305	0.199	0.315	0.175	0.276
	96	0.317	0.383	0.240	0.334	0.431	0.443	0.178	0.285	0.315	0.293	0.188	0.257	0.083	0.208	0.161	0.265
Fully tune	192	0.355	0.413	0.285	0.370	0.479	0.482	0.221	0.323	0.352	0.336	0.238	0.304	0.132	0.277	0.213	0.311
+ GPT-Small	336	0.357	0.414	0.302	0.388	0.486	0.486	0.252	0.346	0.339	0.310	0.275	0.326	0.196	0.336	0.203	0.302
+ TS prompt	720	0.363	0.434	0.361	0.442	0.479	0.483	0.345	0.420	0.350	0.313	0.345	0.384	0.426	0.496	0.224	0.326
	Avg	0.348	0.411	0.297	0.384	0.469	0.473	0.249	0.343	0.339	0.313	0.261	0.317	0.209	0.329	0.200	0.301
	96	0.296	0.371	0.234	0.328	0.309	0.381	0.150	0.263	0.299	0.278	0.175	0.248	0.073	0.204	0.145	0.249
Fully tune	192	0.318	0.386	0.273	0.355	0.301	0.381	0.190	0.293	0.311	0.278	0.212	0.279	0.129	0.271	0.164	0.266
+ Phi-2	336	0.337	0.402	0.311	0.389	0.346	0.419	0.283	0.381	0.323	0.290	0.282	0.345	0.233	0.374	0.179	0.281
+ TS prompt	720	0.372	0.445	0.317	0.407	0.404	0.461	0.439	0.484	0.347	0.305	0.354	0.382	0.404	0.501	0.218	0.319
	Avg	0.331	0.401	0.284	0.370	0.340	0.411	0.265	0.355	0.320	0.288	0.256	0.313	0.210	0.337	0.176	0.279

Table 20: Results of different down-sampling ratios. Experiments with GPT-Medium as backbones, TS prompt, and fully tuning paradigm.

Datasets	8	ET	Γh1	ET	Th2	ET	Γm1	ET	Γm2	Tra	ıffic	Wea	ther	EC	CL
Metric		MSE	MAE												
	96	0.4536	0.4787	0.3395	0.4097	0.7441	0.6068	0.3229	0.4055	0.6619	0.5636	0.3538	0.3979	0.6118	0.499
Downsample Ratio = $40$	192	0.4648	0.4935	0.386	0.4507	0.7659	0.6336	0.3797	0.4604	0.6644	0.5537	0.4106	0.454	0.6756	0.4915
	336	0.6629	0.6167	0.5982	0.5916	1.1366	0.8212	0.6904	0.6455	1.0706	0.7785	0.7359	0.65	1.6485	0.7188
	720	1.0518	0.8133	0.8802	0.7588	1.8609	1.1304	1.2011	0.9073	1.7276	1.062	1.3176	0.9268	3.9988	1.0223
	Avg	0.343	0.408	0.307	0.390	0.353	0.412	0.234	0.331	0.344	0.314	0.253	0.306	0.210	0.330
	96	0.307	0.3767	0.2349	0.3263	0.285	0.3687	0.1555	0.2657	0.3053	0.2778	0.1717	0.2416	0.1447	0.2468
Downsample	192	0.3288	0.3908	0.2826	0.3649	0.3193	0.3925	0.2032	0.3069	0.3132	0.2744	0.2177	0.2782	0.1587	0.2588
Ratio = 20	336	0.346	0.405	0.3198	0.4005	0.3782	0.4254	0.2549	0.3492	0.3263	0.2869	0.2761	0.3287	0.1803	0.2835
Matio = 20	720	0.3704	0.4405	0.378	0.456	0.4638	0.4773	0.3422	0.4172	0.3456	0.301	0.3386	0.3732	0.2151	0.3167
	Avg	0.338	0.404	0.303	0.383	0.316	0.390	0.285	0.360	0.322	0.287	0.252	0.307	0.184	0.315
	96	0.2975	0.3698	0.2268	0.3175	0.2633	0.3462	0.1463	0.2583	0.2961	0.2626	0.2583	0.2247	0.1406	0.2411
Dommoononlo	192	0.3293	0.3896	0.2848	0.3674	0.3286	0.3991	0.1995	0.3014	0.3089	2673	0.2151	0.2786	0.1568	0.2564
Downsample Ratio = 10	336	0.3461	0.4039	0.3097	0.3938	0.3593	0.4198	0.259	0.3505	0.3206	0.2785	0.2651	0.3155	0.1744	0.2747
Ratio = 10	720	0.3676	0.4333	0.4101	0.4738	0.4096	0.4505	0.3632	0.4242	0.3428	0.2983	0.3462	0.38	0.2099	0.3101
	Avg	0.335	0.402	0.321	0.392	0.341	0.399	0.235	0.332	0.312	0.274	0.252	0.308	0.232	0.360