



Less is More: Efficient Time Series Dataset Condensation via Two-fold Modal Matching

Hao Miao

Aalborg University, Denmark
haom@cs.aau.dk

Ziqiao Liu

University of Electronic Science and
Technology of China, China
liuziqiao@std.uestc.edu.cn

Yan Zhao

Aalborg University, Denmark ✉
yanz@cs.aau.dk

Chenjuan Guo

Bin Yang
East China Normal
University, China ✉
{cjguo,byang}@dase.ecnu.edu.cn

Kai Zheng

University of Electronic Science and
Technology of China, China
zhengkai@uestc.edu.cn

Christian S. Jensen

Aalborg University, Denmark
csj@cs.aau.dk

ABSTRACT

The expanding instrumentation of processes throughout society with sensors yields a proliferation of time series data that may in turn enable important applications, e.g., related to transportation infrastructures or power grids. Machine-learning based methods are increasingly being used to extract value from such data. We provide means of reducing the resulting considerable computational and data storage costs. We achieve this by providing means of condensing large time series datasets such that models trained on the condensed data achieve performance comparable to those trained on the original, large data. Specifically, we propose a time series dataset condensation framework, TimeDC, that employs two-fold modal matching, encompassing frequency matching and training trajectory matching. Thus, TimeDC performs time series feature extraction and decomposition-driven frequency matching to preserve complex temporal dependencies in the reduced time series. Further, TimeDC employs curriculum training trajectory matching to ensure effective and generalized time series dataset condensation. To avoid memory overflow and to reduce the cost of dataset condensation, the framework includes an expert buffer storing pre-computed expert trajectories. Extensive experiments on real data offer insight into the effectiveness and efficiency of the proposed solutions.

PVLDB Reference Format:

Hao Miao, Ziqiao Liu, Yan Zhao, Chenjuan Guo, Bin Yang, Kai Zheng, and Christian S. Jensen. Less is More: Efficient Time Series Dataset Condensation via Two-fold Modal Matching. PVLDB, 18(2): 226 - 238, 2024. doi:10.14778/3705829.3705841

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/uestc-liuzq/STdistillation>.

Yan Zhao and Chenjuan Guo are corresponding authors.

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. 18, No. 2 ISSN 2150-8097. doi:10.14778/3705829.3705841

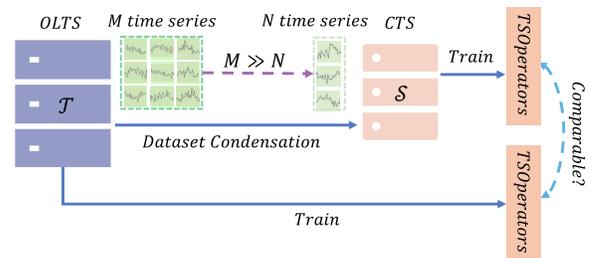


Figure 1: Time Series Condensation

1 INTRODUCTION

With the proliferation of edge computing and mobile sensing, massive volumes of time series data, comprising millions of observations, are being collected and stored into time series database systems (TSMs) [18, 21, 36], enabling various real-world applications [44, 46]. We are seeing impressive advances in machine learning [24, 51], especially in deep learning [10, 44, 54], that are successful at extracting information and creating value from large time series datasets. However, the proposed methods are also resource-intensive. Thus, storing and preprocessing datasets is costly, and model training often calls for the use of specialized equipment and infrastructure [8, 22, 47, 48], limiting the application on edge devices [36, 45].

An effective way to reduce costs associated with the use of large data is coreset construction [1, 7, 17, 22], which often involves clustering [15], Gaussian mixture models [25], and streaming learning [22]. Coreset construction methods [15, 25] often define heuristic criteria to select the most representative subsets from full datasets to form small coresets, such that models trained on the coresets are competitive with those built on the full datasets. Unfortunately, heuristic coreset construction methods guarantee neither optimal solutions nor promise the presence of representative observations in relation to downstream tasks [48]. Given these limitations, a recent alternative approach, dataset condensation [49] aims to directly synthesize small condensed, optimized datasets, not relying on representative subset-selection.

We consider the novel problem of time series dataset condensation, where the goal is to synthesize a small but informative condensed time series dataset \mathcal{S} derived from an original large dataset

\mathcal{T} . A model, composed of stacked time series operators (*TSOperators*) that perform feature extraction, is trained on the condensed time series (CTS) dataset with the objective of achieving performance comparable to that of a model trained on the original large time series (OLTS) dataset for downstream tasks (see Figure 1). The stacked *TSOperators* are multiple *TSOperators* arranged sequentially so that the output of one *TSOperator* is fed into the next *TSOperator*. Stacked *TSOperators* enable more complex feature extraction. Time series are usually stored in TSMs over edge devices [18, 45]. In the setting of continuously produced time series, time series dataset condensation aims to alleviate the storage burden of TSMs considering the limited processing capability of edge devices, which further contributes to the *database community* [2, 3]. Additionally, due to the small-scaled CTS, time series dataset condensation is expected to bring significant efficiency improvements in repeated training scenarios in TSMs, e.g., streaming learning [27] and model selection [42]. Although substantial research has been devoted to inventing effective dataset condensation methods [6, 48, 49], such as gradient matching [48], distribution matching [49], and multi-step parameters matching [6], existing methods target image data and cannot be applied to time series data directly as they are not built to contend with the unique temporal dependencies present in time series [26, 54]. These methods thus fail to capture inherent temporal dependencies such as trends and seasonalities [40]. New methods are needed to enable time series condensation. However, developing such methods is non-trivial, due to the following challenges.

Challenge I: Effectiveness and Generalization Ability. It is challenging to guarantee the effectiveness and generalization ability of the condensed time series dataset [2, 6]. First, typical dataset condensation methods require a bi-level (even triple-level) optimization to jointly learn a minimum of two objectives: model parameters and a condensed time series dataset. Such complex non-convex optimization cannot guarantee optimal solutions, thus significantly limiting its effectiveness as the representative of the large original time series dataset. Second, the condensed time series datasets should be generalized to train different networks. Nonetheless, the condensed data may suffer from various types of overfitting [6], e.g., overfitting to a certain network architecture. Also, the downstream models might overfit the condensed data during training.

Challenge II: Complex Temporal Dependencies. Existing popular dataset condensation methods mainly focus on computer vision without specific modules to capture the unique time series characteristics [4, 44]. Existing dataset condensation methods are ill-equipped for modeling the complex temporal dependencies of time series. For example, traffic conditions during afternoon rush hour may be similar on consecutive workdays. Moreover, morning rush hours may gradually start later when winter arrives as people get up later and later due to the gradual decrease in temperature and later sunrise. It is important to learn complex temporal dependencies in a time series dataset so that a condensed version of the dataset exhibits similar temporal patterns to the original dataset. In addition, time series datasets are often multivariate, encompassing correlated variables, or features (channels). Such features are often coupled to perform feature extraction simultaneously. We argue that analysis in independent channels enables more effective feature extraction. For example, traffic flow is negatively correlated with

vehicle speed, but positively correlated with road occupancy [38]. Coupling features with varying and complex interactions might confuse a model, leading to decreased model performance [38].

Challenge III: Scalability. Existing dataset condensation methods often suffer from poor scalability, as bi-level optimization is generally time-consuming [55]. For example, gradient-matching methods [48] set different hyper-parameters of the outer and the inner loop optimization for different learning settings, which requires extensive cross-validation, rendering condensed dataset synthesis costly. In addition, these methods may incur memory overflow: they require an entire model to reside in memory during training, which is infeasible in many resource-constrained environments [33]. It is highly desirable, but also non-trivial, to develop an efficient time series dataset condensation method.

This study addresses the above challenges by providing an efficient time series dataset condensation (TimeDC) framework, which features a two-fold matching mechanism: time series frequency matching and training trajectory matching. TimeDC synthesizes a small time series dataset that summarizes a large dataset such that models trained on the small dataset achieve comparable performance to models trained on the large dataset. It encompasses three major modules: a time series feature extraction module, a decomposition-driven frequency matching module, and a curriculum training trajectory matching module. Time series of stock prices are among the most common time series in real-world applications. We thus use such time series in a running example to explain and convey the intuition underlying definitions and the proposed modules throughout the paper.

To achieve effective and generalized TS dataset condensation (*Challenge I*), we propose a curriculum training trajectory matching (CT²M) module (see Figure 4), where we consider an original time series dataset as the gold standard and seek to imitate the long-term training dynamics of models trained on it. The parameter trajectories trained on the original dataset are called expert trajectories. CT²M first trains a set of expert trajectories of a model, composed of a set of stacked *TSOperators*, on the large dataset. These are computed offline and are stored in an expert buffer and then serve as guidance for the condensed dataset optimization. Next, we sample trajectories from the expert buffer with curriculum trajectory queries and conduct long-term trajectory matching to align the segments of trajectories between the condensed time series dataset and the large dataset, enabling comprehensive and general knowledge transfer to the condensed dataset.

To support the capture of complex temporal dependencies (*Challenge II*), we propose a time series feature extraction (TSFE) module. In particular, TSFE contains a channel-independent mechanism and a set of stacked *TSOperators*, where a *TSOperator* includes a set of self-attention and fully connected layers. A decomposition block separates and refines the frequency, i.e., trend and seasonality, progressively from the intermediate results extracted by each *TSOperator*. We design a decomposition-driven frequency matching (DDFM) module to facilitate consistent temporal pattern preservation between the condensed time series dataset and the original one. It maps the decomposed frequencies when optimizing the condensed dataset.

To address the issue of high computation costs and poor scalability (*Challenge III*), we propose an expert buffer to save expert

trajectories pre-computed on original time series data to avoid memory overflow and reduce the overall training time. Further, we employ a patching mechanism that splits the time series into patches, thereby accelerating time series feature extraction in the TSFE module and also enabling effective local semantics modeling for time series.

The main contributions are summarized as follows.

- To the best of our knowledge, this is the first study to learn dataset condensation over time series. We propose a framework called TimeDC that aims to condense large time series datasets into small synthetic time series datasets while retaining the expressiveness of the large datasets.
- We design a TSFE module to capture temporal dependencies of time series effectively. We further propose a DDFM module to reduce the discrepancy of frequencies between original and condensed time series.
- A novel curriculum trajectory query and matching module is proposed to penalize condensed time series data based on how far the synthetically trained model deviates from expert trajectories.
- We report on experiments using real data, offering evidence of the effectiveness and efficiency of the proposals.

The remainder of this paper is organized as follows. Section 2 covers preliminary concepts and formalizes the problem of time series dataset condensation. We detail the TimeDC framework in Section 3, followed by the experimental study in Section 4. Section 5 surveys related work, and Section 6 concludes the paper.

2 PRELIMINARIES

We proceed to present the necessary preliminaries and then define the problem addressed.

Definition 2.1 (Time Series). A time series $T = \langle t_1, t_2, \dots, t_n \rangle$ is a time ordered sequence of n observations, where each observation $t_i \in \mathbb{R}^C$ is a C -dimensional vector. If $C = 1$, T is univariate; otherwise, T is multivariate. For example, a series of stock prices with multiple features, e.g., the opening price, closing price, highest price, and lowest price, is an example of multivariate time series.

Definition 2.2 (Time Series Dataset). A time series dataset \mathcal{T} is a set of time series $\mathcal{T} = \{T_1, T_2, \dots, T_M\}$, where M is the cardinality.

2.1 Dataset Condensation over Time Series

Definition 2.3 (Condensed Time Series Dataset). Given a time series dataset \mathcal{T} , a time series dataset $\mathcal{S} = \{\tilde{T}_1, \tilde{T}_2, \dots, \tilde{T}_N\}$ is condensed if N is much smaller than M , where \tilde{T} is the condensed time series, M is the cardinality, and \mathcal{S} is derived from \mathcal{T} .

Note that we can initialize each item in \mathcal{S} with a random time series in the original dataset or a Gaussian noise. After dataset condensation, we expect to learn an optimized \mathcal{S} to replace \mathcal{T} .

Given a large time series dataset $\mathcal{T} = \{T_1, T_2, \dots, T_M\}$, we aim to learn an optimal differentiable function f parameterized by θ that correctly performs time series tasks, e.g., time series forecasting. We use \mathcal{T} as a training set to learn the optimal parameters $\theta^{\mathcal{T}}$ by minimizing an empirical loss term as follows.

$$\theta^{\mathcal{T}} = \arg \min_{\theta} \mathcal{L}^{\mathcal{T}}(\theta), \quad (1)$$

where $\mathcal{L}^{\mathcal{T}}(\theta) = \frac{1}{M} \sum_{T_i \in \mathcal{T}} \ell(f_{\theta}, T_i)$, and $\ell(\cdot)$ is a task specific loss (e.g., Mean Squared Error). We denote the generalization performance of the obtained model $f_{\theta^{\mathcal{T}}}$ by $\mathbb{E}_{T_i \sim P_{\mathcal{T}}}[\ell(f_{\theta^{\mathcal{T}}}, T_i)]$, where $P_{\mathcal{T}}$ is the data distribution of \mathcal{T} .

Dataset Condensation over Time Series. Given a large time series dataset \mathcal{T} , our goal is to synthesize a small time series dataset $\mathcal{S} = \{\tilde{T}_1, \tilde{T}_2, \dots, \tilde{T}_N\}$ that preserves most of the information in \mathcal{T} and $N \ll M$. Once the condensed time series set \mathcal{S} is learned, we can train f on it by replacing \mathcal{T} by \mathcal{S} in Equation 1. As the condensed set \mathcal{S} is significantly smaller than \mathcal{T} , we expect the optimization to be considerably faster. In addition, we expect the generalization performance of $f_{\theta^{\mathcal{S}}}$ to be close to $f_{\theta^{\mathcal{T}}}$, i.e., $\mathbb{E}_{T_i \sim P_{\mathcal{T}}}[\ell(f_{\theta^{\mathcal{T}}}, T_i)] \approx \mathbb{E}_{T_i \sim P_{\mathcal{S}}}[\ell(f_{\theta^{\mathcal{S}}}, T_i)]$, on the real data distribution $P_{\mathcal{T}}$.

Thus, given a learnable function f parameterized by θ , our problem, *dataset condensation over time series*, can be formulated as a bi-level optimization problem [48] as follows.

$$\underbrace{\min_{\mathcal{S}} \mathcal{L}(f_{\theta^{\mathcal{S}}}, \mathcal{T})}_{\text{outer level}}, \text{ s.t. } \theta^{\mathcal{S}} = \underbrace{\arg \min_{\theta} \mathcal{L}(f_{\theta}, \mathcal{S})}_{\text{inner level}}, \quad (2)$$

where the outer loop optimizes the condensed time series set \mathcal{S} , while the inner loop aims to learn an optimal function $f_{\theta^{\mathcal{S}}}$ on \mathcal{S} .

Offline Expert Trajectories. We denote the time sequence of parameters $\{\theta_e^k\}_{e=1}^E$ as expert trajectories, that are obtained during the training of the proposed TSFE module on the full, original time series dataset, where E is the training epoch intervals. To generate these expert trajectories, we train TSFE several times on original datasets and save their snapshot parameters at every epoch. We call these parameter sequences as expert trajectories since they represent the theoretical upper bound for the downstream applications.

3 METHODOLOGY

We proceed to detail the efficient time series dataset condensation framework, TimeDC. We first give an overview of the framework and then provide specifics on each module in the framework.

3.1 Framework Overview

As illustrated in Figure 2, given a model composed of stacked *T*SOperators, we first pre-train a set of training trajectories on the original large time series (OLTS) dataset and store these pre-trained trajectories as expert trajectories in an expert buffer. As the original dataset is used to guide the network training, we denote the parameters trained on the original dataset as expert trajectory. Then, we train the same model on the condensed time series (CTS) dataset and perform two-fold modal matching: trajectory matching and frequency matching. Specifically, we optimize the condensed dataset with respect to the distance between the synthetically training trajectories and the trajectories trained on the original dataset. Further, we align the frequencies (i.e., trend and seasonality) between the original dataset and the condensed dataset to preserve temporal correlations when synthesizing the condensed dataset. The framework encompasses three major modules: time series feature extraction (TSFE), decomposition-driven frequency matching (DDFM), and curriculum training trajectory matching (CT²M).

- *Time Series Feature Extraction.* This module aims to extract effective high-dimensional features from the input time series.

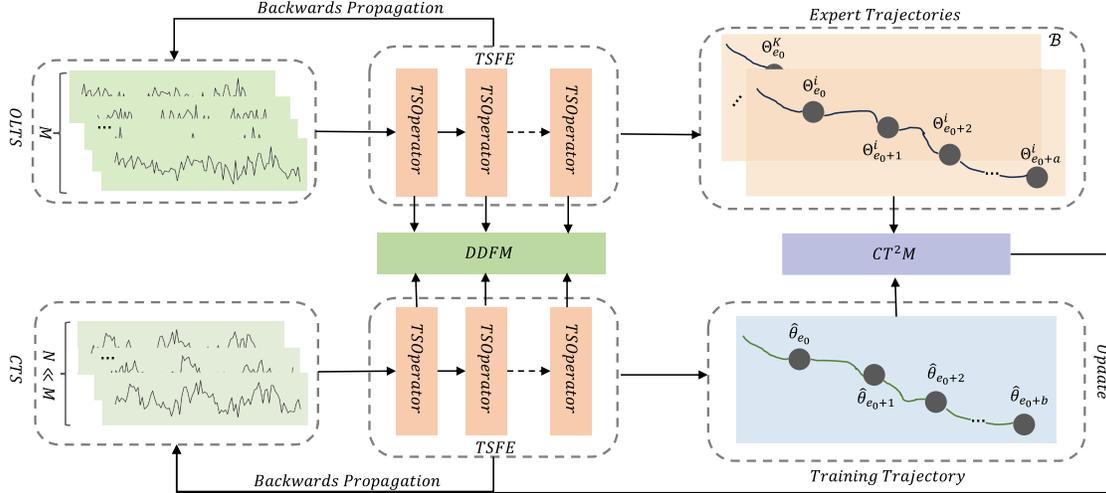


Figure 2: TimeDC Framework Overview

A channel-independent mechanism is adopted to decouple the time series $T_{input} \in \mathbb{R}^{\mathcal{B} \times n \times C}$ ($C \geq 1$) into C univariate time series $\{T_{input}^c\}_{c=1}^C \in \mathbb{R}^{\mathcal{B} \times n \times 1}$ for subsequent feature extraction. Next, the module splits each decoupled univariate (i.e., channel-independent) time series into P patches to learn local semantic information employing the patching mechanism. Then, the module stacks several $TSOperators$, composed of self-attention and fully connected layers, for channel-independent feature extraction. Finally, the channel-independent features are concatenated to enable prediction.

- **Decomposition-Driven Frequency Matching.** This module adopts a decomposition-driven frequency matching mechanism to reduce the discrepancy of decomposed frequencies between the original time series dataset and the condensed time series dataset, aiming to preserve temporal dependencies from the original dataset when synthesizing the condensed dataset. Specifically, it incorporates a moving-average based [12, 40] series decomposition block into each $TSOperator$, which separates the trend and seasonal information progressively from the learned features. This decomposition block enables a $TSOperator$ to alternately decompose and refine the intermediate results during the time series feature extraction.
- **Curriculum Training Trajectory Matching.** This module first trains a set of training trajectories (i.e., parameters) of stacked $TSOperators$ on the large original time series dataset, which are stored in an expert buffer. In addition, we design a curriculum trajectory query and matching mechanism to organize training trajectories matching from similar (easy) to dissimilar (hard). In particular, we sample trajectories from the expert buffer with minimally modified parameters to those with highly modified parameters, iteratively.

We train TimeDC on the original dataset to extract frequencies and compute original parameters denoted as expert trajectories, aiming at guiding the effective condensed dataset synthesis. Thus, extracting frequencies and expert trajectories from the original dataset supports the purpose of TimeDC, which enables models trained on the small dataset to achieve *comparable performance* to models trained on the large dataset. Next, we provide the technical details of each module.

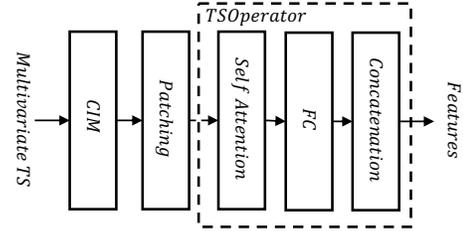


Figure 3: Feature Extraction Module

3.2 Time Series Feature Extraction

We first consider the time series feature extraction (TSFE) module. As shown in Figure 3, TSFE is composed of a channel-independent mechanism, a patching mechanism, and stacked $TSOperators$.

Channel Independent Mechanism. Most existing time series modeling methods [41, 44] perform time series feature extraction by coupling all features and projecting them together into an embedding space. However, simply coupling different features may influence model performance very uncontrollably [38]. To avoid this problem, we employ the channel-independent mechanism $CIM(\cdot)$ to model each time series feature independently.

We consider a time series with C features as a sequence of multi-channel variables. Specifically, given a time series $T_{input} = \langle t_1, t_2, \dots, t_n \rangle \in \mathbb{R}^{\mathcal{B} \times n \times C}$ ($C \geq 1$), the channel independent mechanism separates T_{input} into C univariate time series along the feature (i.e., channel) dimension, where \mathcal{B} is the batch size. For example, in a series of stock prices with multiple features, e.g., opening prices and closing prices, the channel-independent mechanism separates these features by forming univariate time series. This means that the opening and closing prices are considered as two different time series in the subsequent feature extraction. Formally, the channel-independent mechanism works as follows.

$$CIM(T_{input}) = T_{input}^1, \dots, T_{input}^c, \dots, T_{input}^C, \quad (3)$$

where $T_{input}^c \in \mathbb{R}^{\mathcal{B} \times n \times 1}$. The separated univariate time series are then fed independently into the stacked $TSOperators$. The $TSOperators$ share the same architecture, but their forward processes are independent.

Patching Mechanism. Time series feature extraction aims to model the correlations between observations at different time steps. We aggregate observations over several time steps of the separated univariate time series T_{input}^c into subseries-level patches [29] to enable local semantics modeling, which allows the model to see the longer historical sequences to improve feature extraction. For instance, in the context of stock prices, if the original time series represents hourly prices over a week, the patching mechanism may create patches where each patch contains the stock prices for each day, enhancing the capture of local semantics capturing.

We feed the univariate time series obtained from the channel independent mechanism into the patching mechanism that then divides them into patches that can be set to be overlapping or non-overlapping. Given a patch length L and a stride (i.e., the non-overlapping region between two consecutive patches) S , a sequence of patches $T_p^c \in \mathbb{R}^{\mathcal{B} \times L \times \mathcal{P}}$ is generated, where $\mathcal{P} = \lfloor \frac{n-L}{S} + 2 \rfloor$ is the number of patches. The last value t_n repeated $S - 1$ times is padded to the end of the input time series before patching.

The patching decreases the memory usage and computational complexity of the following $TSOperator$, containing self-attention and fully connected (FC) layers quadratically, by a factor of S since the input time series length reduced from n to approximately n/S .

$TSOperator$. We then feed the patches T_p^c into the stacked $TSOperators$ to learn latent features for the following frequency matching and trajectory matching. As illustrated in Figure 3, a $TSOperator$ includes a multi-head self-attention layer followed by a fully-connected (FC) layer. In the context of stock prices, the multi-head self-attention mechanism learns the temporal correlations between different time steps. This may include how the opening prices at the beginning of a day influence the prices in the next few days. The FC layer summarizes these temporal correlations for future prediction. The j -th $TSOperator$ can be formulated as follows.

$$h^j = TSOperator(h^{j-1}) = Norm(FC(MultiHead(h^{j-1}))), \quad (4)$$

where $h^0 = T_p^c$, and $Norm(\cdot)$ is the normalization layer (i.e., Batch-Norm). The multi-head self-attention layer has three components: *Query*, *Key*, and *Value*. The multi-head self-attention is computed by concatenating the output matrix of each attention head. For i -th attention head, the input latent features h^j are transformed as follows.

$$\begin{aligned} A_i^j &= Attention(h^j W_i^Q, h^j W_i^K, h^j W_i^V) \\ Attention(Q, K, V) &= softmax\left(\frac{Q \cdot K^T}{\sqrt{d_k}}\right) \cdot V \\ h^{j+1} &= Norm(h^j + FC_i(A_i^j h^j)), \end{aligned} \quad (5)$$

where d_k is a scaling factor, and W_i^Q , W_i^K , and W_i^V are the linear transformation parameter of the *Query*, *Key*, and *Value*, respectively.

The time series feature extraction is shown in Algorithm 1. Specifically, we first separate the batch of input time series T_{input} into C univariate time series (line 1) and apply the patching mechanism to split the univariate time series into patches (line 2). Then, we input the patches obtained in line 2 into the stacked $TSOperators$ for feature extraction (lines 3–8). Finally, C extracted features are concatenated to get the overall hidden features h (line 9). The space and time complexities of the TSFE module are both $O(\frac{n^2}{S^2})$.

Algorithm 1: Time Series Feature Extraction

Input: a batch of time series: $T_{input} \in \mathcal{B} \times n \times C$; the number of $TSOperators$: N_{op} .
Output: extracted features: h .

- 1 $T_{input}^1, \dots, T_{input}^c, \dots, T_{input}^C \leftarrow$ Separate T_{input} into C univariate channel-independent time series with Equation 3;
- 2 Split each univariate time series into patches;
- 3 **for** $1 \leq c \leq C$ **do**
- 4 $h_c^0 \leftarrow T_{input}^c$;
- 5 **for** $0 < j < N_{op}$ **do**
- 6 Feature extraction with Equation 4;
- 7 $h_c^j \leftarrow TSOperator(h_c^{j-1})$;
- 8 $h_c \leftarrow h_c^{N_{op}-1}$;
- 9 $h \leftarrow$ Concatenate $\{h_c\}_{c=1}^C$ along the channel;
- 10 **return** h .

3.3 Decomposition-Driven Frequency Matching

Time series data usually exhibit distinct frequencies, such as trend and seasonality, separating from other modalities (e.g., images). For instance, the trend in stock prices indicates the prolonged direction of prices over time, which can be upward or downward. Seasonality refers to the regular, periodic fluctuations in stock prices, often influenced by seasonal business cycles or holidays. We argue that it would give more guidance on time series condensation if the condensed and original time series share similar frequencies. Thus, to preserve temporal patterns in the condensed time series dataset, we propose a novel decomposition-driven frequency matching (DDFM) mechanism that incorporates a decomposition block into the stacked $TSOperators$.

As shown in Figure 2, for each $TSOperator$, DDFM decomposes the intermediate hidden features, learned from the original dataset \mathcal{T} and the condensed dataset \mathcal{S} , into their frequencies (i.e., trend and seasonality) progressively. It then aligns the frequencies to reduce the discrepancy of the temporal patterns between the original dataset and the condensed dataset based on the moving average. Concretely, for the hidden features h^j of the j -th $TSOperator$, we adopt the moving average by means of average pooling $AvgPool(\cdot)$ to smooth periodic fluctuations and highlight the trend and seasonality. In particular, the process of series decomposition is as follows.

$$h_{TRE}^j = AvgPool(Padding(h^j)), \quad h_{SEA}^j = h^j - h_{TRE}^j \quad (6)$$

where h_{TRE}^j and h_{SEA}^j are the decomposed trend and seasonality, respectively. We employ the average pooling $AvgPool(\cdot)$ with the padding operation $Padding(\cdot)$ to keep the time series length unchanged. We denote the decomposition process as follows.

$$h_{TRE}^j, h_{SEA}^j = SeriesDecomposition(h^j) \quad (7)$$

Given the extracted features $h_{\mathcal{T}}^j$ and $h_{\mathcal{S}}^j$ of the original dataset \mathcal{T} and the condensed dataset \mathcal{S} , we decompose the corresponding trend and seasonality according to Equation 7. We use cosine similarity $cos(\cdot, \cdot)$ to measure similarities between the trends h_{TRE}^j ,

$h_{TRE_S}^j$ and seasonalities $h_{SEA_{\mathcal{T}}}^j, h_{SEA_S}^j$, formulated as follows.

$$\begin{aligned} \cos(h_{TRE_{\mathcal{T}}}^j, h_{TRE_S}^j) &= \frac{h_{TRE_{\mathcal{T}}}^j}{\|h_{TRE_{\mathcal{T}}}^j\|_2} \cdot \frac{h_{TRE_S}^j}{\|h_{TRE_S}^j\|_2} \\ \cos(h_{SEA_{\mathcal{T}}}^j, h_{SEA_S}^j) &= \frac{h_{SEA_{\mathcal{T}}}^j}{\|h_{SEA_{\mathcal{T}}}^j\|_2} \cdot \frac{h_{SEA_S}^j}{\|h_{SEA_S}^j\|_2}, \end{aligned} \quad (8)$$

where $\|\cdot\|_2$ is l_2 -norm.

We aim to maximize the cosine similarity to reduce the discrepancy of frequencies between the condensed data and the original data. The objective function of DDFM is formulated as follows.

$$L_{Fre} = -\frac{1}{N_{op}} \sum_j^{N_{op}} (\cos(h_{TRE_{\mathcal{T}}}^j, h_{TRE_S}^j) + \cos(h_{SEA_{\mathcal{T}}}^j, h_{SEA_S}^j)), \quad (9)$$

where N_{op} is the number of *TSOperators*.

3.4 Curriculum Training Trajectory Matching

To enable effective dataset condensation, we propose a curriculum training trajectory matching mechanism. Existing popular dataset condensation methods, which are often based on gradient matching [48], conduct online gradient matching step by step. These methods have drawbacks when used for short-range matching, causing short-sightedness issues, thus failing to imitate holistic learning behaviors, and reducing the quality of condensed time series data.

Motivated by an existing study [6], we match the long-term training trajectories of the TTFM module with the offline guidance of expert trajectories, which are pre-trained and stored in the expert buffer. We consider the expert trajectories (i.e., the performance of stacked *TSOperators* trained on the original dataset \mathcal{T}) as the theoretical upper bound for the downstream application tasks. Our goal is to obtain a condensed dataset \mathcal{S} that will induce a similar trajectory as that induced by the real training data \mathcal{T} such that models trained on \mathcal{T} and \mathcal{S} achieve similar performance. In addition, we propose a curriculum training trajectory query to perform trajectory matching from similar to dissimilar expert trajectories, which further accelerates the model training. For example, expert trajectories derived from historical stock prices will then guide the alignment of training trajectories extracted from a condensed dataset. This alignment ensures that the condensed dataset's training trajectories capture similar dynamics (such as relationships across different time steps) to those learned from historical data.

In particular, we first train K time series models with the same architecture, i.e., stacked *TSOperators*, denoted as $f_{\mathcal{T}}$, on the original large time series dataset. Then, we can obtain K numbers of expert training trajectories that have a holistic knowledge of the original time series dataset in terms of $f_{\mathcal{T}}$'s training process. We save these model parameters $\{\Theta_{\mathcal{T}}^k\}_{k=1}^K = \{\theta_e^k\}_{e=1}^E$ at certain epoch intervals E in the expert buffer \mathcal{B} . Finally, we design a curriculum trajectory query mechanism to sample training trajectories of the expert buffer from similar to dissimilar for trajectory matching. Here, the pre-training of $f_{\mathcal{T}}$ is offline and can be separated from the end-to-end time series condensation, reducing the online computation costs.

Training Trajectory Matching. We use the original dataset to guide the network training, and the parameters trained on this

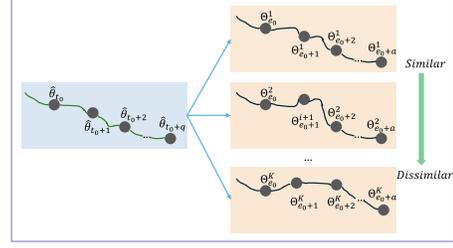


Figure 4: Curriculum Trajectory Query and Matching

dataset are called an expert trajectory. If a condensed time series dataset is capable of forcing network training dynamics to follow expert trajectories, the idea is that the synthetically trained network will be located close to the model trained on the original time series dataset and will achieve similar test performance for downstream applications.

When sampling a pre-computed parameter trajectory θ_e from the expert buffer \mathcal{B} , we aim to minimize the distance between the parameters $\tilde{\theta}_e$ trained on the condensed dataset and θ_e :

$$\arg \min_{\mathcal{S}} \mathbb{E}_{\theta_e \sim \mathcal{B}} [L_{tmm}(\theta_e|_{e=e_0}^a, \tilde{\theta}_e|_{e=e_0}^b)], \quad (10)$$

where \mathbb{E} is the expectation, $\theta_e|_{e=e_0}^a, \tilde{\theta}_e|_{e=e_0}^b$ are the parameters with range $(e_0, e_0 + a)$ and $(e_0, e_0 + b)$, where $e_0 < e_0 + a < E$. More specifically, we update the condensed TS data according to the trajectory matching loss L_{tmm} .

$$L_{tmm} = \frac{\|\tilde{\theta}_{e_0+b} - \theta_{e_0+a}\|_2^2}{\|\tilde{\theta}_{e_0} - \theta_{e_0+a}\|_2^2} \quad (11)$$

Note that we initialize the parameters of f_{θ_S} with those of $f_{\mathcal{T}}$ at e_0 training step, i.e., $\tilde{\theta}_{e_0} = \theta_{e_0}$, for more focused training. This way, we align the learning behavior of $f_{\mathcal{T}}$ with a -steps optimization to f_{θ_S} with b -steps optimization, imitating the long-term learning behavior of time series modeling.

For the inner loop, we train f_{θ_S} on the condensed time series data for optimization until the optimal $\tilde{\theta}^*$. Thus, the objective function of the CT²M is defined as follows.

$$\min_{\mathcal{S}} \mathbb{E}_{\theta_e \sim \mathcal{B}} [L_{tmm}(\theta_e|_{e=e_0}^a, \tilde{\theta}_e|_{e=e_0}^b)] \quad s.t. \tilde{\theta}^*, \mathcal{S} = \arg \min_{\tilde{\theta}} \mathcal{L}(f_{\tilde{\theta}}, \mathcal{S}), \quad (12)$$

where $\mathcal{L}(\cdot)$ is the task-specific loss (e.g., Mean Square Error (MSE)).

Utilizing training trajectory matching, we can reduce computation and memory costs during the condensation process by sampling pre-trained expert trajectories offline. Moreover, long-term trajectory matching provides a more holistic and comprehensive method to imitate the learning behaviors over the original dataset by avoiding that the condensed dataset fits to certain optimization steps short-sightedly.

Curriculum Training Trajectory Query. As shown in Figure 4, we propose a curriculum training trajectory query method to further enhance the generalization of time series condensation and accelerate the model convergence [37, 52]. Curriculum learning aims to train a model with easy samples first, and then gradually increases the difficulty levels. In our setting, we first query similar training trajectories $\theta_e|_{e=e_0}^a$ and $\tilde{\theta}_e|_{e=e_0}^b$, learned on the original and condensed datasets, respectively, to perform trajectory matching.

Algorithm 2: Curriculum Training Trajectory Query and Matching

Input: A buffer \mathcal{B} with a set of trajectories pre-trained on the original TS dataset \mathcal{T} parameterized by $\{\Theta_{\mathcal{T}}^k\}_{k=1}^K$; current model parameters $\tilde{\theta}^S$ on \mathcal{S} .

Output: Trajectory matching loss L_{tmm} .

```

1 Distance list  $DT \leftarrow []$ ;
2 Pre-update  $\tilde{\theta}^S$  for  $a$ -steps with Equation 13;
3  $\tilde{\theta}_{e_0+a} \leftarrow \tilde{\theta}_{e_0} - \sum_{s=1}^a (\alpha \nabla \mathcal{L}(f_{\theta^S}, \mathcal{S}))$ ;
4 for  $\Theta_{\mathcal{T}}^k \in \{\Theta_{\mathcal{T}}^k\}_{k=1}^K$  do
5   Compute the distance  $dis_k$  between  $\tilde{\theta}_{e_0}^a$  and  $\Theta_{\mathcal{T}}^k$  with
   Equation 15;
6    $dis_k \leftarrow -D(\tilde{\theta}_{e=e_0}^a, \theta^k|_{e=e_0})$ ;
7    $DT \leftarrow (k, dis_k)$ ;
8 Rank  $DT$  in a descending order;
9  $\beta \leftarrow 0$ 
10 while  $\beta < K$  do
11    $k \leftarrow DT[\beta][0]$ 
12    $L_{tmm} \leftarrow$  Sample trajectory  $\Theta_{\mathcal{T}}^k$  and match the training
   trajectory according to Equation 11;
13    $\beta \leftarrow \beta + 1$ ;
14 return  $L_{tmm}$ 

```

Then, we gradually increase the dissimilarity of the original data training trajectory.

To quantify the similarity between training trajectories of the original dataset \mathcal{T} and the condensed dataset \mathcal{S} , we design a minimally interfered retrieval sampling strategy. Specifically, given current parameters $\tilde{\theta}_{e_0}$ of f_{θ^S} learned on the condensed dataset, trajectories $\{\Theta_{\mathcal{T}}^k\}_{k=1}^K$ in buffer \mathcal{B} , and a standard task-specific objective function $\min_{\theta^S} \mathcal{L}(f_{\theta^S}, \mathcal{S})$, we retrieve trajectories that will be close to $\tilde{\theta}$ by the update of the foreseen (i.e., future) parameters to select more similar trajectories from \mathcal{B} . We update the parameter $\tilde{\theta}_{e_0}$ for a -steps by gradient matching, as shown in Equation 13.

$$\tilde{\theta}_{e_0+a} = \tilde{\theta}_{e_0} - \sum_{s=1}^a (\alpha \nabla \mathcal{L}(f_{\theta^S}, \mathcal{S})), \quad (13)$$

where α is the learning rate. Then, we compute the distance $D(\cdot, \cdot)$ between the foreseen trajectory $\tilde{\theta}_{e=e_0}^a$ of \mathcal{S} and sub-trajectories $\theta^k|_{e=e_0}$ of $\{\Theta_{\mathcal{T}}^k\}_{k=1}^K$ in \mathcal{B} as follows.

$$dis_k = -D(\tilde{\theta}_{e=e_0}^a, \theta^k|_{e=e_0}) \quad (14)$$

We use the cosine similarity to measure the distance $D(\cdot)$.

$$D(\tilde{\theta}_{e=e_0}^a, \theta^k|_{e=e_0}) = \frac{\tilde{\theta}_{e=e_0}^a \cdot \theta^k|_{e=e_0}}{\|\tilde{\theta}_{e=e_0}^a\|_2 \cdot \|\theta^k|_{e=e_0}\|_2} \quad (15)$$

Finally, we match current model trajectories of \mathcal{S} with the pre-trained trajectories in descending order based on the similarity $\{dis_k\}_{k=1}^K$.

Algorithm 2 shows the process of the proposed curriculum training trajectory matching. Lines 1–8 concern the curriculum training trajectory query, and lines 9–13 concern the trajectory matching.

The space and time complexities of Algorithm 2 are $O(K \cdot n)$ and $O((a+K) \cdot n + K \log K)$, respectively.

3.5 Overall Objective Function

The final loss contains three parts: a task-specific loss \mathcal{L} , a frequency matching loss L_{Fre} , and a trajectory matching loss L_{tmm} . We combine them together and the overall loss is as follows.

$$L_{all} = \mathcal{L} + L_{Fre} + L_{tmm}. \quad (16)$$

The task-specific loss \mathcal{L} is specific to the particular downstream tasks to achieve a better condensed dataset tailored for the intended use. For example, we adopt the cross-entropy loss for task-specific optimization for time series classification tasks to optimize accuracy, while we adopt the mean square error (MSE) loss to minimize the discrepancy between ground truth and predicted values for time series forecasting tasks.

4 EXPERIMENTAL EVALUATION

4.1 Experimental Setup

4.1.1 Datasets. Various time series analytics are on edge devices where the storage is limited. We aim at getting a really small condensed dataset. The experiments are carried out on six widely-used time series datasets, covering four application domains: weather, traffic, economics, and energy.

- **Weather.** The Weather dataset contains 21 indicators of weather (e.g., air temperature and humidity), which are collected in Germany. The data is recorded every 10 minutes.
- **Traffic.** The Traffic dataset contains hourly road occupancy rates obtained from sensors located at San Francisco freeways from 2015 to 2016.
- **Electricity.** The Electricity dataset contains the hourly electricity consumption of 321 clients from 2012 to 2014.
- **ETT.** The ETT dataset includes two hourly-level datasets (ETTh1 and ETTh2) and two 15-minute-level datasets (ETTm1 and ETTm2). Each dataset includes 7 oil and load features of electricity transformers between July 2016 and July 2018.

We choose time series forecasting as a representative downstream task, as it is a popular analytics task. In Section 4.4, we also present the performance comparison on the task of time series classification. The numbers of condensed time series are set to 500 and 50 as default for forecasting and classification tasks, respectively. We employ the proposed stacked *TOperators* as the forecasting and classification models.

4.1.2 Baselines. We compare TimeDC with the following existing methods that include coresets construction methods (i.e., Random, Herding [39], and K-Center [14]), and dataset condensation methods (i.e., DC [48] and MTT [6]).

- **Random.** The Random method randomly selects certain numbers of time series as a coreset.
- **Herding.** The Herding method adds time series observations to coresets greedily [39].
- **K-Center.** The K-Center method first performs K-Center clustering on the original datasets and then chooses observations from each cluster [14].

Table 1: Overall Performance Comparison on Seven Datasets

Baseline		Random		K-Center		Herding		DC		MTT		TimeDC		Whole Dataset	
Dataset	<i>PL</i>	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE
Weather	96	0.731	1.256	0.452	0.687	0.478	0.677	0.361	0.514	<u>0.295</u>	<u>0.244</u>	0.257	0.188	<i>0.239</i>	<i>0.182</i>
	192	0.786	1.302	0.487	0.723	0.512	0.688	0.413	0.527	<u>0.344</u>	<u>0.301</u>	0.285	0.247	<i>0.261</i>	<i>0.195</i>
	336	0.794	1.311	0.524	0.756	0.554	0.712	0.444	0.567	<u>0.368</u>	<u>0.328</u>	0.330	0.287	<i>0.282</i>	<i>0.241</i>
Traffic	96	0.675	1.125	0.503	0.576	0.483	0.554	0.375	0.603	<u>0.279</u>	<u>0.403</u>	0.254	0.375	<i>0.247</i>	<i>0.337</i>
	192	0.712	1.144	0.514	0.604	0.517	0.606	0.432	0.633	<u>0.336</u>	<u>0.442</u>	0.297	0.405	<i>0.265</i>	<i>0.338</i>
	336	0.729	1.117	0.523	0.611	0.553	0.654	0.449	0.676	<u>0.355</u>	<u>0.471</u>	0.312	0.423	<i>0.297</i>	<i>0.360</i>
Electricity	96	0.421	0.669	0.448	0.583	0.501	0.592	0.376	0.513	<u>0.296</u>	<u>0.283</u>	0.274	0.267	<i>0.252</i>	<i>0.268</i>
	192	0.450	0.743	0.476	0.601	0.534	0.628	0.419	0.532	<u>0.315</u>	<u>0.337</u>	0.285	0.294	<i>0.239</i>	<i>0.255</i>
	336	0.491	0.853	0.506	0.622	0.569	0.477	0.436	0.544	<u>0.339</u>	<u>0.356</u>	0.304	0.322	<i>0.271</i>	<i>0.285</i>
ETTh1	96	0.523	0.745	0.554	0.698	0.536	0.656	0.503	<u>0.442</u>	<u>0.456</u>	0.464	0.413	0.401	<i>0.354</i>	<i>0.386</i>
	192	0.557	0.786	0.578	0.722	0.589	0.698	0.552	0.508	<u>0.504</u>	<u>0.471</u>	0.436	0.428	<i>0.362</i>	<i>0.355</i>
	336	0.588	0.802	0.604	0.745	0.603	0.723	0.556	0.513	<u>0.498</u>	<u>0.464</u>	0.447	0.431	<i>0.409</i>	<i>0.387</i>
ETTh2	96	0.487	0.655	0.589	0.711	0.521	0.589	0.463	0.524	<u>0.388</u>	<u>0.342</u>	0.368	0.271	<i>0.324</i>	<i>0.255</i>
	192	0.509	0.673	0.605	0.732	0.553	0.621	0.488	0.536	<u>0.416</u>	<u>0.384</u>	0.389	0.302	<i>0.332</i>	<i>0.257</i>
	336	0.524	0.689	0.623	0.744	0.564	0.640	0.505	0.540	<u>0.435</u>	<u>0.455</u>	0.411	0.334	<i>0.376</i>	<i>0.296</i>
ETTh1	96	0.743	1.124	0.525	0.492	0.607	0.554	0.603	0.665	<u>0.512</u>	<u>0.453</u>	0.503	0.442	<i>0.453</i>	<i>0.403</i>
	192	0.764	1.245	0.566	0.510	0.628	0.571	0.597	0.647	<u>0.563</u>	<u>0.501</u>	0.512	0.465	<i>0.464</i>	<i>0.432</i>
	336	0.801	1.128	0.571	0.523	0.644	0.582	0.624	0.668	<u>0.552</u>	<u>0.488</u>	0.500	0.483	<i>0.477</i>	<i>0.455</i>
ETTh2	96	0.664	0.795	0.486	0.623	0.524	0.558	0.472	0.535	<u>0.376</u>	<u>0.421</u>	0.354	0.391	<i>0.347</i>	<i>0.381</i>
	192	0.687	0.804	0.512	0.643	0.549	0.583	0.488	0.567	<u>0.453</u>	<u>0.479</u>	0.401	0.421	<i>0.358</i>	<i>0.403</i>
	336	0.702	0.823	0.558	0.661	0.598	0.624	0.493	0.556	<u>0.473</u>	<u>0.523</u>	0.453	0.474	<i>0.406</i>	<i>0.435</i>

- **DC.** The DC method employs gradient matching to perform dataset condensation [48].
- **MTT.** The MTT method matches the multi-step training parameters of the condensed data and the original data [6].

4.1.3 *Evaluation Metrics.* Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) are adopted as the evaluation metrics, which are defined as follows.

$$MAE = \frac{1}{M} \sum_{m=1}^M |\hat{y}^m - y^m|, MSE = \frac{1}{M} \sum_{m=1}^M \|\hat{y}^m - y^m\|^2, \quad (17)$$

where M is the testing data size, \hat{y}^t is the prediction and y^t is the ground truth. The smaller the MAE and the RMSE are, the more accurate method is. We also evaluate the efficiency of the models, including the training and dynamic tensor memory cost.

4.1.4 *Implementation Details.* We implement our model using the Pytorch framework on an NVIDIA GTX 3090 GPU. The hyper-parameters in the model are set as follows. The patch length and stride are set to 16 and 8, respectively. The initial learning rate is 0.0001. The number of *TSOperator* layers is 3. The number of heads in the self-attention layer is set to 16. The number of expert trajectories in the expert buffer is set to 10 by default. ETT datasets and other datasets are split into the training data, validation data, and test data by the ratio of 6 : 2 : 2 and 7 : 1 : 2, respectively. The parameters of the baseline methods are set according to their original papers and any accompanying code. All of the models follow the same experimental setup with prediction length $PL \in \{96, 192, 336\}$ on all datasets.

4.2 Experimental Results

4.2.1 *Overall Performance Comparison.* We report the MAE and RMSE values of the methods in Table 1. The best performance by an existing method (Random, Herding, K-Center, DC, and MTT) is underlined, and the overall best performance is marked in bold. Whole Dataset indicates training on the original dataset and serves as an approximate upper-bound performance; is marked in italics. We use a set of stacked TSFE modules as the basic forecasting model for each baseline. The following observations are made.

- TimeDC achieves the best results on all datasets across all prediction lengths ($PL \in \{96, 192, 336\}$). TimeDC performs better than the best among the baselines by up to 13.49% and 26.59% in terms of MAE and RMSE, respectively. We observe that the performance improvements obtained by TimeDC on the Weather dataset exceed those on the Traffic, Electricity, ETTh1, and ETTh2 in most cases. This is because the Weather has much more training data than the other datasets. TimeDC trained with more training data results in better results. Though the ETTm1 and ETTm2 datasets have substantial training data, TimeDC performs slightly better than baselines on these two datasets. This is because the temporal patterns in ETTm1 and ETTm2 are highly regular and can be captured easily by existing methods.
- The coreset methods (i.e., Random, K-Center, and Herding), perform worse than the condensation methods (i.e., DC, MTT, and TimeDC), where Random has the worst performance in most cases. This is because coreset methods are based on heuristic metrics, that makes it hard to select representative observations and guarantee the optimal solutions.
- A popular condensation method, MTT performs the best among the existing methods except when $PL = 96$ on ETTh1, due to

its powerful condensation capability that involves matching multi-step parameters.

The experiment offers evidence that TimeDC is more effective than existing coreset and dataset condensation methods.

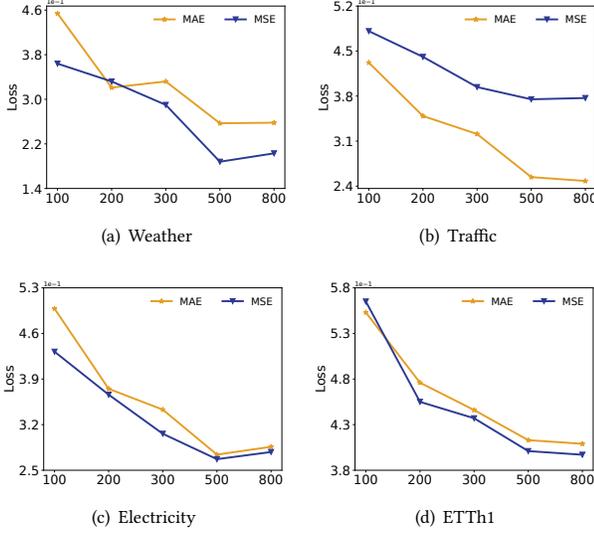


Figure 5: Effect of the Size of Condensed TS Dataset on Four Datasets ($PL = 96$)

4.2.2 Effect of the Size of Condensed Time Series Dataset. To study the effect of the size of a condensed time series dataset, we conduct experiments with 100, 200, 300, 500, and 800 condensed time series. The results are shown in Figure 5. We observe that the curves first drop significantly and then increase slightly (Figures 5(a) and 5(c)) or remain almost the same (Figures 5(b) and 5(d)). Generally, the results demonstrate that the model performance improves with an increase in the condensed time series data, as more condensed data yields more training data. In addition, it shows that using more condensed time series data for training is more likely to lead to better performance because more useful knowledge is learned from more data. One can also observe that TimeDC with 800 condensed time series performs slightly worse than TimeDC with 500 condensed time series data on the Weather and Electricity. This may be because the patterns in these datasets are relatively simple. Additional condensed time series data might introduce recurring patterns, thereby making the model overfit to these patterns and degrading performance on other data.

4.2.3 Ablation Study. To gain insight into the effects of the different components of TimeDC, including the patching mechanism (patch), decomposition-driven frequency matching (DDFM), and curriculum training trajectory matching (CT^2M), we evaluate three variants:

- **w/o_Patch.** TimeDC without the patching mechanism.
- **w/o_DDFM.** TimeDC without the DDFM module.
- **w/o_CT²M.** TimeDC without the CT^2M module.

Figure 6 shows results on Weather, Traffic, Electricity, and ETTh1. Regardless of the datasets, TimeDC outperforms its counterparts without the patching mechanism, the DDFM module, and the CT^2M module. This shows that these three components are all useful for

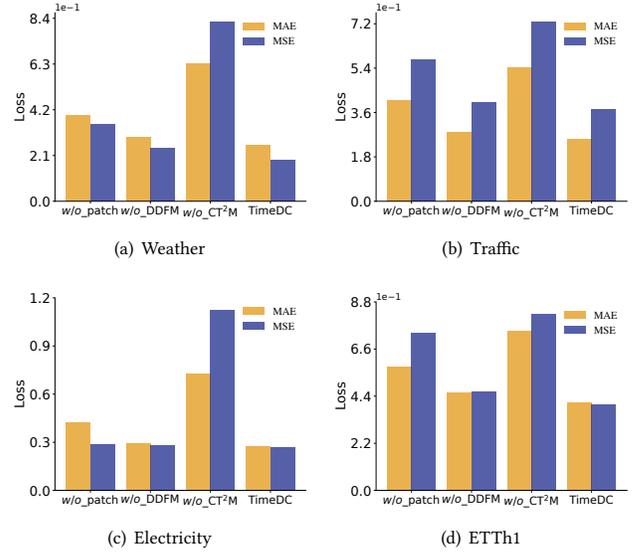


Figure 6: Performance of TimeDC and Its Variants on Four Datasets ($PL = 96$)

effective time series dataset condensation. TimeDC obtains MAE and RMSE reductions by up to 12.88% and 22.95%, respectively, compared with w/o_DDFM. Further, on all datasets, w/o_CT²M performs worst among all variants. TimeDC performs better than w/o_CT²M by at least 44.64% and 48.49% in terms of MAE and MSE, respectively, which indicates the effectiveness of the CT^2M module.

4.2.4 Cross-Architecture Performance. Next, we consider the cross-architecture performance of TimeDC. It is important to determine whether the condensed time series data generated by TimeDC can be used to train an unseen network. To assess such cross-architecture performance comprehensively, we consider three representative state-of-the-art network architectures for time series forecasting, including *Autoformer* [41], *Informer* [53], and *Transformer* [35]. We first synthesize condensed time series data with TimeDC and then train these networks with the condensed time series data. The hyper parameters of these network architectures are set based on their original papers and any accompanying code. For TimeDC, we use a set of stacked TSFE modules as the forecasting network.

The prediction results are given in Table 2. Overall, TimeDC has the best performance in most cases, indicating its stable and superior performance, especially on Traffic and Electricity. It also illustrates the effectiveness of the TSFE module. We observe that Autoformer performs worse than TimeDC but better than Informer and Transformer, while the performances of Informer and Transformer are comparable, especially on ETTh1 and ETTh2. These observations are in line with their performances when trained on the original dataset [41], indicating that TimeDC learns a generalized condensed time series dataset that works across different network architectures. This is because of the powerful feature extraction capabilities of the proposed stacked *TSOperators*, as well as the trajectory matching that imitates the long-term training dynamics of models trained on the original dataset. In addition, training a model on the condensed dataset consumes much less time than

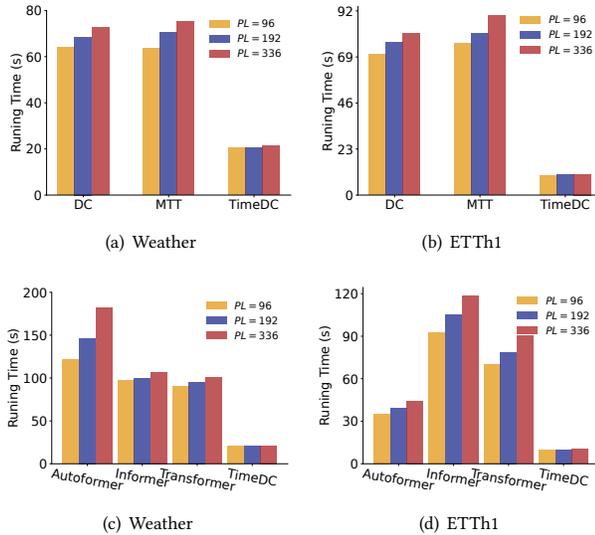
Table 2: Cross-Architecture Performance Comparison

Method	Metric	PL	Weather	Traffic	Electricity	ETT			
						ETTh1	ETTh2	ETTm1	ETTm2
TimeDC	MAE	96	0.257	0.254	0.274	0.413	0.368	0.503	0.354
		192	0.285	0.297	0.285	0.436	0.389	0.512	0.401
	MSE	96	0.188	0.375	0.267	0.401	0.271	0.442	0.391
		192	0.247	0.405	0.294	0.428	0.302	0.465	0.421
Autoformer	MAE	96	0.312	0.370	0.343	0.453	0.473	0.548	0.342
		192	0.381	0.385	0.355	0.478	0.491	0.550	0.334
	MSE	96	0.255	0.597	0.236	0.465	0.412	0.542	0.265
		192	0.334	0.613	0.264	0.493	0.488	0.532	0.287
Informer	MAE	96	0.423	0.430	0.428	0.773	0.842	0.576	0.552
		192	0.482	0.476	0.446	0.788	0.954	0.597	0.532
	MSE	96	0.354	0.643	0.253	0.992	1.032	0.624	0.402
		192	0.478	0.710	0.271	0.987	1.055	0.653	0.432
Transformer	MAE	96	0.389	0.412	0.398	0.632	0.506	0.563	0.555
		192	0.588	0.431	0.422	0.612	0.513	0.588	0.576
	MSE	96	0.344	0.578	0.267	0.785	0.579	0.615	0.479
		192	0.524	0.567	0.288	0.732	0.542	0.643	0.455

Table 3: Dynamic Tensor Memory Cost on Four Datasets

Dataset	DC	MTT	TimeDC
Weather	10.0 GB	8.9 GB	3.3 GB
Traffic	17.8 GB	13.7 GB	10.9 GB
Electricity	8.5 GB	932.5 MB	516.0 MB
ETTh1	1.9 GB	845.5 MB	280.9 MB

training that on the original dataset (see Table 4). This makes it possible to train different candidate models using the condensed data set when performing model selection in TSMs, thus saving considerable training time.


Figure 7: Training Time Comparison

4.2.5 Training Time. As resource efficiency is important in dataset condensation to enable scalability, especially on resource-constrained edge computing devices, we study training time (of an epoch) for

the condensation methods. Figures 7(a) and 7(b) report the training time on Weather and ETTh1. We see that the training time of TimeDC is much lower than those of DC and MTT, which is largely because of the expert buffer in the CT²M module that stores the pre-computed trajectories. This indicates the feasibility of TimeDC for model deployment in large time series dataset reduction scenarios. We also compare the training time of TimeDC and coresets methods, which are included in the code repository, showing TimeDC achieves training times comparable to those of coresets methods. But coresets methods, especially Herding, need more time to construct coresets while having worse performance. Moreover, we compare the memory used by the dynamic (online) tensor across DC, MTT, and TimeDC in Table 3. TimeDC is able to reduce markedly the online memory and computation costs thanks to the training trajectories precomputed offline.

Table 4: Training Time of TimeDC and Training Time on Condensed and Original Datasets (s/epoch)

Dataset	TimeDC	Condensed Dataset	Original Dataset
Weather	22.39	4.31	35.26
Traffic	232.34	61.94	346.76
Electricity	314.56	41.14	522.85
ETTh1	14.38	4.93	20.43

We also study training time across different network architectures on Weather and ETTh1—see Figures 7(c) and 7(d). It is clear that TimeDC consumes the least training time. TimeDC is faster than the other methods by at least 73.0%, due to its patching mechanism, which reduces the complexity of the self-attention mechanism through input data simplification. Thus, TimeDC has lower training time across different networks, which still offering better performance. Finally, we compare the training time of TimeDC and the training time on original time series datasets based on the stacked *TSOperators*, as shown in Table 4. One can see that the training time of TimeDC is notably lower than when using original datasets,

showing the efficiency and practicality of time series dataset condensation. For example, the training time of TimeDC for dataset condensation is reduced by 39.84% compared to using the original dataset on Electricity.

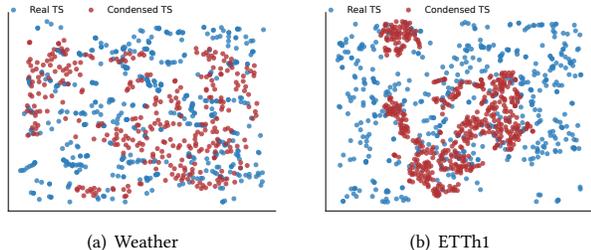


Figure 8: Dataset Distribution Comparison on Two Datasets

4.2.6 Case Study on Dataset Distribution. To observe the effectiveness of TimeDC on synthesizing condensed time series datasets that cover the original time series distribution well, we visualize the t-SNE [34] graphs of the original time series dataset and the condensed time series dataset on Weather and ETTh1. Figure 8 compares the dataset distributions. Blue and red dots represent the original (i.e., real) time series dataset and the condensed time series dataset, respectively. We sample 500 time series from the original time series dataset for the visualization of the original time series. We observe that the red dots are integrated with the blue dots, indicating that the original time series and condensed time series datasets exhibit similar distributions. This indicates that the condensed dataset is of high quality and that the proposed method is effective. Moreover, the condensed dataset can cover the original dataset evenly on Weather, showing that TimeDC achieves a more diverse and robust condensation result on Weather that exhibits less complicated temporal patterns.

Table 5: Performance on Streaming Learning ($PL = 96$)

Dataset	Metric	MAE		RMSE	
	Method	Autoformer _f	TimeDC	Autoformer _f	TimeDC
Weather	\mathcal{B}_0	0.603	0.456	0.625	0.477
	\mathcal{B}_1	0.694	0.512	0.873	0.533
	\mathcal{I}	0.823	0.656	1.014	0.512
Traffic	\mathcal{B}_0	0.541	0.465	0.657	0.489
	\mathcal{B}_1	0.635	0.502	0.829	0.611
	\mathcal{I}	0.735	0.624	0.929	0.784
Electricity	\mathcal{B}_0	0.557	0.437	0.601	0.542
	\mathcal{B}_1	0.628	0.504	0.875	0.599
	\mathcal{I}	0.844	0.689	1.046	0.702
ETTh1	\mathcal{B}_0	0.655	0.573	0.812	0.705
	\mathcal{B}_1	0.712	0.604	0.933	0.742
	\mathcal{I}	1.297	0.762	2.015	0.979

4.3 Application on Streaming Time Series

In real-world scenarios, time series data is often generated incrementally at edge devices that are distributed geographically. It may be preferable to process such streaming time series data on the edge directly to enable compliance with data access restrictions and exploit the potential for more efficient processing and lower

Table 6: Storage Comparison on Four Datasets

Storage	Weather	Traffic	Electricity	ETTh1
Whole Dataset	2.9 GB	20.0 GB	11.2 GB	313.1 MB
Condensed TS	38.5 MB	827.5 MB	308.2 MB	12.8 MB

latencies. However, time series data keeps growing while the capacities of edge devices are limited. We perform a performance, storage, and parameter comparison to determine whether our condensed time series data can be a key ingredient for streaming time series learning, thereby addressing also the catastrophic forgetting problem [11, 22], due to its condensed nature.

4.3.1 Performance Comparison. A naive solution to streaming time series learning is to use only newly arrived data, not all the data collected so far, to update model parameters continuously, called fine-tuning. We compare the performance of TimeDC and fine-tuning to test the effectiveness of TimeDC at streaming learning. For the fine-tuning, we use Autoformer as the basic model, entitled Autoformer_f. We split the original dataset into a base set \mathcal{B} and an incremental set \mathcal{I} with the ratio 7 : 3. \mathcal{B} and \mathcal{I} are further split into training data, validation data, and test data with the ratios 7 : 1 : 2. We first train and test the model on \mathcal{B} . The result is denoted as \mathcal{B}_0 . Then, we update the learned model with \mathcal{I} and test on \mathcal{B} and \mathcal{I} simultaneously, with the results denoted as \mathcal{B}_1 and \mathcal{I} , respectively. For the model update stage of TimeDC, we condense the base set and incorporate the condensed data into \mathcal{I} for subsequent model training to address concept drift in evolving time series datasets. The results in Table 5 show that TimeDC has the best performance in terms of MAE and RMSE when compared with Autoformer_f on four datasets. For example, TimeDC outperforms Autoformer_f by 18.36% ~ 21.54% and 9.82% ~ 32.89% in terms of MAE and RMSE on Electricity, respectively. Autoformer_f provides acceptable MAE and RMSE results on the base set, but their performance deteriorates on incremental set, especially on ETTh1. This indicates that a simple fine-tuning method is insufficient, due to forgetting problems caused by the possible concept drift. TimeDC achieves relatively stable performance on all base sets and incremental sets, demonstrating its superiority.

Table 7: Parameter Comparison on Four Architectures

Dataset	Autoformer	Informer	Transformer	TimeDC
Weather	10.6 M	11.4 M	10.6 M	1.8 M
Etth1	10.5 M	11.3 M	10.5 M	0.2 M

4.3.2 Storage and Parameter Comparison. As storage size is a main concern in time series streaming learning on edge devices, we compare the storage cost of the pre-processed whole dataset and the condensed time series dataset using four datasets, as shown in Table 6. Generally, the condensed time series data of TimeDC reduces the storage space substantially compared with the whole dataset. The condensed time series data is only 4.09% of the whole dataset on ETTh1. Thus, TimeDC enables much lower storage costs while achieving more promising performance on time series streaming learning.

Another characteristic of edge devices is limited computational capabilities. We thus compare the number of parameters of different network architectures including TimeDC (i.e., TSFE), Autoformer, Informer, and Transformer—see Table 7. TimeDC achieves much

fewer parameters, indicating that the TSFE module of TimeDC (mainly the patching mechanism) can significantly reduce computational costs. Tables 5, 6, and 7 indicate the feasibility and scalability of TimeDC for streaming learning on edge devices.

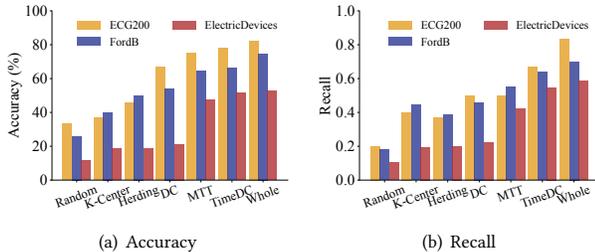


Figure 9: Time Series Classification Performance

4.4 Performance on Time Series Classification

It is important to show that the proposed TimeDC generalizes and can be used in other time series tasks. We conduct experiments on time series classification, considering three datasets (i.e., ECG200, ElectricDevices, and FordB) of the UCR time series archive[†]. A set of stacked TSFE modules with an FC is used as the basic classification model. CrossEntropy Loss is used as the objective of time series classification. *Accuracy* and *Recall* are adopted as evaluation metrics. The number of condensed time series is set to 50 for each dataset. The overall performance results are provided in Figure 9. The whole (dataset) indicates training on the whole original dataset, serving as an approximate upper-bound performance. TimeDC performs better than the best among the baselines. All coreset construction methods perform worse than dataset condensation methods. Herding has the best performance in most cases among the coreset methods. Overall, TimeDC achieves the best results on three time series classification datasets among the baselines, which shows that TimeDC can be extended to other time series tasks.

5 RELATED WORK

5.1 Time Series Modeling

Time series modeling attracts increasing interest due to the growing availability of time series data and rich downstream applications [13, 19, 20, 50, 51], such as traffic prediction [31, 44], electricity prediction [29, 53], and anomaly detection [32, 46]. Traditional time series analytic models are mostly based on statistical models [31]. However, the statistical models cannot capture complex temporal correlations of time series data effectively due to their limited learning capacity. Recent advances in deep learning techniques have sparked a surge of interest in applying neural network architectures for time series modeling [4, 43, 44, 53] outperforming traditional statistical models, including temporal convolutional network (TCN) based methods [9, 23], recurrent neural network (RNN) based methods [30], and transformer based methods [53, 54]. However, these methods are mostly supervised, and large training data is required resulting in high computational cost. At such scales, it becomes burdensome to store and preprocess the data and calls for specialized equipment and infrastructure to train machine learning models on them.

[†]https://www.cs.ucr.edu/~eamonn/time_series_data_2018/

5.2 Coreset and Dataset Condensation

Coreset construction [1, 5, 7, 16, 22] is the traditional dataset reduction approach that works by identifying the most representative training samples in an original dataset, aiming at achieving models trained on the coreset that are provably competitive with those built on the full dataset. Typically, coreset construction methods choose samples that are representative for training based on heuristic criteria [1, 7, 22], e.g., minimizing distance between coreset and whole-dataset centers [1]. Nonetheless, these heuristic methods cannot guarantee optimal solutions and the presence of representative samples for the downstream task. A recent approach, dataset condensation (or distillation) [6, 48], is proposed to address these limitations by learning a small typical dataset that distills the most important knowledge from a given large dataset, such that a model trained on it can obtain comparable testing accuracy to that trained on the original training set. Existing dataset distillation methods have demonstrated superior performance, which can be categorized into matching-based methods [6, 48, 49] and kernel-based methods [28, 55]. Matching-based methods generate synthetic datasets by matching gradients [48], multi-step parameters [6], and distributions [49] between two surrogate models trained on the synthetic dataset and the original dataset. Kernel-based methods [28, 55] treat the synthetic dataset as the parameters to be optimized inspired by kernel functions e.g., neural tangent kernel. However, most of the above methods are designed for computer vision, and cannot be applied to time series directly due to complex temporal correlations such as seasonality and trend. In addition, these methods require substantial storage costs.

6 CONCLUSION

We present TimeDC, a new efficient time series dataset condensation framework that aims to synthesize a small but informative condensed time series dataset summarizing an original large time series dataset. To capture complex temporal dependencies, we design a time series feature extraction module with stacked *TSo*perators. In addition, decomposition-driven frequency matching is proposed to ensure similar temporal patterns between the condensed and original time series datasets. To enable effective and generalized dataset condensation, we propose a curriculum training trajectory matching module with an expert buffer that aims to decrease the training cost. Comprehensive experiments on original datasets offer evidence that TimeDC achieves state-of-the-art accuracy and requires fewer computational and storage resources.

7 ACKNOWLEDGMENTS

This work was supported in part by the Independent Research Fund Denmark under agreement 8048-00038B, the Villum Fonden under agreement 40567, the Innovation Fund Denmark project DIREC (9142-00001B), the National Natural Science Foundation of China (62372179, 62472068), Shenzhen Municipal Science and Technology R&D Funding Basic Research Program (JCYJ20210324133607021), Municipal Government of Quzhou (2023D044), and Key Laboratory of Data Intelligence and Cognitive Computing, Longhua District, Shenzhen.

REFERENCES

- [1] Pankaj K Agarwal, Sarel Har-Peled, and Kasturi R Varadarajan. 2004. Approximating extent measures of points. *JACM* 51, 4 (2004), 606–635.
- [2] Charu C Aggarwal and Philip S Yu. 2008. On static and dynamic methods for condensation-based privacy-preserving data mining. *TODS* 33, 1 (2008), 1–39.
- [3] Fabrizio Angiulli. 2007. Fast nearest neighbor condensation for large data sets classification. *TKDE* 19, 11 (2007), 1450–1464.
- [4] Angela Bonifati, Francesco Del Buono, Francesco Guerra, and Donato Tiano. 2022. Time2Feat: learning interpretable representations for multivariate time series clustering. *PVLDB* 16, 2 (2022), 193–201.
- [5] 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. *PVLDB* 17, 11 (2024), 2708–2721.
- [6] George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A Efros, and Jun-Yan Zhu. 2022. Dataset distillation by matching training trajectories. In *CVPR*. 4750–4759.
- [7] Chengliang Chai, Jiabin Liu, Nan Tang, Ju Fan, Dongjing Miao, Jiayi Wang, Yuyu Luo, and Guoliang Li. 2023. GoodCore: Data-effective and Data-efficient Machine Learning through Coreset Selection over Incomplete Data. *SIGMOD* 1, 2 (2023), 1–27.
- [8] 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. In *ICLR*.
- [9] Yunyao Cheng, Peng Chen, Chenjuan Guo, Kai Zhao, Qingsong Wen, Bin Yang, and Christian S Jensen. 2024. Weakly guided adaptation for robust time series forecasting. *PVLDB* 17, 4 (2024), 766–779.
- [10] Yunyao Cheng, Chenjuan Guo, Bin Yang, Haomin Yu, Kai Zhao, and Christian S. Jensen. 2024. A Memory Guided Transformer for Time Series Forecasting. *PVLDB* 18 (2024).
- [11] Zhixuan Chu, Ruopeng Li, Stephen Rathbun, and Sheng Li. 2023. Continual causal inference with incremental observational data. In *ICDE*. 3430–3439.
- [12] 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.
- [13] Yuchen Fang, Jiandong Xie, Yan Zhao, Lu Chen, Yunjun Gao, and Kai Zheng. 2024. Temporal-Frequency Masked Autoencoders for Time Series Anomaly Detection. In *ICDE*. 1228–1241.
- [14] Reza Zanjirani Farahani and Masoud Hekmatfar. 2009. *Facility location: concepts, models, algorithms and case studies*. Springer Science & Business Media.
- [15] Dan Feldman and Michael Langberg. 2011. A unified framework for approximating and clustering data. In *STOC*. 569–578.
- [16] Dan Feldman, Melanie Schmidt, and Christian Sohler. 2020. Turning big data into tiny data: Constant-size coresets for k-means, PCA, and projective clustering. *SICOMP* 49, 3 (2020), 601–657.
- [17] Sona Hasani, Saravanan Thirumuruganathan, Abolfazl Asudeh, Nick Koudas, and Gautam Das. 2018. Efficient construction of approximate ad-hoc ML models through materialization and reuse. *PVLDB* 11, 11 (2018), 1468–1481.
- [18] Søren Kejsler Jensen, Torben Bach Pedersen, and Christian Thomsen. 2018. Modelardb: Modular model-based time series management with spark and cassandra. *PVLDB* 11, 11 (2018), 1688–1701.
- [19] Duc Kieu, Tung Kieu, Peng Han, Bin Yang, Christian S. Jensen, and Bac Le. 2024. TEAM: Topological Evolution-aware Framework for Traffic Forecasting. *PVLDB* 18 (2024).
- [20] 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.
- [21] Zhichen Lai, Dalin Zhang, Huan Li, Christian S Jensen, Hua Lu, and Yan Zhao. 2023. LightCTS: A Lightweight Framework for Correlated Time Series Forecasting. *SIGMOD* 1, 2 (2023), 1–26.
- [22] Yiming Li, Yanyan Shen, and Lei Chen. 2022. Camel: Managing Data for Efficient Stream Learning. In *SIGMOD*. 1271–1285.
- [23] Chi Harold Liu, Chengzhe Piao, Xiaoxin Ma, Ye Yuan, Jian Tang, Guoren Wang, and Kin K Leung. 2021. Modeling citywide crowd flows using attentive convolutional LSTM. In *ICDE*. 217–228.
- [24] Ziqiao Liu, Hao Miao, Yan Zhao, Chenxi Liu, Kai Zheng, and Huan Li. 2024. LightTR: A Lightweight Framework for Federated Trajectory Recovery. In *ICDE*. 4422–4434.
- [25] Mario Lucic, Matthew Faulkner, Andreas Krause, and Dan Feldman. 2018. Training gaussian mixture models at scale via coresets. *J. Mach. Learn. Res* 18, 160 (2018), 1–25.
- [26] Hao Miao, Jiaying Shen, Jiannong Cao, Jiangnan Xia, and Senzhang Wang. 2023. MBA-STNet: Bayes-Enhanced Discriminative Multi-Task Learning for Flow Prediction. *TKDE* 35, 7 (2023), 7164–7177.
- [27] Hao Miao, Yan Zhao, Chenjuan Guo, Bin Yang, Kai Zheng, Feiteng Huang, Jiandong Xie, and Christian S. Jensen. 2024. A Unified Replay-Based Continuous Learning Framework for Spatio-Temporal Prediction on Streaming Data. In *ICDE*. 1050–1062.
- [28] Timothy Nguyen, Zhoung Chen, and Jaehoon Lee. 2020. Dataset Meta-Learning from Kernel Ridge-Regression. In *ICLR*.
- [29] Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *ICLR*.
- [30] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. 2020. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *Int J Forecast* 36, 3 (2020), 1181–1191.
- [31] Shashank Shekhar and Billy M Williams. 2007. Adaptive seasonal time series models for forecasting short-term traffic flow. *TRR* 2024, 1 (2007), 116–125.
- [32] Emmanouil Sylligardos, Paul Boniol, John Paparrizos, Panos Trahanias, and Themis Palpanas. 2023. Choose wisely: An extensive evaluation of model selection for anomaly detection in time series. *PVLDB* 16, 11 (2023), 3418–3432.
- [33] Kai Sheng Tai, Vatsal Sharan, Peter Bailis, and Gregory Valiant. 2018. Sketching linear classifiers over data streams. In *SIGMOD*. 757–772.
- [34] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *J Mach Learn Res* 9, 11 (2008).
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *NeurIPS* 30 (2017).
- [36] Chen Wang, Xiangdong Huang, Jialin Qiao, et al. 2020. Apache IoTDB: Time-series database for internet of things. *PVLDB* 13, 12 (2020), 2901–2904.
- [37] Qitong Wang, Stephen Whitmarsh, Vincent Navarro, and Themis Palpanas. 2022. iEDeAL: A Deep Learning Framework for Detecting Highly Imbalanced Interictal Epileptiform Discharges. *PVLDB* 16, 3 (2022), 480–490.
- [38] Senzhang Wang, Meiyue Zhang, Hao Miao, Zhaohui Peng, and Philip S Yu. 2022. Multivariate correlation-aware spatio-temporal graph convolutional networks for multi-scale traffic prediction. *TIST* 13, 3 (2022), 1–22.
- [39] Max Welling. 2009. Herding dynamical weights to learn. In *ICML*. 1121–1128.
- [40] Qingsong Wen, Kai He, Liang Sun, Yingying Zhang, Min Ke, and Huan Xu. 2021. RobustPeriod: Robust time-frequency mining for multiple periodicity detection. In *SIGMOD*. 2328–2337.
- [41] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. 2021. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *NeurIPS* 34 (2021), 22419–22430.
- [42] Xinle Wu, Xingjian Wu, Bin Yang, Lekui Zhou, Chenjuan Guo, Xiangfei Qiu, Jilin Hu, Zhenli Sheng, and Christian S. Jensen. 2024. AutoCTS++: zero-shot joint neural architecture and hyperparameter search for correlated time series forecasting. *Vldb* 33, 5 (2024), 1743–1770.
- [43] Xinle Wu, Xingjian Wu, Dalin Zhang, Miao Zhang, Chenjuan Guo, Bin Yang, and Christian S. Jensen. 2024. Fully Automated Correlated Time Series Forecasting in Minutes. *PVLDB* 18 (2024).
- [44] 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. *SIGMOD* 1, 1 (2023), 1–26.
- [45] Jinzhao Xiao, Yuxiang Huang, Changyu Hu, Shaoxu Song, Xiangdong Huang, and Jianmin Wang. 2022. Time series data encoding for efficient storage: A comparative analysis in apache IoTDB. *PVLDB* 15, 10 (2022), 2148–2160.
- [46] Ronghui Xu, Hao Miao, Senzhang Wang, Philip S Yu, and Jianxin Wang. 2024. PeFAD: A Parameter-Efficient Federated Framework for Time Series Anomaly Detection. In *SIGKDD*. 3621–3632.
- [47] Xinyang Yu, Yanqing Peng, Feifei Li, Sheng Wang, Xiaowei Shen, Huijun Mai, and Yue Xie. 2020. Two-level data compression using machine learning in time series database. In *ICDE*. 1333–1344.
- [48] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. 2021. Dataset Condensation with Gradient Matching. In *ICLR*.
- [49] Ganlong Zhao, Guanbin Li, Yipeng Qin, and Yizhou Yu. 2023. Improved distribution matching for dataset condensation. In *CVPR*. 7856–7865.
- [50] Kai Zhao, Chenjuan Guo, Yunyao Cheng, Peng Han, Miao Zhang, and Bin Yang. 2023. Multiple Time Series Forecasting with Dynamic Graph Modeling. *PVLDB* 17, 4 (2023), 753–765.
- [51] 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.
- [52] Kaiping Zheng, Gang Chen, Melanie Herschel, Kee Yuan Ngiam, Beng Chin Ooi, and Jinyang Gao. 2021. PACE: learning effective task decomposition for human-in-the-loop healthcare delivery. In *SIGMOD*. 2156–2168.
- [53] 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.
- [54] 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.
- [55] Yongchao Zhou, Ehsan Nezhdaryya, and Jimmy Ba. 2022. Dataset distillation using neural feature regression. *NeurIPS* 35 (2022), 9813–9827.