



TFB: Towards Comprehensive and Fair Benchmarking of Time Series Forecasting Methods

Xiangfei Qiu
East China Normal
University, China

Jilin Hu
East China Normal
University, China ✉

Lekui Zhou
Huawei Cloud Algorithm
Innovation Lab, China

Xingjian Wu
East China Normal
University, China

Junyang Du
East China Normal
University, China

Buang Zhang
East China Normal
University, China

Chenjuan Guo
East China Normal
University, China

Aoying Zhou
East China Normal
University, China

Christian S. Jensen
Aalborg University,
Denmark

Zhenli Sheng
Huawei Cloud Algorithm
Innovation Lab, China

Bin Yang
East China Normal
University, China

ABSTRACT

Time series are generated in diverse domains such as economic, traffic, health, and energy, where forecasting of future values has numerous important applications. Not surprisingly, many forecasting methods are being proposed. To ensure progress, it is essential to be able to study and compare such methods empirically in a comprehensive and reliable manner. To achieve this, we propose **TFB**, an automated benchmark for Time Series Forecasting (TSF) methods. TFB advances the state-of-the-art by addressing shortcomings related to datasets, comparison methods, and evaluation pipelines: 1) insufficient coverage of data domains, 2) stereotype bias against traditional methods, and 3) inconsistent and inflexible pipelines. To achieve better domain coverage, we include datasets from 10 different domains: traffic, electricity, energy, the environment, nature, economic, stock markets, banking, health, and the web. We also provide a time series characterization to ensure that the selected datasets are comprehensive. To remove biases against some methods, we include a diverse range of methods, including statistical learning, machine learning, and deep learning methods, and we also support a variety of evaluation strategies and metrics to ensure a more comprehensive evaluation of different methods. To support the integration of different methods into the benchmark and enable fair comparisons, TFB features a flexible and scalable pipeline that eliminates biases. Next, we employ TFB to perform a thorough evaluation of 21 Univariate Time Series Forecasting (UTSF) methods on 8,068 univariate time series and 14 Multivariate Time Series Forecasting (MTSF) methods on 25 datasets. The results offer a deeper understanding of the forecasting methods, allowing us to better select the ones that are most suitable for particular datasets and settings. Overall, TFB and this evaluation provide researchers with improved means of designing new TSF methods.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 17, No. 9 ISSN 2150-8097.
doi:10.14778/3665844.3665863

PVLDB Reference Format:

Xiangfei Qiu, Jilin Hu, Lekui Zhou, Xingjian Wu, Junyang Du, Buang Zhang, Chenjuan Guo, Aoying Zhou, Christian S. Jensen, Zhenli Sheng and Bin Yang. TFB: Towards Comprehensive and Fair Benchmarking of Time Series Forecasting Methods. PVLDB, 17(9): 2363 - 2377, 2024.
doi:10.14778/3665844.3665863

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/decisionintelligence/TFB>.

1 INTRODUCTION

As part of the ongoing digitalization, time series are generated in a variety of domains, such as economic [36, 75], traffic [30, 33–35, 49, 51, 52, 62, 79, 85, 93, 95], health [44, 61, 83, 88], energy [1, 29], and AIOps [7, 8, 41, 72, 87, 103]. Time Series Forecasting (TSF) is essential in key applications in these domains [28, 67, 94, 97]. Given historical observations, it is valuable if we can know the future values ahead of time. Correspondingly, TSF has been firmly established as an active research field, witnessing the proposal of numerous methods.

Time series organize data points chronologically and are either univariate or multivariate depending on the number of variables in each data point. Accordingly, TSF methods can be classified as either Univariate Time Series Forecasting (UTSF) or Multivariate Time Series Forecasting (MTSF) methods. Among early methods, Autoregressive Integrated Moving Average (ARIMA) [4] and Vector Autoregression (VAR) [82] are arguably the most popular univariate and multivariate forecasting methods, respectively. Subsequent methods that exploit machine learning, e.g., XGBoost [11, 99] and Random Forest [5, 59] offer better performance than the early methods. Most recently, methods based on deep learning have demonstrated state-of-the-art (SOTA) forecasting performance on a variety of datasets [10, 12, 14–16, 50, 60, 64, 70, 89, 91, 92, 96, 101, 102, 104].

As more and more methods are being proposed for different datasets and settings, there is an increasing need for fair and comprehensive empirical evaluations. To achieve this, we identify and address three issues in existing evaluation frameworks, thereby advancing our evaluation capabilities.

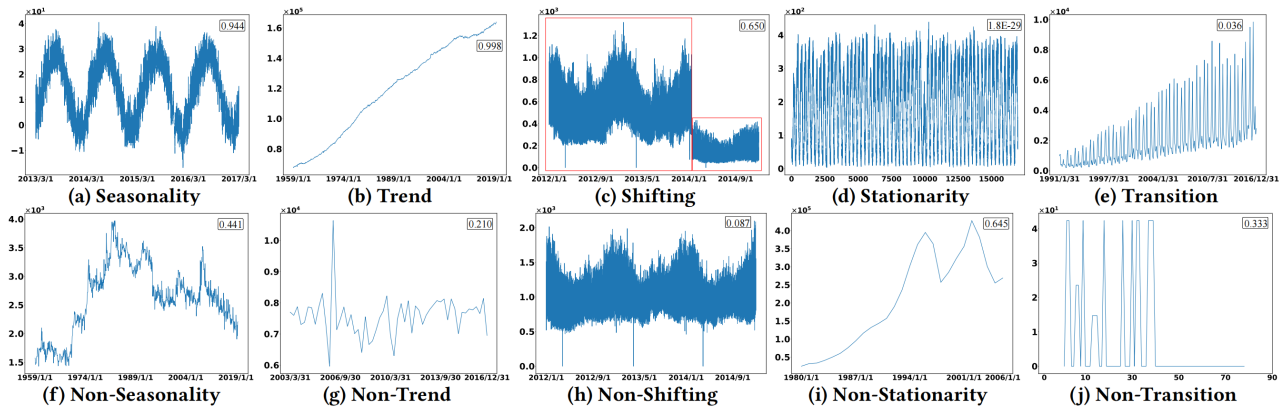


Figure 1: Visualization of data with different characteristics.




Figure 2: Statistics of data domains covered by existing multivariate time series benchmarks.




Figure 3: Box plot of the variations in normalized values of characteristics across the multivariate datasets in the TFB and TSlib.

Issue 1. Insufficient Coverage of Data Domains. Time series from different domains may exhibit diverse characteristics. Figure 1a depicts a time series from the environment domain called AQShunyi [100] that records temperature information at hourly intervals, exhibiting a distinct seasonal pattern. This pattern is reasonable in this scenario because temperatures in nature often cycle around the year. Figure 1b shows a time series from FRED-MD [58] belongs to economic domain that describes the monthly macroeconomic from 114 regional, national, and international sources with a clear increasing tendency. This may be attributed to overall

Table 1: VAR, LR versus other methods, using MAE as the evaluation metric and a forecasting horizon of 24 steps.

Datasets	VAR	LR	PatchTST	NLinear	FEDformer	Crossformer
NASDAQ	0.462	0.616	0.567	<u>0.522</u>	0.547	0.745
Wind	0.620	0.583	0.652	0.640	0.697	<u>0.590</u>
ILI	1.012	4.856	0.835	<u>0.919</u>	1.020	1.096

economic stability with minimal fluctuations, reflecting sustained growth in the macroeconomic indicators. Figure 1c depicts a series among Electricity [84] which comes from electricity domain and has a significant change in the data at a certain point in time, which might indicate an abrupt event, etc. However, these simple patterns are only the tip of the iceberg, and time series from different domains may exhibit much more complex patterns that either combine the above characteristics or are entirely different. Therefore, using only limited domains results in limited coverage of time series characteristics, which cannot offer a full picture.

However, few empirical studies and benchmarks cover a wide variety of data domains. Figure 2 summarizes the multivariate data domains used in existing forecasting benchmarks which include MTSF. We observe that TSlib [89], LTSF-Linear [98], BasicTS [48], and BasicTS+ [76] only include around 10 datasets, covering less than or equal to 5 domains. We observe that these datasets are concentrated in mainly two domains, namely traffic and electricity. Since the multivariate time series datasets in TSlib are the most used, we investigate the variations in the values of the characteristics of datasets in TSlib and TFB—see Figure 3. We observe that the TFB datasets exhibit more diverse distributions than those of TSlib across the six characteristics. *We argue that it is beneficial to broaden the coverage of domains, thereby enabling a more extensive assessment of method performance.*

Issue 2. Stereotype bias against traditional methods. It is difficult for a single method to exhibit the best performance across all datasets. Methods exhibit varying performance across different datasets. To illustrate the issue, we conduct experiments on three datasets (NASDAQ [23], Wind [46], and ILI [90]) from different domains (stock markets, energy, health) on methods VAR [82], PatchTST [64], LinearRegression (LR) [32, 40], NLinear [98], FEDformer [106], and Crossformer [101]. Results are shown in Table 1.

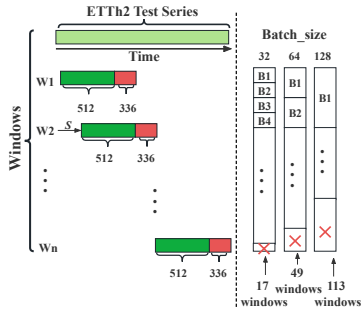


Table 2: Impact of batch sizes with “drop last.”

Size	PatchTST	DLinear	FEDformer
1	0.4203	0.4874	0.4120
32	0.4138	0.4831	0.4084
64	0.3999	0.4726	0.4022
128	0.3750	0.4539	0.3921
256	0.3561	0.4360	0.3825
512	0.3483	0.4251	0.3736

Figure 4: “Drop last” illustration.

Surprisingly, VAR outperforms all recently proposed SOTA methods on NASQAD and is better than FEDformer and Crossformer on ILL. Furthermore, LR performs better than recently proposed SOTA methods on Wind. However, the experimental studies in their original papers [64, 101, 106] do not include VAR and LR in their baselines and rather assume that traditional methods are incapable of obtaining competitive performance. From Table 3, it follows that no existing MTSF benchmark has evaluated statistical methods. Moreover, as the training mechanisms for statistical methods differ from those of deep learning-based methods, it is difficult for existing benchmarks to accommodate statistical methods. *We argue that it is beneficial to eliminate stereotype bias against traditional methods by comparing a wide range of methods.*

Issue 3. Lack of consistent and flexible pipelines. The performance of different methods varies with changing experimental settings, e.g., splits among training/validation/testing data, the choice of normalization method, and the selection of hyper-parameter settings. For example, implementations of existing methods often employ a “Drop Last” trick in the testing phase [64, 86, 106]. To accelerate the testing, it is common to split the data into batches. However, if we discard the last incomplete batch with fewer instances than the batch size, this causes unfair comparisons. For example, in Figure 4, the ETTh2 [104] has a testing series of length 2,880, and we need to predict 336 future time steps using a look-back window of size 512. If we select the batch sizes to be 32, 64, and 128, the number of samples in the last batch are 17, 49, and 113, respectively. Discarding those last-batch testing samples is inappropriate unless all methods use the same strategy. Table 2 shows the testing results of PatchTST [64], DLinear [98] and FEDformer [106] on the ETTh2 with different batch sizes and the “Drop Last” trick turned on. We observe that the performance of the methods changes when varying the batch size. In addition, the pipelines in most benchmarks are inflexible and do not support simultaneous evaluation of statistical learning, machine learning, and deep learning methods. *We argue that it is crucial to ensure a consistent and flexible pipeline so that methods are evaluated in the same setting, thereby improving the fairness of the findings.*

Robust and extensive benchmarks enable researchers to evaluate proposals for new methods more rigorously across diverse range of datasets, which is crucial for advancing the state-of-the-art [73, 81]. For example, ImageNet [20], that encompasses a substantial dataset has been instrumental in ensuring progress in computer vision. ImageNet has established itself as a standard for assessing methods

in image processing due to its support for rigorous evaluation. Table 3 compares existing benchmarks for TSF according to seven properties. No existing benchmark possesses all properties.

We propose the Time series Forecasting Benchmark (TFB) to facilitate the empirical evaluation and comparison of TSF methods more comprehensively across application domains and methods, and with improved fairness. TFB contributes a collection of challenging and realistic datasets and provides a user-friendly, flexible, and scalable evaluation pipeline that offers robust evaluation support. TFB possesses the following key characteristics:

- A comprehensive collection of datasets organized according to a taxonomy (to address **Issue 1**): Its collection of datasets offer diverse characteristics, encompassing time series from multiple domains and complex settings. This contributes to ensuring more robust and extensive evaluations.
- Broad coverage of existing methods and extended support for evaluation strategies and metrics (to address **Issue 2**): TFB covers a diverse range of methods, including statistical learning, machine learning, and deep learning methods, accompanied by a variety of evaluation strategies and metrics. This richness enables more comprehensive evaluations across methods and evaluation settings.
- Flexible and scalable pipeline (to address **Issue 3**): By its design, TFB improves the fairness of comparisons of methods. Methods are evaluated using a unified pipeline, consistent and standardized evaluation strategy and datasets are employed, eliminating biases and enabling more accurate performance comparisons. This enables more fair and meaningful conclusions regarding the effectiveness and efficiency of methods.

Based on the experiments conducted using TFB, we make the following key observations: (1) The statistical methods VAR [82] and LinearRegression (LR) [32, 40] perform better than recently proposed SOTA methods on some datasets—see Table 8. (2) Linear-based methods perform well when datasets exhibit an increasing trend or significant shifts. (3) Transformer-based methods outperform linear-based methods on datasets with marked seasonality, and nonlinear patterns, as well as more pronounced patterns or strong internal similarities. (4) Methods that consider dependencies between channels can enhance the performance of MTSF substantially, particularly on datasets with strong correlations, compared to methods that assume channel independence.

We make the following main contributions.

- We present TFB which is specifically designed to further improve fair comparisons in TSF, include UTSF and MTSF. TFB facilitates comparisons with 20+ UTSF methods on 8,068 univariate time series and 14 MTSF methods on 25 multivariate datasets.
- We identify, collect, and process previously proposed TSF datasets to determine a comprehensive dataset covering diverse domains and characteristics and organize them in a standardized format. Then, we design experiments to study the performance of different methods over different characteristics of datasets.
- TFB offers an automated end-to-end pipeline for evaluating forecasting methods, streamlining and standardizing the steps of loading time series datasets, configuring experiments, and evaluating methods. This simplifies the evaluation process for

Table 3: Time series forecasting benchmark comparison.

Benchmark \ Property	Univariate forecasting	Multivariate forecasting	Statistical method	Machine learning method	Deep learning method	Taxonomy of data	Flexible & scalable pipeline
M3 [56]	✓	×	✓	✓	×	×	×
M4 [57]	✓	×	✓	✓	✓	×	×
LTSF-Linear [98]	×	✓	×	×	✓	×	○
TSlib [89]	✓	✓	×	×	✓	×	○
BasicTS [48]	×	✓	×	✓	✓	×	○
BasicTS+ [76]	×	✓	×	×	✓	○	○
Monash [27]	✓	×	✓	✓	×	×	○
Libra [3]	✓	×	✓	✓	×	×	○
TFB (Ours)	✓	✓	✓	✓	✓	✓	✓

× indicates absent, ✓ indicates present, ○ indicates incomplete.

researchers. Additionally, all datasets and code are available at <https://github.com/decisionintelligence/TFB>.

- We use TFB to evaluate and compare a diverse set of methods, covering statistical learning, machine learning, and deep learning methods and a rich array of evaluation tasks and strategies. We summarize the evaluation results into some key findings.

The rest of the paper is structured as follows. We review related work in Section 2 and introduce terms and definitions in Section 3. In Section 4, we cover the design of TFB, and in Section 5, we use TFB to benchmark existing forecasting methods. We offer conclusions in Section 6.

2 RELATED WORK

We proceed to first cover approaches to TSF and then review existing benchmarking proposals.

Time series forecasting: Existing methods for TSF can be categorized broadly into three main categories: statistical learning, machine learning, and deep learning methods.

Early proposals primarily employed statistical learning methods. ARIMA [4], ETS [37], Theta [22], VAR [27], and Kalman Filter (KF) [31] are classical and widely utilized methods. These forecasting methods are based on the idea that future values of time series can be predicted from observed, past values. With the rapid advances in machine learning technologies, machine learning methods for TSF have emerged [24]. Notably, XGBoost [11, 99], Gradient Boosting Regression Trees (GBRT) [25], Random Forests [5, 59] and LightGBM [39] have been applied extensively to better accommodate nonlinear relationships and complex patterns. Machine learning methods are flexible at handling different types and lengths of time series and generally offer better forecasting accuracy than traditional methods.

However, these methods still require manual feature engineering and model design. Taking advantage of the representation learning capabilities offered by deep neural networks (DNNs) on rich data, numerous deep learning methods have been proposed. In many cases, these methods outperform traditional techniques in terms of forecasting accuracy. TCN [2] and DeepAR [74] treat time series as sequences of vectors and utilize CNNs or RNNs to capture temporal dependencies. Additionally, Transformer architectures, including Informer [104], FEDformer [106], Autoformer [90], Triformer [13], and PatchTST [64] can capture more complex temporal

dynamics, significantly improving forecasting performance. MLP-based models such as N-HITS [9], N-BEATS [66], NLinear [98], and DLinear [98] employ a simple architecture with relatively few parameters and have also demonstrated good forecasting accuracy. **Benchmarks:** Several benchmarks for TSF have been proposed: Libra [3], BasicTS [48], BasicTS+ [76], Monash [27], M3 [56], M4 [57], LTSF-Linear [98] and TSlib [89]. However, these benchmarks fall short in different aspects—see Table 3.

First, most benchmarks consider either univariate or multivariate time series forecasting. Early works, i.e., M3 [56], M4 [57], Libra [3] and Monash [27], only pay attention to univariate time series. And recent works, e.g., LTSF-Linear [98], BasicTS [48] and BasicTS+ [76], just consider multivariate. Only TSlib [89] considers both.

Second, a notable issue arises regarding the assessment of method diversity. In the early times, the deep learning methods are not dominant in TSF, so most early benchmarks do not have deep learning methods in their comparison, e.g., M3 [56], M4 [57] and Libra [3]. On the contrary, with more deep learning methods demonstrating their performance in TSF, recent benchmarks focus solely on deep learning methods, e.g., LTSF-Linear [98], TSlib [89], BasicTS [48] and BasicTS+ [76]. As mentioned in Section 1, this limited coverage of different method paradigms is a concern.

Moreover, there is a challenge regarding scalable and unified pipelines. In the earliest works, M3 [56] and M4 [57] do not provide a pipeline in their benchmark at all, making it hard to integrate new methods. Monash [27] and Libra [3] have pipelines but it is designed for statistical methods, which are unable to integrate deep learning methods. On the contrary, LTSF-Linear [98] and TSlib [89] have pipelines specific to deep learning methods, which are unable to evaluate statistical methods. We consider these pipelines to be incomplete, reducing scalability and hindering their broader application of these benchmarks.

Furthermore, the majority of benchmarks do not categorize the dataset; only BasicTS+ [76] provides a coarse-grained classification. We consider this taxonomy to be incomplete. In contrast, TFB involves a fine-grained classification of the dataset.

Generally, TFB aims to enable a more reliable, thorough, and user-friendly evaluation. Thus, the proposed benchmark encompasses a broader range of methods and utilizes more scalable implementation pipelines than existing benchmarks, ensuring a fair comparison of TSF methods and promoting progress in the field of TSF.

Algorithm 1 Calculating Shifting Values of Time Series

Input: Time series $X \in \mathbb{R}^{T \times 1}$ **Output:** Shifting value $\delta \in (0, 1)$ of X

- 1: Normalize X by calculating the z-score to obtain $Z \in \mathbb{R}^{T \times 1}$
 - 2: $Z_{\min} \leftarrow \min(Z)$, $Z_{\max} \leftarrow \max(Z)$
 - 3: $S = \{s_i \mid s_i \leftarrow Z_{\min} + (i - 1) \frac{Z_{\max} - Z_{\min}}{m}, 1 \leq i \leq m\}$ where m is the number of thresholds
 - 4: **for** s_i in S **do**
 - 5: $K \leftarrow \{j \mid Z_j > s_i, 1 \leq j \leq T\}$, $M_i \leftarrow \text{median}(K)$, $1 \leq i \leq m$
 - 6: **end for**
 - 7: $M' \leftarrow \text{Min-Max Normalization}(M)$
 - 8: **return** $\delta \leftarrow \text{median}(\{M'_1, M'_2, \dots, M'_m\})$
-

Algorithm 2 Calculating Transition Values of Time Series

Input: Time series $X \in \mathbb{R}^{T \times 1}$ **Output:** Transition value $\Delta \in (0, \frac{1}{3})$ of X

- 1: Calculate the first zero crossing of the autocorrelation function:
 $\tau \leftarrow \text{firstzero_ac}(X)$
 - 2: Generate $Y \in \mathbb{R}^{T' \times 1}$ by downsampling X with stride τ
 - 3: Define index $I = \text{argsort}(Y) \in \mathbb{R}^{T' \times 1}$, then characterize Y to obtain $Z \in \mathbb{R}^{T' \times 1}$:
 - 4: **for** $j \in [0 : T']$ **do**
 - 5: $Z[j] \leftarrow \text{floor}(I[j] / \frac{1}{3} T')$
 - 6: **end for**
 - 7: Generate a transition matrix $M \in \mathbb{R}^{3 \times 3}$:
 - 8: **for** $j \in [0 : T']$ **do**
 - 9: $M[Z[j] - 1][Z[j + 1] - 1]++$
 - 10: **end for**
 - 11: $M' \leftarrow \frac{1}{T'} M$
 - 12: Compute the covariance matrix C between the columns of M'
 - 13: **return** $\Delta \leftarrow \text{tr}(C)$
-

3 PRELIMINARIES

We provide definitions of time series and time series forecasting, and we cover key dataset characteristics, including trend, seasonality, stationarity, shifting, transition, and correlation.

DEFINITION 1 (TIME SERIES). A time series $X \in \mathbb{R}^{T \times N}$ is a time-oriented sequence of N -dimensional time points, where T is the number of time points, and N is the number of variables. When $N = 1$, a time series is called univariate. When $N > 1$, it is called multivariate.

DEFINITION 2 (TIME SERIES FORECASTING). Given a historical time series $X \in \mathbb{R}^{H \times N}$ of H time points, time series forecasting aims to predict the next F future time points, i.e., $Y \in \mathbb{R}^{F \times N}$, where F is called the forecasting horizon.

DEFINITION 3 (TREND). The trend of a time series refers to the long-term changes or patterns that occur over time. Intuitively, it represents the general direction in which the data is moving. Referring to the explained variance [65], Trend Strength can be defined as follows.

$$\text{Trend_Strength} = \max \left(0, 1 - \frac{\text{var}(R)}{\text{var}(X - S)} \right), \quad (1)$$

where X is a time series that can be decomposed into a trend (T), seasonality (S), and the remainder (R): $X = S + T + R$ by employing Seasonal and Trend decomposition using Loess, which is a highly versatile and robust method for time series decomposition [17].

DEFINITION 4 (SEASONALITY). Seasonality refers to the phenomenon where changes in a time series repeat at specific intervals. Similar to the measurement of trend, the strength of Seasonality of a time

series X can be estimated as follows.

$$\text{Seasonality_Strength} = \max \left(0, 1 - \frac{\text{var}(R)}{\text{var}(X - T)} \right) \quad (2)$$

DEFINITION 5 (STATIONARITY). Stationarity refers to the mean of any observation in a time series $X = \langle x_1, x_2, \dots, x_n \rangle$ is constant, and the variance is finite for all observations. Also, the covariance $\text{cov}(x_i, x_j)$ between any two observations x_i and x_j depends only on their distance $|j - i|$, i.e., $\forall i + r \leq n, j + r \leq n$ ($\text{cov}(x_i, x_j) = \text{cov}(x_{i+r}, x_{j+r})$). Strictly stationary time series are rare in practice. Therefore, weak stationarity conditions are commonly applied [45] [63]. In our paper, we also exclusively focus on weak stationarity.

We adopt the Augmented Dick-Fuller (ADF) test statistic [21] to quantify stationarity. The stationarity can be calculated as follows.

$$\text{Stationarity} = \begin{cases} \text{True}, & \text{ADF}(X) \leq 0.05 \\ \text{False}, & \text{ADF}(X) > 0.05 \end{cases} \quad (3)$$

DEFINITION 6 (SHIFTING). Shifting refers to the phenomenon where the probability distribution of time series changes over time. This behavior can stem from structural changes within the system, external influences, or the occurrence of random events. As the value approaches 1, the degree of shifting becomes more severe. Algorithm 1 details the calculation process.

DEFINITION 7 (TRANSITION). Transition refers to the trace of the covariance of transition matrix between symbols in a 3-letter alphabet [55]. It captures the regular and identifiable fixed features present in a time series, such as the clear manifestation of trends, periodicity, or the simultaneous presence of both seasonality and trend. Algorithm 2 details the calculation process.

DEFINITION 8 (CORRELATION). Correlation refers to the possibility that different variables in a multivariate time series may share common trends or patterns, indicating that they are influenced by similar factors or have some underlying relationship. Correlation is calculated as follows.

First, use Catch22 [55] to extract 22 features for each variable in X , serving as a representation for each variable. $F = \langle F^1, F^2, \dots, F^N \rangle \in \mathbb{R}^{22 \times N}$ is obtained using the formulation specified in Equation 4.

$$F = \text{Catch22}(X) \quad (4)$$

Second, according to Equation 5 we calculate the Pearson correlation coefficient between all variable pairs.

$$P = \{r(F^i, F^j) \mid 1 \leq i \leq N, i + 1 \leq j \leq N, i, j \in N^*\} \quad (5)$$

Third, compute the mean and variance of all Pearson correlation coefficients (PCCs) [18] to obtain the correlation, as shown in Equation 6.

$$\text{Correlation} = \text{mean}(P) + \frac{1}{1 + \text{var}(P)} \quad (6)$$

Figure 1 shows time series exhibit diverse characteristics in real-world scenarios. For each characteristic, the calculated characteristic values in the upper right corner of the image corresponding to series with distinct and indistinct characteristic are significantly different. In summary, the time series characteristics define here enable a deep understanding of time series. Using these characteristics for classifying time series, we can more selectively choose forecasting methods that are appropriate for specific scenarios.

Table 4: Statistics of univariate datasets.

Frequency	#Series	Seasonality	Trend	Shifting	Transition	Stationarity	$ TS < 300^1$	F^2
Yearly	1,500	611	1,086	978	633	354	1,499	6
Quarterly	1,514	486	933	889	894	471	1,508	8
Monthly	1,674	883	884	778	1,212	667	1,026	18
Weekly	805	253	330	445	407	372	473	13
Daily	1,484	374	502	487	1,176	714	442	14
Hourly	706	435	276	284	680	472	0	48
Other	385	75	248	236	195	124	215	8
Total	8,068	3,117	4,259	4,097	5,197	3,174	5,163	

¹ $|TS| < 300$: the length of univariate time series < 300

² F : the forecasting horizon

Table 5: Statistics of multivariate datasets.

Dataset	Domain	Frequency	Lengths	Dim	Split
METR-LA [47]	Traffic	5 mins	34,272	207	7:1:2
PEMS-BAY [47]	Traffic	5 mins	52,116	325	7:1:2
PEMS04 [77]	Traffic	5 mins	16,992	307	6:2:2
PEMS08 [77]	Traffic	5 mins	17,856	170	6:2:2
Traffic [90]	Traffic	1 hour	17,544	862	7:1:2
ETTh1 [104]	Electricity	1 hour	14,400	7	6:2:2
ETTh2 [104]	Electricity	1 hour	14,400	7	6:2:2
ETTm1 [104]	Electricity	15 mins	57,600	7	6:2:2
ETTm2 [104]	Electricity	15 mins	57,600	7	6:2:2
Electricity [84]	Electricity	1 hour	26,304	321	7:1:2
Solar [43]	Energy	10 mins	52,560	137	6:2:2
Wind [46]	Energy	15 mins	48,673	7	7:1:2
Weather [90]	Environment	10 mins	52,696	21	7:1:2
AQShunyi [100]	Environment	1 hour	35,064	11	6:2:2
AQWan [100]	Environment	1 hour	35,064	11	6:2:2
ZafNoo [71]	Nature	30 mins	19,225	11	7:1:2
CzeLan [71]	Nature	30 mins	19,934	11	7:1:2
FRED-MD [58]	Economic	1 month	728	107	7:1:2
Exchange [43]	Economic	1 day	7,588	8	7:1:2
NASDAQ [23]	Stock	1 day	1,244	5	7:1:2
NYSE [23]	Stock	1 day	1,243	5	7:1:2
NN5 [80]	Banking	1 day	791	111	7:1:2
ILI [90]	Health	1 week	966	7	7:1:2
Covid-19 [68]	Health	1 day	1,392	948	7:1:2
Wike2000 [26]	Web	1 day	792	2,000	7:1:2

4 TFB: BENCHMARK DETAILS

We proceed to cover the design of the TFB benchmark. First, we provide a detailed overview of the dataset in TFB, including the process of collection, accompanied by a comprehensive demonstration. (Section 4.1). Second, we categorize the existing methods supported (Section 4.2). Third, we cover evaluation strategies and metrics (Section 4.3). Finally, we describe the full benchmark pipeline (Section 4.4).

4.1 Datasets

We equip TFB with a set of 25 multivariate and 8,068 univariate datasets with the following desirable properties. All datasets are formatted consistently. The collection is comprehensive, covering a wide range of domains and characteristics. The values of key characteristics of the multivariate and univariate datasets, as well as their classification based on the characteristics, can be found in our code repository. This marks an improvement, addressing challenges such as different formats, varied documentation, and the time-consuming nature of dataset collection. We will provide a brief introduction to these datasets (Section 4.1.1), demonstrating their comprehensiveness (Section 4.1.2).

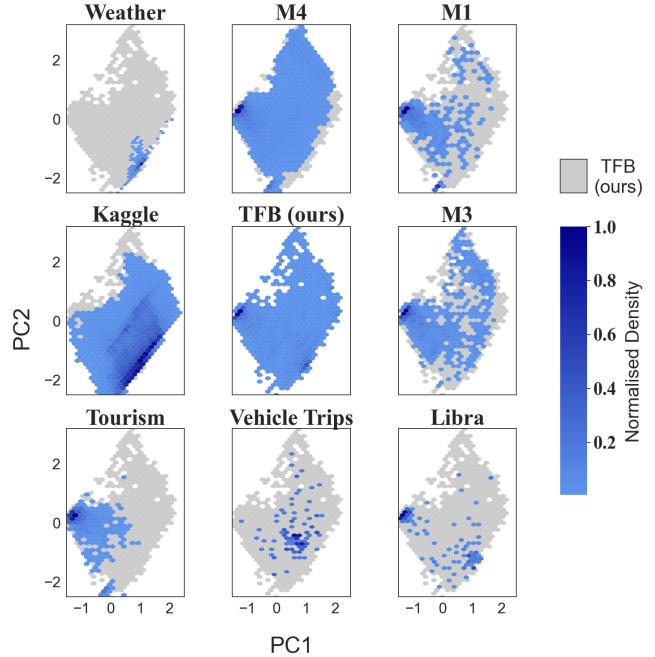


Figure 5: Hexbin plots showing normalised density values of the low-dimensional feature spaces generated by PCA across trend, seasonality, stationarity, shifting, and transition for 9 univariate datasets.

4.1.1 Dataset overview.

Univariate time series. The univariate datasets are carefully curated from 16 open-source datasets, thus covering dozens of domains. To fully reflect the complexity of real-world time series, we employed Principal Feature Analysis (PFA) [54], a variation of Principal Component Analysis (PCA) [6]. PFA preserves the original values of individual time series data points. We employ the concept of explained variance, representing the ratio between the variance of a single time series and the sum of variances across all individual time series. A threshold t for the explained variance is set to 0.9. This implies that for each dataset, we choose to retain the minimum number of time series required to encompass 90% of the variance contributed by the remaining time series. As a result, the selected data exhibits pronounced heterogeneity. Compared to datasets with strong homogeneity, it can better reflect methods performance. In the end, we select 8,068 time series to enable the combined dataset to capture the diversity of real-world time series. Statistical information is reported in Table 4.

Multivariate time series. Table 5 lists statistics of the 25 multivariate time series datasets, which cover 10 domains. The frequencies vary from 5 minutes to 1 month, the range of feature dimensions varies from 5 to 2,000, and the sequence length varies from 728 to 57,600. This substantial diversity of the datasets enables comprehensive studies of forecasting methods. To ensure fair comparisons, we choose a fixed data split ratio for each dataset chronologically, i.e., 7:1:2 or 6:2:2, for training, validation and testing.

4.1.2 *Dataset comprehensiveness.* We proceed to investigate the coverage of the selected univariate and multivariate datasets.

Univariate time series. Since time series have different lengths, we first represent time series as vectors that consist of five feature indicators: trend, seasonality, stationarity, shifting, and transition. For ease of visualization, we adopt PCA [6] to reduce the dimensionality from five to two and visualize the eight most widely distributed univariate time series datasets in hexbin—see Figure 5. We observe that TFB and M4 cover the most cells, while all other benchmarks are smaller than TFB. This emphasizes the coverage of our dataset in terms of diversity in characteristic distribution. In addition, compared to M4, our dataset covers a wider range of domains. Further, we note that M4 has a much larger sample size, totaling 100,000, compared to our dataset that contains only 8,068 time series. We believe that testing on diverse datasets is essential to better reflect the practical performance of methods. Yet, we can run much fewer experiments with TFB than the M4 dataset, i.e., around 8%.

Multivariate time series. Figure 2 shows a comparison of TFB and existing multivariate time series benchmarks according to dataset domain distribution. Our benchmark includes more diverse datasets in terms of both quantity and domain categories. Next, we select TSlib, whose multivariate time series datasets are the most popular, to compare the characteristic distributions with TFB, which are shown in Figure 3. We can observe that the datasets in TFB also represent a more diverse characteristic distribution.

4.2 Comparison Methods

To investigate the strengths and limitations of different forecasting methods, we include 22 methods that can be categorized as statistical learning, machine learning, and deep learning methods. In terms of statistical learning methods, we include ARIMA [4], ETS [37], Kalman Filter (KF) [31], and VAR [82]. Among the machine learning methods, we include XGBModel (XGB) [32], LinearRegression (LR) [32, 40], and Random Forest (RF) [5]. Finally, we further split the deep learning methods into RNN-based models (RNN [42]), CNN-based models (MICN [86], TimesNet [89], and TCN [2]), MLP-based models (NLinear [98], DLinear [98], TiDE [19], N-HiTS [9], and N-BEATS [66]), Transformer-based models (PatchTST [64], Crossformer [101], and FEDformer [106], Non-stationary Transformer (Stationary) [53], Informer [104], and Triformer [13]), and Model-Agnostic models (FiLM [105]).

4.3 Evaluation Settings

4.3.1 *Evaluation strategies.* To assess the forecasting accuracy of a method, TFB implements two distinct evaluation strategies: 1) *Fixed forecasting*; and 2) *Rolling forecasting*.

Fixed forecasting. Given a time series on length n , f future time points are predicted from $n - f$ historical time points, which is illustrated in Figure 6a.

Rolling forecasting. In Rolling forecasting, illustrated in Figure 6b, the blue squares represent historical data, the green squares represent the forecasting horizon, and the white squares represent the remaining data in the time series. During rolling forecast, excluding the last iteration, the historical data expands by a fixed number of steps (called the stride) in each iteration. In the final iteration, the historical data is extended to cover the entire time series along with

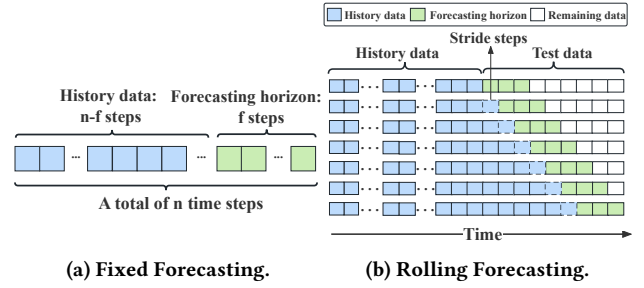


Figure 6: Time series forecasting evaluation strategies.

the forecasting horizon. In each iteration of the inference process, the forecasting model is applied based on the historical data to forecast the designated forecasting horizon. Then, we calculate the average of the error metrics for each iteration. In statistical learning methods (e.g., ARIMA [4], ETS [22]), it is common to use the entire or a subset of the historical data for training and then make forecasting during each iteration. In contrast, in deep learning or machine learning methods, each iteration often involves using only the last portion of the historical data, with a length equal to the look-back windows size, for inference to make forecasting. The current evaluation strategy for TSF, such as TimesNet [89], aligns with our defined criteria.

From a practical perspective, despite some limitations in the generalization capabilities of statistical learning methods, their relatively short runtimes prompts us to adopt the approach of retraining and then making inference predictions to support rolling forecasting. However, in the case of deep learning and machine learning methods, each iteration of retraining typically takes much longer time. Therefore, to balance timeliness and prediction accuracy, machine learning and deep learning methods opt for the strategy of reinferring during each iteration in rolling forecasting.

4.3.2 *Evaluation metrics.* To achieve an extensive evaluation of forecasting performance, we employ eight error metrics, namely the Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), Mean Squared Error (MSE), Symmetric Mean Absolute Percentage Error (SMAPE), Root Mean Squared Error (RMSE), Weighted Absolute Percent Error (WAPE), Modified Symmetric Mean Absolute Percentage Error (MSMAPE) [78], and Mean Absolute Scaled Error (MASE) [38], which are defined in Equations 7–14.

$$MAE = \frac{1}{h} \sum_{k=1}^h |F_k - Y_k| \quad (7) \quad MAPE = \frac{1}{h} \sum_{k=1}^h \frac{|Y_k - F_k|}{Y_k} \times 100\% \quad (8)$$

$$MSE = \frac{1}{h} \sum_{k=1}^h (F_k - Y_k)^2 \quad (9) \quad SMAPE = \frac{100\%}{h} \sum_{k=1}^h \frac{|F_k - Y_k|}{\frac{|Y_k| + |F_k|}{2}} \quad (10)$$

$$RMSE = \sqrt{\frac{1}{h} \sum_{k=1}^h (F_k - Y_k)^2} \quad (11) \quad WAPE = \frac{\sum_{k=1}^h |Y_k - F_k|}{\sum_{k=1}^h |Y_k|} \quad (12)$$

$$MSMAPE = \frac{100\%}{h} \sum_{k=1}^h \frac{|F_k - Y_k|}{\max(|Y_k| + |F_k| + \epsilon, 0.5 + \epsilon)/2} \quad (13)$$

$$MASE = \frac{\sum_{k=M+1}^{M+h} |F_k - Y_k|}{\frac{h}{M-S} \sum_{k=S+1}^M |Y_k - Y_{k-S}|}, \quad (14)$$

where M is the length of the training series, S is the seasonality of the dataset, h is the forecasting horizon, F_k are the generated

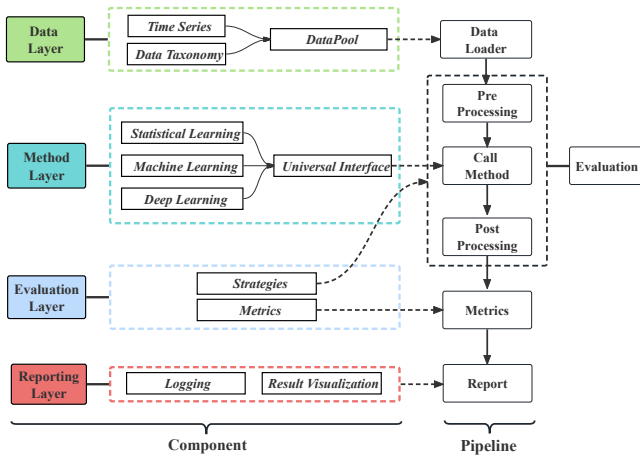


Figure 7: The TFB pipeline.

forecasts, and Y_k are the actual values. We set the parameter ϵ in Equation 13 to its proposed default of 0.1. For rolling forecasting, we further calculate the average of error metrics for all samples (windows) on each time series to assess the method performance.

4.4 Unified Pipeline

As mentioned in Section 1, small implementation differences can impact evaluation results profoundly. To enable fair and comprehensive comparisons of methods, we introduce a unified pipeline that is structured into a data layer, a method layer, an evaluation layer, and a reporting layer—see Figure 7. The details of each component are listed as follows.

- The data layer is a repository of univariate and multivariate time series from diverse domains, structured according to their distinct characteristics, frequencies, and sequence lengths. The data is uniformly according to a standardized format. When a new dataset becomes available, this layer can assess whether the distribution of existing datasets across the six features can be expanded. If yes, it is accepted as a new dataset.
- The method layer supports embedding statistical learning, machine learning and deep learning methods. However, other benchmarks have not achieved this, most of them can only embed deep learning methods. Additionally, TFB is designed to be compatible with any third-party TSF library, such as Darts [32], TSlib [89]. Users can easily integrate forecasting methods implemented in third-party libraries into TFB by writing a simple Universal Interface, facilitating fair comparisons. TFB goes beyond benchmarks that only support Direct Multi-Step (DMS) forecasting to support also Iterative Multi-Step (IMS) forecasting. The method layer thus contributes to the applicability of TFB by offering support for a broad range of methods.
- The evaluation layer offers support for a diverse range of evaluation strategies and metrics. It thus supports both fixed and rolling forecasting strategies, increasing again the applicability of the benchmark to a wider range of methods and applications. The layer also covers evaluation metrics found in other studies

and enables the use of customized metrics for a more comprehensive assessment of method performance. Besides, for each evaluation strategy, TFB offers standardized dataset handling, splitting, and normalization. Additionally, it provides a standard configuration file that can be customized by users. This aims to facilitate deeper understanding of method performance in different settings.

- The reporting layer encompasses a logging system for tracking information, enabling the capture of experimental settings to enable traceability. Further, it encompasses a visualization module to facilitate a clear understanding of method performance. This design is to offer exhaustive support and transparency throughout the evaluation process.

Users only need to deploy their method architecture at the method layer and choose or configure the configuration file, then TFB can automatically run the pipeline in Figure 7.

The TFB pipeline supports various extensible features. Compatibility with CPU and GPU hardware enables evaluation in different computing environments. TFB also supports both sequential and parallel program execution, providing users with multiple options.

In summary, TFB is a unified, flexible, scalable, and user-friendly benchmarking tool for TSF methods. It enables to better understand, compare, and select TSF methods for specific application scenarios.

5 EXPERIMENTS

We report on experiments with the 14 multivariate and 22 univariate forecasting methods covered in Section 4.2 on all the datasets covered in Section 4.1. We adopt the pipeline covered in Section 4.4 to do so. For each method, we conduct comprehensive hyper-parameters selection, such that its performance result approach or surpasses the performance reported in its original paper. Our goal is to showcase TFB as an easy-to-use and powerful resource, providing a convenient and fair evaluation platform for TSF methods.

5.1 Experimental Setup

5.1.1 Datasets and Comparison methods. We include all the datasets included in TFB and all methods mentioned in Section 4.2.

5.1.2 Implementation details. For multivariate forecasting, we use the rolling forecasting strategy. We consider four forecasting horizon F : 24, 36, 48, and 60, for FredMd, NASDAQ, NYSE, NN5, ILI, Covid-19, and Wike2000, and we use another four forecasting horizon, 96, 192, 336, and 720, for all other datasets which have longer lengths. The look-back window H underwent testing with lengths 36 and 104 for FredMd, NASDAQ, NYSE, NN5, ILI, Covid-19, and Wike2000, and 96, 336, and 512 for all other datasets. For univariate forecasting, we adopt a fixed forecasting strategy to maintain consistency with the setting of the M4 competition [57], forecasting horizons go from 6 to 48, with the look-back window H set to 1.25 times forecasting horizon F .

For each method, we adhere to the hyper-parameter as specified in their original papers. Additionally, we perform hyper-parameter searches across multiple sets, with a limit of 8 sets. The optimal result is then selected from these evaluations, contributing to a comprehensive and unbiased assessment of each method’s performance. Due to space constraints, we only report the results of a

subset of metrics in the paper, additional metrics results can be found in the code repository.

All experiments are conducted using PyTorch [69] in Python 3.8 and execute on an NVIDIA Tesla-A800 GPU. The training process is guided by the L2 loss, employing the ADAM optimizer. Initially, the batch size is set to 32, with the option to reduce it by half (to a minimum of 8) in case of an Out-Of-Memory (OOM) situation. We do not use the "Drop Last" operation during testing. To ensure reproducibility and facilitate experimentation, datasets and code are available at: <https://github.com/decisionintelligence/TFB>.

5.2 Experimental Results

5.2.1 Univariate time series forecasting. Table 6 reports the average results of UTSF in terms of the metrics MASE, MSMAPE, and the Ranks in MSMAPE that indicates how many times the best performance is achieved on the datasets. We observe that the recently proposed deep learning methods, including TimesNet [89], PatchTST [64], and N-HiTS [9], exhibit substantially better average performance on univariate datasets in terms of MASE and MSMAPE. However, when considering Ranks, the (non-deep) machine learning methods LinearRegression (LR) [32, 40] and RandomForest (RF) [5] outperform all competitors. This suggests that the machine learning methods may be more suitable in specific scenarios. In this scenario, each individual univariate time series is adopted to train a separate model, and deep learning methods require large amount of training data to be effective. Therefore, the performance of the deep learning methods falls short. Next, we can observe that LR performs better on the time series with seasonality, trend, and shifting characteristics, while RF is better when those patterns are absent. Further, we notice that LR is more suitable for data without stationarity than RF. Finally, both LR and RF are sensitive to the transition characteristic: the stronger the characteristic, the better. These results offer guidance on choosing the right method for a specific setting.

5.2.2 Multivariate time series forecasting. Due to the large number of results, we report them in two tables, Tables 7 and 8. The datasets in these tables are ordered based on their trend characteristic, where datasets with weaker trend are occur first. In both tables, we report the MAE and MSE on normalized data considering four different forecasting horizons for each dataset. "nan" represents the inability of a method to generate effective predictions, while "inf" represents infinite results. We see that no single method achieves the best performance on all datasets. We also see that the Transformer-based methods generally outperform other methods on datasets with weak trend. Next, Linear-based methods tend to perform moderately better on datasets with strong trend. Surprisingly, we observe that recent methods not consistently outperform earlier studies, such as Informer, LR, and VAR. This discovery highlights the need for evaluating method performance across diverse datasets. Evaluating methods on relatively few datasets render it difficult to accurately assess their universality and overall performance. Therefore, it is crucial to expand the range of datasets used in evaluations.

5.2.3 Performance on different characteristics. We proceed to study the performance of different deep learning methods with respect to different characteristics. First, we score the multivariate time




Figure 8: MAE results of methods across six characteristics.

series with respect to the six characteristics we consider. Then, we select the dataset with the highest score for each characteristic, which are FRED-MD on trend, Electricity on seasonality, PEMS08 on transition, NYSE on shifting, PEMS-BAY on correlation, and Solar on stationarity. Next, we show the best MAE results for the methods in a radar figure. We use a forecasting horizon of 24 for FRED-MD and NYSE and of 96 for other datasets—see Figure 8. We see that no deep learning method excels on all datasets. In particular, Crossformer demonstrates exceptional performance on datasets where transition is highly pronounced (PEMS08), the data is most stationary (Solar), and the data is most correlated (PEMS-BAY). However, Crossformer’s performance is noticeably inferior to other methods on time series with other characteristics. Next, PatchTST achieves optimal performance on datasets with strong seasonality (Electricity). Similarly, NLinear delivers outstanding results on time series with the most significant trend (FRED-MD) and severe drift (NYSE). Both PatchTST and NLinear perform well consistently, without any notably poor outcomes.

5.3 Hints to Method Design

5.3.1 Transformers vs. linear methods. To study the impact of different data characteristics on these two types of methods, we consider the best MAE results obtained by CNN, Linear and Transformer. We use a forecasting horizon of 24 for NASDAQ, NYSE, and NN5 and of 96 for other datasets—see Figure 9. We choose CNN as a reference to gain a more comprehensive understanding of the performance of Transformer and Linear methods. We have the following observations. First, each method exhibits distinct advantages on datasets with different characteristics. Second, Linear-based methods excel when the dataset shows an increasing trend or significant shifts. This can be attributed to the linear modeling capabilities of the Linear model, making it well-suited for capturing linear trends and shifts. Third, Transformer-based methods outperform Linear-based methods on datasets exhibiting marked seasonality, stationarity, and nonlinear patterns, as well as more pronounced patterns or strong internal similarities. This superiority may stem from the enhanced nonlinear modeling capabilities of Transformer-based

Table 6: Univariate forecasting results.

Dataset	Metric	PatchTST	Crossformer	FEDformer	Stationary	Informer	Triformer	DLinear	NLinear	TiDE	N-BEATS	N-HiTS	TimesNet	TCN	RNN	FILM	LR	RF	XGB	ARIMA	ETS	KF
Seasonality	mase	1.660	29.704	2.100	2.384	2.390	19.378	2.409	2.850	2.074	2.081	2.189	1.446	24.159	30.231	1.882	7.8e+9	1.649	1.715	2.830	4.091	9.002
	mssape	12.263	161.074	19.041	15.190	14.301	107.957	19.628	20.938	14.566	15.794	13.557	10.927	121.335	149.830	13.537	19.183	12.790	13.404	19.868	26.386	57.409
	Ranks	167	14	59	40	85	15	22	45	162	168	213	314	14	8	141	603	425	250	225	92	43
Trend	mase	1.639	23.704	1.879	1.638	1.601	16.496	1.949	1.733	2.526	1.677	1.678	1.478	15.441	23.889	1.769	2.5e+10	1.731	1.814	1.496	1.544	3.318
	mssape	21.671	166.859	26.766	22.050	21.413	138.945	26.966	27.154	30.968	25.705	24.127	20.497	140.174	162.648	22.331	33.413	25.316	26.838	24.491	25.190	44.551
	Ranks	214	35	208	200	303	72	61	157	189	383	556	371	64	40	138	326	423	273	355	268	292
Stationarity	mase	2.220	41.287	2.758	2.651	2.658	28.091	3.016	2.914	3.316	2.512	2.553	1.911	28.716	44.365	2.492	7.9e+9	2.271	2.355	2.822	3.615	8.486
	mssape	10.679	184.125	13.917	11.709	11.216	133.424	14.435	13.463	13.747	11.583	11.090	9.247	136.243	180.218	11.442	16.920	10.832	11.221	11.686	12.986	50.062
	Ranks	201	2	114	113	188	20	51	113	201	270	375	402	6	2	162	737	298	246	403	222	123
Transition	mase	1.007	8.941	1.077	1.127	1.064	5.878	1.131	1.326	1.272	1.073	1.116	0.968	7.718	8.282	1.052	3.1e+10	1.059	1.127	1.104	1.311	2.185
	mssape	26.261	142.824	34.808	27.894	26.993	119.760	34.972	37.374	36.802	33.385	30.053	25.243	129.153	136.076	27.307	40.210	21.262	33.307	35.019	39.814	48.911
	Ranks	180	47	153	127	200	67	32	89	150	281	394	283	72	46	117	192	550	277	177	138	212
Stationarity	mase	1.004	9.380	1.057	1.132	1.133	6.309	1.139	1.290	1.343	1.162	1.212	0.961	7.870	8.305	1.066	15.848	1.043	1.100	1.257	1.618	3.172
	mssape	27.024	135.888	35.122	28.539	27.546	114.323	35.434	37.306	37.594	33.519	31.080	26.120	122.194	129.738	28.172	38.320	32.232	34.281	36.212	41.513	52.234
	Ranks	154	45	128	105	177	67	24	69	124	214	285	242	66	42	99	197	444	219	150	117	180
Transition	mase	1.397	19.759	1.571	1.744	1.779	11.972	1.820	1.985	1.827	1.543	1.601	1.282	13.744	19.901	1.505	5.7e+4	1.380	1.474	1.930	2.723	3.998
	mssape	21.932	155.700	24.803	23.707	22.978	117.240	25.741	29.013	27.664	23.031	22.672	20.869	125.136	150.624	22.973	29.179	23.002	24.545	25.020	28.725	45.521
	Ranks	242	44	187	153	263	79	44	121	230	373	527	464	73	48	187	560	533	382	304	186	166
Shifting	mase	2.102	37.372	2.676	2.277	2.136	27.831	2.682	2.489	3.303	2.358	2.371	1.799	27.987	39.651	2.369	5.3e+10	2.278	2.323	2.157	2.172	8.259
	mssape	10.984	180.775	21.932	11.399	10.860	144.591	21.216	17.039	19.143	19.785	15.284	9.435	146.942	173.676	11.622	25.629	15.905	16.404	18.520	20.089	56.754
	Ranks	139	5	80	87	125	8	39	81	121	178	242	221	5	0	92	369	315	141	276	174	169
Stationarity	mase	2.138	36.092	2.646	2.507	2.314	25.570	2.823	2.747	2.975	2.289	2.345	1.857	24.925	37.921	2.352	3.7e+10	2.224	2.306	2.331	2.799	6.862
	mssape	13.453	173.924	19.874	14.454	13.509	133.554	19.930	19.013	17.860	16.573	14.248	11.873	137.113	167.496	14.074	21.775	14.877	16.159	17.550	20.517	50.844
	Ranks	181	13	107	86	194	27	47	123	173	268	380	384	14	18	143	556	401	220	370	227	157
Transition	mase	1.142	15.639	1.262	1.355	1.485	9.401	1.410	1.563	1.709	1.363	1.390	1.062	12.499	15.066	1.257	7.4e+4	1.159	1.229	1.681	2.248	4.122
	mssape	22.763	155.031	27.811	24.258	23.983	120.183	28.465	30.674	31.618	27.347	26.022	21.881	128.544	148.943	23.945	34.251	26.254	27.310	28.022	30.950	48.150
	Ranks	200	36	160	154	194	60	36	79	178	283	389	301	64	30	136	373	447	303	210	133	178

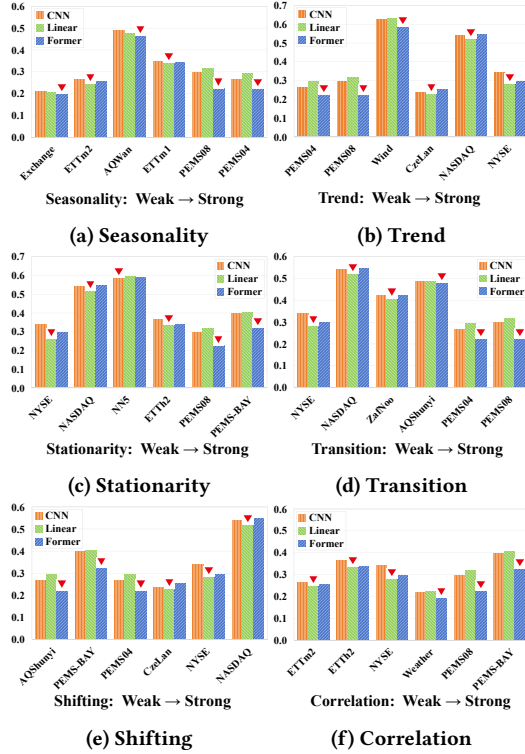


Figure 9: Comparison between Transformer-based, CNN-based, and Linear-based methods. Red triangles indicate methods with the best accuracy (minimum MAE).

methods, enabling them to flexibly adapt to complex time series patterns and intrinsic correlations.




Figure 10: Method performance for varying correlation within datasets.

Therefore, the observed phenomenon underscores the inherent differences between Transformer-based and Linear-based methods at addressing diverse characteristics of time series. To achieve optimal performance, we recommend selecting the appropriate method based on the characteristics of the relevant time series, allowing the method to fully leverage its strengths.

5.3.2 Channel independence vs. channel dependence. In multivariate datasets, variables are occasionally referred to as channels. To study the impact of channel dependency in multivariate time series, we compare PatchTST and Crossformer on ten datasets with dependencies ranging from weak to strong. We report the MAE for forecasting horizon F is 96 in Figure 10. We observe that as the correlations within a dataset increase, the performance of Crossformer gradually surpasses that of PatchTST, suggesting that it is better to consider channel dependencies when correlations are strong. However, when correlations among variables are not pronounced, PatchTST that does not consider channel dependencies is better.

REFERENCES

- [1] Francisco Martinez Alvarez, Alicia Troncoso, Jose C Riquelme, and Jesus S Aguilar Ruiz. 2010. Energy time series forecasting based on pattern sequence similarity. *IEEE Transactions on Knowledge and Data Engineering* 23, 8 (2010), 1230–1243.
- [2] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. 2018. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271* (2018).
- [3] André Bauer, Marwin Züfle, Simon Eismann, Johannes Grohmann, Nikolas Herbst, and Samuel Kounev. 2021. Libra: A benchmark for time series forecasting methods. In *ICPE*. 189–200.
- [4] George EP Box and David A Pierce. 1970. Distribution of residual autocorrelations in autoregressive-integrated moving average time series models. *Journal of the American statistical Association* 65, 332 (1970), 1509–1526.
- [5] Leo Breiman. 2001. Random forests. *Machine learning* 45 (2001), 5–32.
- [6] Rasmus Bro and Age K Smilde. 2014. Principal component analysis. *Analytical methods* 6, 9 (2014), 2812–2831.
- [7] David Campos, Tung Kieu, Chenjuan Guo, Feiteng Huang, Kai Zheng, Bin Yang, and Christian S. Jensen. 2022. Unsupervised Time Series Outlier Detection with Diversity-Driven Convolutional Ensembles. *Proc. VLDB Endow.* 15, 3 (2022), 611–623.
- [8] David Campos, Bin Yang, Tung Kieu, Miao Zhang, Chenjuan Guo, and Christian S Jensen. 2024. QCore: Data-Efficient, On-Device Continual Calibration for Quantized Models—Extended Version. *arXiv preprint arXiv:2404.13990* (2024).
- [9] Cristian Challu, Kin G Olivares, Boris N Oreshkin, Federico Garza Ramirez, Max Mergenthaler Canseco, and Artur Dubrawski. 2023. Nhits: Neural hierarchical interpolation for time series forecasting. In *AAAI*, Vol. 37. 6989–6997.
- [10] Peng Chen, Yingying Zhang, Yunyao Cheng, Yang Shu, Yihang Wang, Qingsong Wen, Bin Yang, and Chenjuan Guo. 2024. Pathformer: Multi-scale transformers with Adaptive Pathways for Time Series Forecasting. *arXiv preprint arXiv:2402.05956* (2024).
- [11] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *SIGKDD*. 785–794.
- [12] Yunyao Cheng, Peng Chen, Chenjuan Guo, Kai Zhao, Qingsong Wen, Bin Yang, and Christian S Jensen. 2023. Weakly guided adaptation for robust time series forecasting. *Proc. VLDB Endow.* 17, 4 (2023), 766–779.
- [13] Razvan-Gabriel Cirstea, Chenjuan Guo, Bin Yang, Tung Kieu, Xuanyi Dong, and Shirui Pan. 2022. Triformer: Triangular, Variable-Specific Attentions for Long Sequence Multivariate Time Series Forecasting. In *IJCAI*. 1994–2001.
- [14] Razvan-Gabriel Cirstea, Bin Yang, Chenjuan Guo, Tung Kieu, and Shirui Pan. 2022. Towards Spatio-Temporal Aware Traffic Time Series Forecasting. In *ICDE*. 2900–2913.
- [15] Razvan-Gabriel Cirstea, Tung Kieu, Chenjuan Guo, Bin Yang, and Sinno Jialin Pan. 2021. EnhanceNet: Plugin Neural Networks for Enhancing Correlated Time Series Forecasting. In *ICDE*. 1739–1750.
- [16] Razvan-Gabriel Cirstea, Bin Yang, and Chenjuan Guo. 2019. Graph Attention Recurrent Neural Networks for Correlated Time Series Forecasting. In *MileTS19@KDD*.
- [17] Robert B Cleveland, William S Cleveland, Jean E McRae, and Irma Terpenning. 1990. STL: A seasonal-trend decomposition. *J. Off. Stat* 6, 1 (1990), 3–73.
- [18] Israel Cohen, Yiteng Huang, Jingdong Chen, Jacob Benesty, Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. 2009. Pearson correlation coefficient. *Noise reduction in speech processing* (2009), 1–4.
- [19] Abhimanyu Das, Weihao Kong, Andrew Leach, Rajat Sen, and Rose Yu. 2023. Long-term Forecasting with TiDE: Time-series Dense Encoder. *arXiv preprint arXiv:2304.08424* (2023).
- [20] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*. 248–255.
- [21] Graham Elliott, Thomas J Rothenberg, and James H Stock. 1992. Efficient tests for an autoregressive unit root.
- [22] Cristian Challú Kin G. Olivares Federico Garza, Max Mergenthaler Canseco. 2022. StatsForecast: Lightning fast forecasting with statistical and econometric models. PyCon Salt Lake City, Utah, US 2022.
- [23] Fuli Feng, Xiangnan He, Xiang Wang, Cheng Luo, Yiqun Liu, and Tat-Seng Chua. 2019. Temporal relational ranking for stock prediction. *ACM Transactions on Information Systems* 37, 2 (2019), 1–30.
- [24] Jan Alexander Fischer, Philipp Pohl, and Dietmar Ratz. 2020. A machine learning approach to univariate time series forecasting of quarterly earnings. *Review of Quantitative Finance and Accounting* 55 (2020), 1163–1179.
- [25] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* (2001), 1189–1232.
- [26] Jan Gasthaus, Konstantinos Benidis, Yuyang Wang, Syama Sundar Rangapuram, David Salinas, Valentin Flunkert, and Tim Januschowski. 2019. Probabilistic forecasting with spline quantile function RNNs. In *AISTATS*. 1901–1910.
- [27] Rakshitha Godaheva, Christoph Bergmeir, Geoffrey I Webb, Rob J Hyndman, and Pablo Montero-Manso. 2021. Monash time series forecasting archive. *arXiv preprint arXiv:2105.06643* (2021).
- [28] Chenjuan Guo, Christian S. Jensen, and Bin Yang. 2014. Towards Total Traffic Awareness. *SIGMOD Record* 43, 3 (2014), 18–23.
- [29] Chenjuan Guo, Bin Yang, Ove Andersen, Christian S Jensen, and Kristian Torp. 2015. Ecomark 2.0: empowering eco-routing with vehicular environmental models and actual vehicle fuel consumption data. *Geoinformatica* 19 (2015), 567–599.
- [30] Chenjuan Guo, Bin Yang, Jilin Hu, Christian S. Jensen, and Lu Chen. 2020. Context-aware, preference-based vehicle routing. *The VLDB Journal* 29, 5 (2020), 1149–1170.
- [31] Andrew C Harvey. 1990. Forecasting, structural time series models and the Kalman filter. (1990).
- [32] Julien Herzen, Francesco Lässig, Samuele Giuliano Piazzetta, Thomas Neuer, Léo Tafti, Guillaume Raille, Tomas Van Pottelbergh, Marek Pasielka, Andrzej Skrodzki, Nicolas Huguenin, et al. 2022. Darts: User-friendly modern machine learning for time series. *The Journal of Machine Learning Research* 23, 1 (2022), 5442–5447.
- [33] Jilin Hu, Chenjuan Guo, Bin Yang, and Christian S Jensen. 2019. Stochastic weight completion for road networks using graph convolutional networks. In *ICDE*. 1274–1285.
- [34] Jilin Hu, Bin Yang, Chenjuan Guo, and Christian S Jensen. 2018. Risk-aware path selection with time-varying, uncertain travel costs: a time series approach. *The VLDB Journal* 27 (2018), 179–200.
- [35] Jilin Hu, Bin Yang, Christian S. Jensen, and Yu Ma. 2017. Enabling time-dependent uncertain eco-weights for road networks. *Geoinformatica* 21, 1 (2017), 57–88.
- [36] Xuanwen Huang, Yang Yang, Yang Wang, Chunping Wang, Zhisheng Zhang, Jiarong Xu, Lei Chen, and Michalis Vazirgiannis. 2022. Dgraph: A large-scale financial dataset for graph anomaly detection. *Advances in Neural Information Processing Systems* 35 (2022), 22765–22777.
- [37] Rob Hyndman, Anne B Koehler, J Keith Ord, and Ralph D Snyder. 2008. *Forecasting with exponential smoothing: the state space approach*.
- [38] Rob J Hyndman and Anne B Koehler. 2006. Another look at measures of forecast accuracy. *International journal of forecasting* 22, 4 (2006), 679–688.
- [39] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems* 30 (2017).
- [40] Benjamin Kedem and Konstantinos Fokianos. 2005. *Regression models for time series analysis*.
- [41] Tung Kieu, Bin Yang, Chenjuan Guo, Christian S. Jensen, Yan Zhao, Feiteng Huang, and Kai Zheng. 2022. Robust and Explainable Autoencoders for Unsupervised Time Series Outlier Detection. In *ICDE*. 3038–3050.
- [42] Taesung Kim, Jinhee Kim, Yunwon Tae, Cheonbok Park, Jang-Ho Choi, and Jaegul Choo. 2021. Reversible instance normalization for accurate time-series forecasting against distribution shift. In *ICLR*.
- [43] Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. 2018. Modeling long-and short-term temporal patterns with deep neural networks. In *SIGIR*. 95–104.
- [44] Chonho Lee, Zhaojing Luo, Kee Yuan Ngiam, Meihui Zhang, Kaiping Zheng, Gang Chen, Beng Chin Ooi, and Wei Luen James Yip. 2017. Big healthcare data analytics: Challenges and applications. *Handbook of large-scale distributed computing in smart healthcare* (2017), 11–41.
- [45] Doyup Lee. 2017. Anomaly detection in multivariate non-stationary time series for automatic DBMS diagnosis. In *ICMLA*. 412–419.
- [46] Yan Li, Xinjiang Lu, Yaqing Wang, and Dejing Dou. 2022. Generative time series forecasting with diffusion, denoise, and disentanglement. *Advances in Neural Information Processing Systems* 35 (2022), 23009–23022.
- [47] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2017. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv preprint arXiv:1707.01926* (2017).
- [48] Yubo Liang, Zezhi Shao, Fei Wang, Zhao Zhang, Tao Sun, and Yongjun Xu. 2022. BasicTS: An Open Source Fair Multivariate Time Series Prediction Benchmark. In *International Symposium on Benchmarking, Measuring and Optimization*. 87–101.
- [49] Yan Lin, Jilin Hu, Shengnan Guo, Bin Yang, Christian S Jensen, Youfang Lin, and Huaiyu Wan. 2024. GenSTL: General Sparse Trajectory Learning via Autoregressive Generation of Feature Domains. *arXiv preprint arXiv:2402.07232* (2024).
- [50] Yan Lin, Huaiyu Wan, Shengnan Guo, Jilin Hu, Christian S Jensen, and Youfang Lin. 2023. Pre-Training General Trajectory Embeddings With Maximum Multi-View Entropy Coding. *IEEE Transactions on Knowledge and Data Engineering* (2023).
- [51] Yan Lin, Huaiyu Wan, Shengnan Guo, and Youfang Lin. 2021. Pre-training context and time aware location embeddings from spatial-temporal trajectories for user next location prediction. In *AAAI*, Vol. 35. 4241–4248.
- [52] Yan Lin, Huaiyu Wan, Jilin Hu, Shengnan Guo, Bin Yang, Youfang Lin, and Christian S Jensen. 2023. Origin-destination travel time oracle for map-based services. *Proceedings of the ACM on Management of Data* 1, 3 (2023), 1–27.

- [53] Yong Liu, Haixu Wu, Jianmin Wang, and Mingsheng Long. 2022. Non-stationary networks: Exploring the stationarity in time series forecasting. *Advances in Neural Information Processing Systems* 35 (2022), 9881–9893.
- [54] Yijuan Lu, Ira Cohen, Xiang Sean Zhou, and Qi Tian. 2007. Feature selection using principal feature analysis. In *ACM MM*. 301–304.
- [55] Carl H Lubba, Sarab S Sethi, Philip Knaute, Simon R Schultz, Ben D Fulcher, and Nick S Jones. 2019. catch22: CAnonical Time-series CHaracteristics: Selected through highly comparative time-series analysis. *Data Mining and Knowledge Discovery* 33, 6 (2019), 1821–1852.
- [56] Spyros Makridakis and Michele Hibon. 2000. The M3-Competition: results, conclusions and implications. *International journal of forecasting* 16, 4 (2000), 451–476.
- [57] Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. 2018. The M4 Competition: Results, findings, conclusion and way forward. *International Journal of Forecasting* 34, 4 (2018), 802–808.
- [58] Michael W McCracken and Serena Ng. 2016. FRED-MD: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics* 34, 4 (2016), 574–589.
- [59] Jie Mei, Dawei He, Ronald Harley, Thomas Habetler, and Guannan Qu. 2014. A random forest method for real-time price forecasting in New York electricity market. In *2014 IEEE PES General Meeting| Conference & Exposition*. 1–5.
- [60] Hao Miao, Yan Zhao, Chenjuan Guo, Bin Yang, Zheng Kai, Feiteng Huang, Jiandong Xie, and Christian S. Jensen. 2024. A Unified Replay-based Continuous Learning Framework for Spatio-Temporal Prediction on Streaming Data. *ICDE* (2024).
- [61] Xiaoye Miao, Yangyang Wu, Jun Wang, Yunjun Gao, Xudong Mao, and Jianwei Yin. 2021. Generative semi-supervised learning for multivariate time series imputation. In *AAAI*, Vol. 35. 8983–8991.
- [62] Xian Mo, Jun Pang, and Zhiming Liu. 2022. THS-GWNN: a deep learning framework for temporal network link prediction. *Frontiers of Computer Science* 16, 2 (2022), 162304.
- [63] Guy P Nason. 2006. Stationary and non-stationary time series. (2006).
- [64] Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2022. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730* (2022).
- [65] Kevin E O’Grady. 1982. Measures of explained variance: Cautions and limitations. *Psychological Bulletin* 92, 3 (1982), 766.
- [66] Boris N Oreshkin, Dmitri Carpov, Nicolas Chapados, and Yoshua Bengio. 2019. N-BEATS: Neural basis expansion analysis for interpretable time series forecasting. *arXiv preprint arXiv:1905.10437* (2019).
- [67] Zhicheng Pan, Yihang Wang, Yingying Zhang, Sean Bin Yang, Yunyao Cheng, Peng Chen, Chenjuan Guo, Qingsong Wen, Xiduo Tian, Yunliang Dou, et al. 2023. Magicscaler: Uncertainty-aware, predictive autoscaling. *Proc. VLDB Endow.* 16, 12 (2023), 3808–3821.
- [68] George Panagopoulos, Giannis Nikolentzos, and Michalis Vazirgiannis. 2021. Transfer graph neural networks for pandemic forecasting. In *AAAI*, Vol. 35. 4838–4845.
- [69] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems* 32 (2019).
- [70] Simon Aagaard Pedersen, Bin Yang, and Christian S. Jensen. 2020. Anytime Stochastic Routing with Hybrid Learning. *Proc. VLDB Endow.* 13, 9 (2020), 1555–1567.
- [71] Rafael Poyatos, Victor Granda, Victor Flo, Mark A Adams, Balázs Adorján, David Aguadé, Marcos PM Aidar, Scott Allen, M Susana Alvarado-Barrientos, Kristina J Anderson-Teixeira, et al. 2020. Global transpiration data from sap flow measurements: the SAPFLUXNET database. *Earth System Science Data Discussions* 2020 (2020), 1–57.
- [72] Xuecheng Qi, Huiqi Hu, Jinwei Guo, Chenchen Huang, Xuan Zhou, Ning Xu, Yu Fu, and Aoying Zhou. 2023. High-availability in-memory key-value store using RDMA and Optane DCPMM. *Frontiers of Computer Science* 17, 1 (2023), 171603.
- [73] Zhongzheng Qiao, Quang Pham, Zhen Cao, Hoang H Le, PN Suganthan, Xudong Jiang, and Ramasamy Savitha. 2024. Class-incremental Learning for Time Series: Benchmark and Evaluation. *arXiv preprint arXiv:2402.12035* (2024).
- [74] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. 2020. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting* 36, 3 (2020), 1181–1191.
- [75] Omer Berat Sezer, Mehmet Ugur Gudelek, and Ahmet Murat Ozbayoglu. 2020. Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. *Applied soft computing* 90 (2020), 106181.
- [76] Zezhi Shao, Fei Wang, Yongjun Xu, Wei Wei, Chengqing Yu, Zhao Zhang, Di Yao, Guangyin Jin, Xin Cao, Gao Cong, et al. 2023. Exploring Progress in Multivariate Time Series Forecasting: Comprehensive Benchmarking and Heterogeneity Analysis. *arXiv preprint arXiv:2310.06119* (2023).
- [77] Chao Song, Youfang Lin, Shengnan Guo, and Huaiyu Wan. 2020. Spatial-temporal synchronous graph convolutional networks: A new framework for spatial-temporal network data forecasting. In *AAAI*, Vol. 34. 914–921.
- [78] A Suilin. 2017. *kaggle-web-traffic*. <https://github.com/Arturus/kaggle-web-traffic>
- [79] Chenchen Sun, Yan Ning, Derong Shen, and Tiezheng Nie. 2023. Graph Neural Network-Based Short-Term Load Forecasting with Temporal Convolution. *Data Science and Engineering* (2023), 1–20.
- [80] Souhaib Ben Taieb, Gianluca Bontempi, Amir F Atiya, and Antti Sorjamaa. 2012. A review and comparison of strategies for multi-step ahead time series forecasting based on the NN5 forecasting competition. *Expert systems with applications* 39, 8 (2012), 7067–7083.
- [81] Chang Wei Tan, Christoph Bergmeir, François Petitjean, and Geoffrey I. Webb. 2020. Monash University, UEA, UCR Time Series Regression Archive. *arXiv preprint arXiv:2006.10996* (2020).
- [82] Hiro Y Toda and Peter CB Phillips. 1994. Vector autoregression and causality: a theoretical overview and simulation study. *Econometric reviews* 13, 2 (1994), 259–285.
- [83] Luan Tran, Manh Nguyen, and Cyrus Shahabi. 2019. Representation learning for early sepsis prediction. In *2019 Computing in Cardiology (CinC)*. 1–4.
- [84] Artur Trindade. 2015. ElectricityLoadDiagrams20112014. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C58C86>.
- [85] Feng Wan, Linsen Li, Ke Wang, Lu Chen, Yunjun Gao, Weihao Jiang, and Shiliang Pu. 2022. MTPRE: a multi-scale spatial-temporal model for travel time prediction. In *SIGSPATIAL*. 1–10.
- [86] Huiqiang Wang, Jian Peng, Feihu Huang, Jince Wang, Junhui Chen, and Yifei Xiao. 2022. Micn: Multi-scale local and global context modeling for long-term series forecasting. In *ICLR*.
- [87] Jiaqi Wang, Tianyi Li, Anni Wang, Xiaozhe Liu, Lu Chen, Jie Chen, Jianye Liu, Junyang Wu, Feifei Li, and Yunjun Gao. 2023. Real-time Workload Pattern Analysis for Large-scale Cloud Databases. *arXiv preprint arXiv:2307.02626* (2023).
- [88] Kaimin Wei, Tianqi Li, Feiran Huang, Jinpeng Chen, and Zefan He. 2022. Cancer classification with data augmentation based on generative adversarial networks. *Frontiers of Computer Science* 16 (2022), 1–11.
- [89] Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. 2022. Timesnet: Temporal 2d-variation modeling for general time series analysis. *arXiv preprint arXiv:2210.02186* (2022).
- [90] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. 2021. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in Neural Information Processing Systems* 34 (2021), 22419–22430.
- [91] Xinle Wu, Dalin Zhang, Chenjuan Guo, Chaoyang He, Bin Yang, and Christian S Jensen. 2021. AutoCTS: Automated correlated time series forecasting. *Proc. VLDB Endow.* 15, 4 (2021), 971–983.
- [92] Xinle Wu, Dalin Zhang, Miao Zhang, Chenjuan Guo, Bin Yang, and Christian S Jensen. 2023. AutoCTS+: Joint Neural Architecture and Hyperparameter Search for Correlated Time Series Forecasting. *Proceedings of the ACM on Management of Data* 1, 1 (2023), 1–26.
- [93] Ronghui Xu, Meng Chen, Yongshun Gong, Yang Liu, Xiaohui Yu, and Liqiang Nie. 2023. TME: Tree-guided Multi-task Embedding Learning towards Semantic Venue Annotation. *ACM Transactions on Information Systems* 41, 4 (2023), 1–24.
- [94] Sean Bin Yang, Chenjuan Guo, Jilin Hu, Jian Tang, and Bin Yang. 2021. Unsupervised Path Representation Learning with Curriculum Negative Sampling. In *IJCAL*. 3286–3292.
- [95] Sean Bin Yang, Jilin Hu, Chenjuan Guo, Bin Yang, and Christian S Jensen. 2023. Lightpath: Lightweight and scalable path representation learning. In *SIGKDD*. 2999–3010.
- [96] Yuanyuan Yao, Dimeng Li, Hailiang Jie, Hailiang Jie, Tianyi Li, Jie Chen, Jiaqi Wang, Feifei Li, and Yunjun Gao. 2023. SimpleTS: An efficient and universal model selection framework for time series forecasting. *Proc. VLDB Endow.* 16, 12 (2023), 3741–3753.
- [97] Haomin Yu, Jilin Hu, Xinyuan Zhou, Chenjuan Guo, Bin Yang, and Qingyong Li. 2023. CGF: A Category Guidance Based PM_{2.5} Sequence Forecasting Training Framework. *IEEE Transactions on Knowledge and Data Engineering* (2023).
- [98] Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. 2023. Are transformers effective for time series forecasting?. In *AAAI*, Vol. 37. 11121–11128.
- [99] Lingyu Zhang, Wenjie Bian, Wenyi Qu, Liheng Tuo, and Yunhai Wang. 2021. Time series forecast of sales volume based on XGBoost. In *Journal of Physics: Conference Series*, Vol. 1873. 012067.
- [100] Shuyi Zhang, Bin Guo, Anlan Dong, Jing He, Ziping Xu, and Song Xi Chen. 2017. Cautionary tales on air-quality improvement in Beijing. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 473, 2205 (2017), 20170457.
- [101] Yunhao Zhang and Junchi Yan. 2022. Crossformer: Transformer utilizing cross-dimension dependency for multivariate time series forecasting. In *ICLR*.
- [102] Kai Zhao, Chenjuan Guo, Yunyao Cheng, Peng Han, Miao Zhang, and Bin Yang. 2023. Multiple time series forecasting with dynamic graph modeling. *Proc. VLDB Endow.* 17, 4 (2023), 753–765.
- [103] Yan Zhao, Xuanhao Chen, Liwei Deng, Tung Kieu, Chenjuan Guo, Bin Yang, Kai Zheng, and Christian S Jensen. 2022. Outlier detection for streaming task

- assignment in crowdsourcing. In *WWW*. 1933–1943.
- [104] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *AAAI*, Vol. 35. 11106–11115.
- [105] Tian Zhou, Ziqing Ma, Qingsong Wen, Liang Sun, Tao Yao, Wotao Yin, Rong Jin, et al. 2022. Film: Frequency improved legendre memory model for long-term time series forecasting. *Advances in Neural Information Processing Systems* 35 (2022), 12677–12690.
- [106] Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. 2022. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *ICML*. 27268–27286.