

Learning Multi-Pattern Normalities in the Frequency Domain for Efficient Time Series Anomaly Detection

Feiyi Chen¹, Yingying Zhang², Zhen Qin¹, Lunting Fan², Renhe Jiang³, Yuxuan Liang⁴,
Qingsong Wen⁵, Shuiguang Deng^{1†}

¹Zhejiang University ²Alibaba Group ³The University of Tokyo

⁴The Hong Kong University of Science and Technology (Guangzhou) ⁵Squirrel AI

¹{chenfeiyi,zhenqin,dengsg}@zju.edu.cn

²congrong.zyy@alibaba-inc.com ²lunting.fan@taobao.com

³jiangrh@csis.u-tokyo.ac.jp ⁴yuxliang@outlook.com ⁵qingsongedu@gmail.com

Abstract—Anomaly detection significantly enhances the robustness of cloud systems. While neural network-based methods have recently demonstrated strong advantages, they encounter practical challenges in cloud environments: the contradiction between the impracticality of maintaining a unique model for each service and the limited ability to deal with diverse normal patterns by a unified model, as well as issues with handling heavy traffic in real time and short-term anomaly detection sensitivity. Thus, we propose MACE, a multi-normal-pattern accommodated and efficient anomaly detection method in the frequency domain for time series anomaly detection. There are three novel characteristics of it: (i) a pattern extraction mechanism excelling at handling diverse normal patterns with a unified model, which enables the model to identify anomalies by examining the correlation between the data sample and its service normal pattern, instead of solely focusing on the data sample itself; (ii) a dualistic convolution mechanism that amplifies short-term anomalies in the time domain and hinders the reconstruction of anomalies in the frequency domain, which enlarges the reconstruction error disparity between anomaly and normality and facilitates anomaly detection; (iii) leveraging the sparsity and parallelism of frequency domain to enhance model efficiency. We theoretically and experimentally prove that using a strategically selected subset of Fourier bases can not only reduce computational overhead but is also profitable to distinguish anomalies, compared to using the complete spectrum. Moreover, extensive experiments demonstrate MACE’s effectiveness in handling diverse normal patterns with a unified model and it achieves state-of-the-art performance with high efficiency.

Index Terms—Anomaly detection, multiple normal patterns, efficiency

I. INTRODUCTION

Anomaly detection is an extensively researched issue crucial for bolstering cloud system reliability and curbing labor expenses [1], [2]. Reconstruction-based methods, in particular, have demonstrated state-of-the-art performance in this domain [3]–[5]. Despite these advancements, several significant challenges persist, as outlined below:

- **C1. Limited capacity to accommodate diverse normal patterns:** In practical scenarios, cloud centers host millions

of services concurrently, and each service exhibits a unique normal pattern. To illustrate this, we project the normal data of each service in the Server Machine Dataset onto a two-dimensional space and visualize their distribution in Fig. 1(a), where the data points are scattered randomly. It is impractical and excessively expensive to maintain a tailored model for each service [9]. However, most state-of-the-art methods train a tailored model for each cloud service [3], [4], [7]. Moreover, it is reported that many reconstruction-based methods can only effectively capture a few dominant normal patterns in the training set for each trained model [10]. We further verify this by conducting an experiment comparing performances of some state-of-the-art methods when training a unified model for ten services and tailoring ten models for ten services, as shown in Fig.1(b). The performance of the unified model is substantially lower than the one of the tailored models.

- **C2. Inefficiency in handling heavy traffic in real time:** In large cloud centers, the volume of service traffic can escalate to hundreds of thousands of requests per second. In these high-demand scenarios, numerous deep learning-based methods face challenges in efficiently handling peak traffic in real time. Furthermore, a notable complication arises from the incorporation of recurrent networks in several anomaly detection neural networks, such as VRNN [11], omniAnomaly [3], and MSCRED [7]. This inclusion hampers operator parallelization, as recurrent networks cannot be effectively parallelized across different recurrent steps.
- **C3. Insensitivity to single point anomalies:** Encoder-decoder neural networks, which are used by many reconstruction-based methods, are insensitive to single point anomalies and often overlook them [8], [12].

Tackling these challenges is imperative for enhancing the efficacy and applicability of deep learning-based anomaly detection methods in practical cloud environments. Therefore,

[†]Corresponding author.

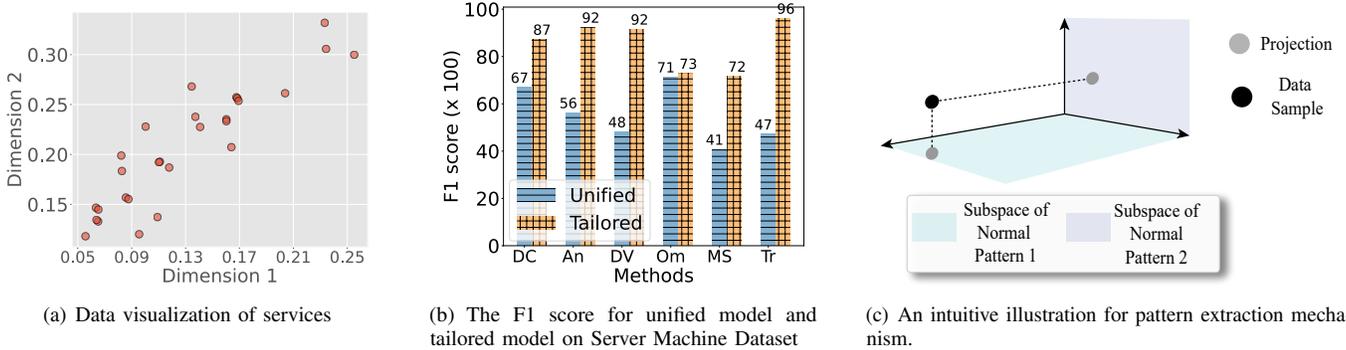


Fig. 1. (a) The normal data of each service is compressed into a two-dimensional vector, which is scattered randomly. (b) The figure shows F1 score of some SOTA methods: DCdetector [6], AnomalyTransformer [4], DVGCRN [5], OmniAnomaly [3], MSCRED [7], TranAD [8]. (c) A data sample is projected to a normal pattern subspace in a pattern extraction mechanism. When the data sample is closer to the normal pattern subspace, it is easier to reconstruct it from its projection with less reconstruction error. Thus, the data sample is more likely to be inferred as normality for the normal pattern 1 than normal pattern 2. anomalies.

we design innovative mechanisms to address these issues, as outlined below:

- **S1.** To enhance the model’s ability to accommodate diverse normal patterns (C1), we propose a pattern extraction mechanism. *The most challenging problem of dealing with diverse normal patterns is that an anomaly for one normal pattern could be a normality for another.* Thus, we detect anomalies according to the correlation between the data sample and its service normal pattern, instead of the data sample itself. The pattern extraction mechanism identifies a normal pattern subspace in the frequency domain for each service. Subsequently, it tailors a representation for each data sample from the sample’s projection on its service normal pattern subspace. In this way, when the data sample is closer to its service normality subspace, it is easier to reconstruct from the representation with less reconstruction error and is more likely to be inferred as normality, as shown in Fig. 1(c).
- **S2.** To enhance model efficiency and parallelism (C2), we introduce a frequency-domain-based approach. Anomaly detection in the frequency domain can leverage sparsity to reduce computational overhead and enhance fine-grained parallelism by eliminating temporal dependencies. To exploit frequency domain sparsity, we propose a strategy for selecting a subset of Fourier bases for each service in the pattern extraction process. We theoretically demonstrate that using only this subset of Fourier bases not only reduces computational overhead but also improves anomaly detection performance compared to utilizing the complete spectrum. For effective anomaly detection in the frequency domain, we introduce a dualistic convolution mechanism to replace the standard convolution in the auto-encoder. This mechanism hinders the reconstruction of anomalies while keeping the reconstruction of normalities easy.
- **S3.** To enhance the sensitivity to short-term anomalies (C3), we introduce the dualistic convolution mechanism to the time domain, which amplifies the anomalies,

facilitating their detection while maintaining the similarity of normality to the original time series. As illustrated in Fig. 3(b), the dualistic convolution mechanism extends anomalies while maintaining the similarity of normality to the original time series.

Accordingly, this work makes the following novel and unique contributions to the field of anomaly detection:

- We propose a novel pattern extraction mechanism to deal with diverse normal patterns by facilitating the model to detect anomalies from the correlation between the data sample and its service normal pattern, instead of only the data sample.
- We propose a dualistic convolution mechanism. In the time domain, it amplifies the anomalies. In the frequency domain, it hinders the reconstruction of anomalies while keeping the reconstruction of normalities easy.
- We leverage the sparsity and parallelism of the frequency domain to improve model efficiency. It is theoretically and experimentally proved that using just a subset of Fourier bases can not only reduce computational overhead but also achieve better anomaly detection performance compared with using the complete spectrum.

Moreover, we conduct extensive experiments on four real-world datasets to demonstrate that MACE can effectively capture diverse normal patterns with a unified model and averagely improve F1 score by 8.7%, compared with the strongest method in baselines, while achieving $4\times$ faster.

II. RELATED WORK

Anomaly detection is a crucial task that focuses on identifying outliers within time series data and has been the subject of extensive research. The existing body of work can be broadly categorized into three main groups: classical machine learning [13]–[15] and statistical methods [16]–[19], signal processing-based methods [20]–[22], and deep learning-based methods [23]–[28]. In the following, we provide a brief overview of each category and a review of multi-task learning since multiple

normal pattern learning is highly correlated to the aim of this work.

A. Anomaly Detection

Classical methods. Conventional statistical and machine learning methods, as highlighted in earlier works [29]–[31], operate without the need for extensive training data and remain unaffected by the challenge of diverse normal patterns. Furthermore, they typically incur little computational overhead. Despite these advantages, these approaches are contingent upon certain assumptions and, in real-world applications, exhibit constrained robustness [32].

Signal-processing-based methods. Signal processing-based methods leverage the advantages of fine-grained parallelism and the sparsity inherent in the frequency domain. Despite these advantages, they encounter difficulties in simultaneously capturing both global and subtle features while maintaining manageable computational overhead. For instance, the Fourier transform [33] excels at capturing global information but struggles with subtle local features. In contrast, context-aware Discrete Fourier Transform (DFT) and context-aware Inverse DFT (IDFT) select Fourier bases based on a given normal pattern and integrate with dualistic convolution mechanism, enhancing the capacity to extract subtle features. The theoretical evidence supports that this approach widens the reconstruction error gap between anomalies and normal patterns. JumpStarter [32], a recent state-of-the-art method in this category, suffers from significant inference time overhead and struggles to handle heavy traffic loads in real time.

Deep learning-based methods. Deep learning-based methods prove to be particularly effective for variable time series [8]. These methods can be broadly categorized into prediction-based approaches [34], [35], reconstruction-based methods [4], [5], and classifier-based methods [36]–[38]. Prediction-based methods, such as LSTM-NDT [34] and DAGMM [35], incorporate recurrent networks, which are non-parallelizable and inefficient. Similarly, reconstruction-based methods like Donut [39] and OmniAnomaly [3] also rely on recurrent neural networks. Recent advancements, exemplified by USAD [40] and GDN [41], replace recurrent neural networks with attention-based architectures to expedite the training process. However, these methods face challenges in effectively capturing long-term dependencies due to the removal of recurrent networks and the use of small input windows [8]. In contrast to these approaches, anomaly detection in the frequency domain eliminates temporal dependencies without sacrificing global information. Consequently, there is no need for recurrent neural networks, yet the model can still effectively capture long-term features. The most recent works leverage the power of transformers, exemplified by AnomalyTransformer [4] and TranAD [8], enabling fine-grained parallelism. However, these methods encounter challenges in handling diverse normal patterns with a unified model.

B. Multi-task Learning

Multitask Learning (MTL) is a machine learning paradigm in which a model is trained to concurrently address multiple interrelated tasks. Unlike the conventional approach of training separate models for individual tasks, MTL facilitates the sharing of specific model parameters. This sharing mechanism allows the model to glean common representations across the spectrum of tasks it is designed to handle. Within the realm of MTL research, two principal categories emerge: hard-sharing methods and soft-sharing methods. Hard-sharing methods involve the sharing of common low-level hidden layers among the tasks. In contrast, soft-sharing methods adopt a more nuanced strategy. They foster the exchange of general knowledge among multiple models by incorporating regularization techniques into neural network parameters [42] or establishing connections across networks [43], [44]. However, both soft-sharing and hard-sharing methods necessitate the maintenance of task-specific neural network layers. This maintenance, despite the advancements in MTL, remains computationally expensive, particularly when dealing with a multitude of diverse services and tasks in real-world applications. Finding more efficient ways to share knowledge across tasks without compromising performance is an ongoing challenge.

III. PRELIMINARY

A normal multivariate time series is denoted by $X_N \in \mathbb{R}^{T \times m}$, where T denotes the sliding window length and m denotes the number of feature dimensions. Similarly, an anomalous multivariate time series is denoted by $X_A \in \mathbb{R}^{T \times m}$. A normal pattern is a distribution that determines the possibility $p(X_N[t]|X_N[1:t])$, where $X_N[t]$ denotes the t^{th} element of X_N and $X_N[1:t]$ denotes the subsequence of X_N from the first slot to the $(t-1)^{th}$ slot. The reconstruction-based method compresses X_N and X_A separately and reconstructs them. The reconstructed normal time series is denoted by \tilde{X}_N , while the reconstructed anomalous one is denoted by \tilde{X}_A . The objective of reconstruction-based methods is shown in Eq. 1, which enlarges the gap between the reconstruction error of normalities and anomalies. Subsequently, it determines a threshold by some strategies, such as POT [18]. When the reconstruction error exceeds the threshold, the input data is inferred as an anomaly.

$$\max. |\tilde{X}_A - X_A| - |\tilde{X}_N - X_N| \quad (1)$$

Additionally, other symbols utilized in this paper are outlined in Table I.

IV. PROPOSED METHOD

A. Overview

The overview of MACE is depicted in Figure 2. There are roughly four stages in MACE. In the first stage, it amplifies anomalies and makes them easier to detect. In the second stage, it extracts the normal pattern subspace of each service, transforms the time series into the frequency domain, and obtains its frequency representation. In the third stage, it

TABLE I
THE DEFINITION OF SYMBOLS.

Symbol	Definition
γ	The power in dualistic convolution
σ	The scaling factor in dualistic convolution
s	The stride length of convolution
ω_i	The frequency of i^{th} strongest signal of normal pattern
$\mathcal{A}_N(\omega_i)$	The amplitude of normality spectrum corresponding to ω_i
$\mathcal{A}_A(\omega_i)$	The amplitude of anomaly spectrum corresponding to ω_i
$q_N(\omega_i)$	The normalized value of $\mathcal{A}_N(\omega_i)$
$q_A(\omega_i)$	The normalized value of $\mathcal{A}_A(\omega_i)$
\mathcal{A}_i	The i^{th} amplitude of a spectrum
$\Delta\mathcal{A}_i$	The shift variable adding to normal spectrum
ΔA	The expectation of shift variable
α_i	The i^{th} element in convolution kernel
$\mathcal{F}_{i,j}$	The Fourier result of j^{th} base of i^{th} feature
$\omega_{i,j}$	The frequency of Fourier base corresponding to $\mathcal{F}_{i,j}$
Σ	The correlation matrix of joint distribution of amplitudes
ν_i	The i^{th} row and i^{th} column of Σ
m	The number of feature dimensions
n	The number of signals in a spectrum
k	The number of signals in the context-aware selecting subset
μ	The expectation of joint distribution of amplitudes
μ_i	The i^{th} element of μ

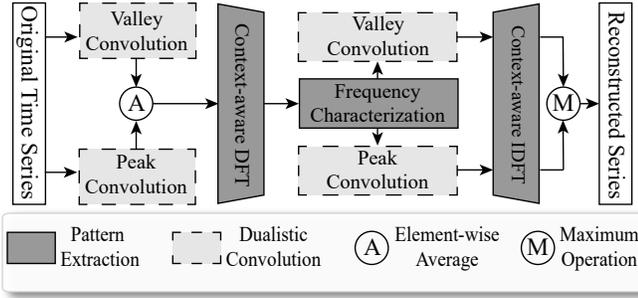


Fig. 2. The model architecture of MACE

reconstructs the representation in the frequency domain. In the final stage, it transforms the reconstructed spectrum back into the time domain. More specifically, the details are provided in the following:

- **Stage 1. Amplify anomalies:** MACE initially employs *dualistic convolution* in the time domain to amplify anomalies. The dualistic convolution is armed with a peak convolution and a valley convolution, which are used to emphasize the upward deviations and downward deviations respectively. We average the results of peak convolution and valley convolution in an element-wise manner.
- **Stage 2. Time domain \rightarrow Frequency representation:** We present a new *pattern extraction* mechanism to leverage the sparsity inherent in the frequency domain, thereby augmenting the model’s capacity for generalization across various normal patterns. The *pattern extraction* mechanism consists of a context-aware Discrete Fourier Transform (DFT), a frequency characterization module and a context-aware Inverse Discrete Fourier Transform (IDFT). The context-aware DFT identifies a normal pattern subspace in the frequency domain by selecting a subset of Fourier

bases for each service and projects the services’ data sample to the subspace in the frequency domain. After that, the frequency characterization module learns a frequency representation for the sample. In this way, we learn a correlation between the data sample and a normal frequency subspace. In the following, we detect anomalies from this correlation instead of the data sample itself.

- **Stage 3. Reconstruction:** MACE replaces the vanilla convolution in the auto-encoder with the peak convolution and valley convolution separately to reconstruct the frequency representation, which enlarges the reconstruction error disparity between normality and anomalous.
- **Stage 4. Reconstructed spectrum \rightarrow Time domain:** MACE uses the context-aware IDFT in pattern extraction to transform the spectrums reconstructed by peak convolution and valley convolution back to the time domain. Finally, it selects the time series with the highest reconstruction error as the final reconstructed time series.

During the training process, we use stochastic gradient descent to minimize the reconstruction error obtained from stage 4, which requires determining a hyperparameter – learning rate.

There may be a concern about why we use different operations in stage 1 and stage 2 to the results of peak and valley convolution. It is worth noting that in stage 1, what we average is the manipulated data, while in stage 4, from what we select the maximum is the reconstruction error. In stage 1, we use average operation. In this way, the amplified downward and upward deviations are fused into the following networks. In stage 4, since the peak convolution and valley convolution enlarge the reconstruction error of upward deviations and downward deviations separately, for each time slot we compare the reconstruction error of them and pick the maximum one. In this way, we can detect both the upward deviation and the downward deviation.

B. Dualistic Convolution

The dualistic convolution consists of a peak convolution and a valley convolution, which targets at upward deviation and downward deviation respectively. It exhibits different effects in the time domain and frequency domain. In the time domain, it is reported that the short-term anomalies are easily neglected by encoder-decoder models [8]. Thus, we propose the dualistic convolution mechanism to extend the anomaly in the time domain, which makes the anomalies more conspicuous and easily detected. In the frequency domain, it hinders the reconstruction of anomalies, while keeping the reconstruction of normal samples easy to facilitate the model to identify anomalies.

The dualistic convolution mechanism is depicted in Eq.2, where γ is a hyperparameter to make the convolution pay more attention to the deviations, σ is a scaling hyperparameter and $\text{Conv}_{1 \times L}(x, s)$ denotes applying standard convolution to x with stride s and kernel length L . The peak convolution and valley convolution are obtained by setting different γ . The peak convolution is the dualistic convolution with γ greater than 1,

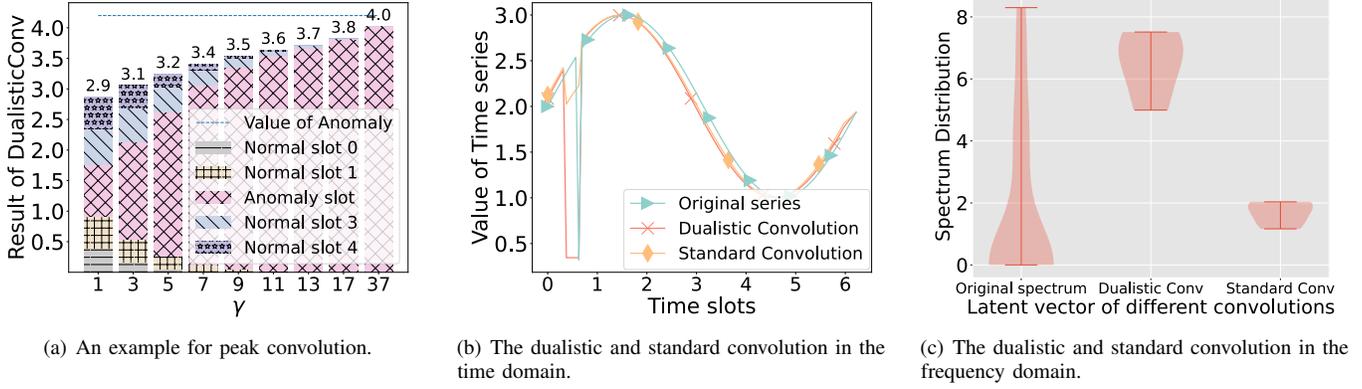


Fig. 3. (a) The figure shows the contributions of different time slots in peak convolution result in a convolution window when specifying different γ . As γ grows, the contribution of deviations increases significantly. (b)-(c) The utility of dualistic convolution, compared with standard convolution.

while the valley convolution is the dualistic convolution with γ less than -1 . The larger the absolute value of γ is, the more dominant the deviations are in the convolution result, as an example shown in Fig.3(a).

$$\text{DualisticConv}(x) = \sqrt[\gamma]{\text{Conv}_{1 \times L}\left(\frac{x^\gamma}{\sigma}, s\right)} \quad (2)$$

$$\gamma \in \{2g + 1 | g \in \mathbb{Z} \wedge g \neq 0\}, \sigma > 0$$

The dualistic convolution yields distinct outcomes in both the time and frequency domains. To illustrate this visually, the convolution results are depicted in Fig.3. In the time domain, standard convolution mitigates deviations, while dualistic convolution amplifies them. In the frequency domain, standard convolution compresses the spectrum to a low-dimensional space around the main body of the spectrum distribution, whereas dualistic convolution compresses it to the space around the "long tail" of the spectrum distribution. As confirmed later, anomalous data typically exhibits a long-tail distribution in the frequency spectrum, while normal data rarely shows such a tail. Considering that reconstruction is easier from latent vectors close to the main body of the distribution than from those close to the tail, dualistic convolution hinders the reconstruction of anomalies while keeping the process for normal data straightforward.

In the Time Domain, the stride s of dualistic convolution is set to 1. In this way, the dualistic convolution functions as a weighted summation operator for each kernel-length sliding window, which emphasizes the deviations in its results. Thus, once an anomaly is included in a convolution window, the convolution result of this window will be dominated by the anomaly. Consequently, short-term anomaly will be extended by the kernel length, as shown in Fig.3(b), which makes short-term anomaly easier to detect.

In the Frequency Domain, it has been reported that most anomalies manifest themselves as strong signals with high-energy components [32], which renders their spectrum higher variability. This can be proven by statistical data in Table II, where the amplitude variances of anomalies are greater than those of normal patterns. Consequently, we set the stride of

dualistic convolution in the frequency domain to the size of the convolution kernel. In this way, the dualistic convolution actually picks the min (valley convolution) or the max (peak convolution) amplitude in each kernel-length segment to comprise the latent vector, as shown in Fig.4(a). Intuitively, the dualistic convolution in the frequency domain hinders the reconstruction of anomalies and keeps the reconstruction of normalities easy, because when the components in a spectrum are highly variable, the dualistic convolution tends to pick the high-energy components to comprise its latent vector, which can obviously deviate from other components and are difficultly reconstructed from. In contrast, when the components in a spectrum are closer to each other, the latent vector obtained by dualistic convolution does not deviate significantly from the original spectrum and can represent it better. Furthermore, we conduct a theoretical comparison of the challenges involved in reconstructing a spectrum from its latent vector for both normal and anomalous cases. The level of reconstruction difficulty is directly associated with the gap between the latent vector and the original spectrum. In Theorem 1, we examine the upper bound of this gap for both normal and anomalous scenarios. Our analysis reveals that the constraints on the gap in normal cases are more stringent when compared to those in anomalous cases.

Definition 1. The gap between the convolution result and the original spectrum of each convolution window is defined as $\sum_{j=1}^n \mathbb{E}(|\text{DualisticConv}(\mathcal{A}) - \mathcal{A}_j|)$, where \mathcal{A} is the amplitudes of spectrum in the convolution window, \mathcal{A}_j is the j^{th} element in \mathcal{A} and n is the kernel length.

Theorem 1. When the amplitudes follow a Gaussian joint distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, the distance between the latent vector and the original spectrum is upper bounded by $2^{\frac{\gamma-1}{\gamma}} n \sqrt{\sum_{i=1}^n (\gamma-1)!! \nu_i^\gamma |\alpha_i| + |\alpha_i \mu_i^\gamma|} - \sum_{j=1}^n \mu_j$, where α_i is the i^{th} element in the kernel divided by σ , μ_i is the i^{th} element of $\boldsymbol{\mu}$, ν_i is the i^{th} row and i^{th} column element of $\boldsymbol{\Sigma}$ and $n!!$ denotes $[n(n-2)(n-4) \cdots 1]$.

Proof skeleton. We first transform the gap formation to $\left| \sum_{j=1}^n \mathbb{E}(\text{DualisticConv}(\mathcal{A}) - \mathcal{A}_j) \right|$. That is because

DualisticConv(\mathcal{A}) - $\mathcal{A}_j \geq 0, \forall j$ as long as we choose a big enough γ . Taking the peak convolution as an example, $\lim_{\gamma \rightarrow \infty} \text{DualisticConv}(\mathcal{A}) = \max(\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n)$. Thus, when we take a big γ , we can approximately confirm that $\mathcal{A} - \mathcal{A}_j \geq 0, \forall j$. Subsequently, we further transform the gap as shown in Eq.4-Eq.5. Since the function $f(x) = \sqrt[\gamma]{x}$ is a concave function when $x \geq 0$, it can be further scaled by Jensen inequality [45] as shown in Eq.6. Let $\mathcal{A}_i = \bar{\mathcal{A}}_i + \mu_i$, where $\bar{\mathcal{A}}_i \sim \mathcal{N}(\mathbf{0}, \Sigma)$, and we obtain Eq.7. The equation is further scaled by power mean inequality [46] and we obtain Eq.8. Facilitated by the property of Gamma Function [47], it can be computed that $\mathbb{E}(|\bar{\mathcal{A}}_i^\gamma|) = (\gamma - 1)!! \nu_i^\gamma$. Thus, we get the conclusion, as shown in Eq.9.

$$\left| \sum_{j=1}^n \mathbb{E}(\text{DualisticConv}(\mathcal{A}) - \mathcal{A}_j) \right| \quad (3)$$

$$= \left| \mathbb{E} \left(n \sqrt[\gamma]{\sum_{i=1}^n \alpha_i \mathcal{A}_i^\gamma} - \sum_{j=1}^n \mu_j \right) \right| \quad (4)$$

$$\leq \left| \mathbb{E} \left(n \sqrt[\gamma]{\sum_{i=1}^n |\alpha_i| |\mathcal{A}_i^\gamma|} - \sum_{j=1}^n \mu_j \right) \right| \quad (5)$$

$$\leq n \sqrt[\gamma]{\sum_{i=1}^n |\alpha_i| \mathbb{E}(|\mathcal{A}_i^\gamma|)} - \sum_{j=1}^n \mu_j \quad (6)$$

$$= n \sqrt[\gamma]{\sum_{i=1}^n |\alpha_i| \mathbb{E}(|(\bar{\mathcal{A}}_i + \mu_i)^\gamma|)} - \sum_{j=1}^n \mu_j \quad (7)$$

$$\leq 2^{\frac{\gamma-1}{\gamma}} n \sqrt[\gamma]{\sum_{i=1}^n |\alpha_i| \mathbb{E}(|\bar{\mathcal{A}}_i^\gamma|) + |\alpha_i \mu_i^\gamma|} - \sum_{j=1}^n \mu_j \quad (8)$$

$$= 2^{\frac{\gamma-1}{\gamma}} n \sqrt[\gamma]{\sum_{i=1}^n |\alpha_i| (\gamma - 1)!! \nu_i^\gamma + |\alpha_i \mu_i^\gamma|} - \sum_{j=1}^n \mu_j \quad (9)$$

It is worth noting that the upper bound is primarily determined by $\nu_i, i \in [1, n]$, whereas the influence of μ is negligible. This is because, regardless of the value of μ , the expression $2^{\frac{\gamma-1}{\gamma}} n \sqrt[\gamma]{\sum_{i=1}^n |\alpha_i| (\gamma - 1)!! \nu_i^\gamma + |\alpha_i \mu_i^\gamma|} - \sum_{j=1}^n \mu_j$ is always greater than $2^{\frac{\gamma-1}{\gamma}} n \sqrt[\gamma]{\sum_{i=1}^n |\alpha_i| (\gamma - 1)!! \nu_i^\gamma}$, which is solely related to $\nu_i, i \in [1, n]$. This can be proven using the power mean inequality. As a result, the upper bound for the gap is primarily influenced by the standard deviation of amplitudes and positively correlated with it. Consequently, the gap between the latent vector and the original amplitudes for normal distributions is more rigorously constrained, implying that they are easier to reconstruct.

C. Pattern Extraction

The most challenging problem in tackling diverse normal patterns with a unified model is that an anomaly for one

¹The Server Machine Dataset [3]

²A service dataset from one of global top 10 internet company [32]

³A service dataset from one of global top 10 internet company [32]

TABLE II
THE AVERAGE SPECTRUM VARIANCES OF ANOMALOUS AND NORMAL PATTERNS IN DIFFERENT DATASETS.

	SMD ¹	J-D1 ²	J-D2 ³
Anomaly	4.55	12.38	15.64
Normality	3.36	11.74	14.13

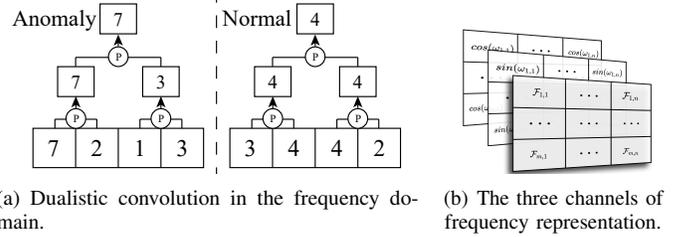


Fig. 4. (a) The dualistic convolution applying to the frequency domain actually picks the prominent deviation in each compression step. (b) The figure shows the three channels of frequency representation in the frequency characterization module. The first channel is the result of Fourier transformation, the second is corresponding *sin* Fourier bases and the third is corresponding *cos* Fourier bases.

normal pattern can be normal for another. To overcome this issue, we propose a pattern extraction mechanism to detect anomalies by the correlation between the data sample and its service normal pattern, instead of the data sample itself, which allows us to handle various normal patterns and reduce computational overhead by capitalizing on the sparsity inherent in the frequency domain. Our pattern extraction mechanism comprises three key components: a context-aware discrete Fourier transformation (DFT) module, a frequency characterization module, and a context-aware inverse discrete Fourier transformation (IDFT) module. In the preprocessing stage, we analyze each service in the frequency domain and identify a normal pattern subspace containing most of the normalities as the service normal pattern subspace by establishing a compact set of dominant Fourier bases for each normal pattern. During both the training and testing phases, we employ the context-aware DFT module to project time series data to its service normal pattern subspace by approximating it with a linear combination of the bases within the relevant normal pattern subspace. This procedure effectively compresses the spectral volume and minimizes computational demands. Subsequently, the frequency characterization module is used to create a tailored frequency representation for the time series data from its projection. After undergoing reconstruction through a dualistic convolution-based auto-encoder, the spectrum is transformed back into time series data using the context-aware IDFT. Furthermore, we theoretically prove that utilizing only the dominant bases for each normal pattern yields superior performance in distinguishing anomalies from normal patterns compared to using the complete spectrum. This is further supported by experimental evidence in Section V-D.

Context-aware DFT and IDFT. It is assumed that each service or server exhibits its unique normal pattern. Consequently, during the preprocessing stage, we process the training

TABLE III
THE EXPECTATION OF AMPLITUDES FOR ANOMALIES AND NORMALITIES
ON DIFFERENT DATASETS.

	SMD	J-D1	J-D2
Anomaly	0.36	0.74	0.81
Normality	0.23	0.72	0.77

dataset in the frequency domain and count the occurrences of each Fourier base as the first k strongest signals across all sliding windows. Subsequently, we select the top k bases with the highest incidence to serve as Fourier bases for their service normal patterns subspace. In both the training and testing phases, the context-aware DFT transforms the time series data exclusively using the bases from the corresponding normal pattern subspace through the DFT process. Likewise, the context-aware IDFT processes the spectrum using only the corresponding bases through the IDFT. Furthermore, we conduct a theoretical comparison of the reconstruction error between anomalies and normal patterns, demonstrating that the context-aware DFT can significantly widen the gap in reconstruction errors between anomalies and normal patterns.

Definition 2 (Spectrum). Given a DFT spectrum of a normal pattern $\mathcal{A}_N(\omega_0) > \dots > \mathcal{A}_N(\omega_n)$, where $\mathcal{A}_N(\omega_i)$ is the amplitude of the signal and ω_i is its corresponding frequency, we compute normalized value of them as follows: $q_N(\omega_i) = \frac{\mathcal{A}_N(\omega_i)}{\sum_{i=1}^n \mathcal{A}_N(\omega_i)}$. The spectrum of anomalies is denoted by $\mathcal{A}_A(\omega_0), \dots, \mathcal{A}_A(\omega_n)$, where ω_i is exactly the ω_i in spectrum of normal pattern. Similarly, the spectrum of anomalies is normalized and denoted by $q_A(\omega_i)$. The normalized spectrum for normalities and anomalies obtained by context-aware DFT are denoted by $\bar{q}_N(\omega_i)$ and $\bar{q}_A(\omega_i)$ respectively.

Definition 3 (Reconstruction error). The reconstruction error of context-aware DFT is defined as the KL divergence between the spectrum obtained by context-aware DFT and the original spectrum, i.e. $\text{KL}(\bar{q}|q)$.

Assumption 1. The anomalies manifest themselves by adding a shift variable to the spectrum of normalities, whose expectation is greater than 0, i.e. $\mathcal{A}_A(\omega_i) = \mathcal{A}_N(\omega_i) + \Delta\mathcal{A}_i$, where $\forall i, \Delta\mathcal{A}_i$ are independent identically distributed and follow an unknown distribution with expectation ΔA , $\Delta A > 0$. It is reasonable to assume the expectation of shift variable is bigger than 0 since it is reported that the anomalies have stronger signals than normalities [32], which implies higher amplitude expectations of anomalies. Moreover, we statistically collect the expectation of anomalies and normalities across three real-world datasets and verify this point, as shown in Table III.

Theorem 2. The gap of reconstruction error between the anomaly and normality is $\log \frac{\sum_{i=1}^k q_N(\omega_i)}{\sum_{i=1}^k q_A(\omega_i)}$.

Proof. Using normality as an example, we derive the expression for its reconstruction error. The expression for an anomaly can be derived in a similar manner. We start by representing \bar{q}_N as shown in Eq. 10. Subsequently, the reconstruction error for it is obtained in Eq. 11. Therefore, the difference in the reconstruction error between anomalies and normal patterns is given by $\text{KL}(\bar{q}_A|q_A) - \text{KL}(\bar{q}_N|q_N) = \log \frac{\sum_{i=1}^k q_N(\omega_i)}{\sum_{i=1}^k q_A(\omega_i)}$.

$$\bar{q}_N(\omega_i) = \begin{cases} \frac{q_N(\omega_i)}{\sum_{i=1}^k q_N(\omega_i)} & , i \leq k \\ 0 & , i > k \end{cases} \quad (10)$$

$$\begin{aligned} \text{KL}(\bar{q}_N|q_N) &= \sum_{j=1}^k \frac{q_N(\omega_j)}{\sum_{i=1}^k q_N(\omega_i)} \log \frac{1}{\sum_{t=1}^k q_N(\omega_t)} \\ &= - \frac{\sum_{j=1}^k q_N(\omega_j)}{\sum_{t=1}^k q_N(\omega_t)} \log \sum_{i=1}^k q_N(\omega_i) \\ &= - \log \sum_{i=1}^k q_N(\omega_i) \end{aligned} \quad (11)$$

Intuitively, the gap is greater than 0 because the numerator $\sum_{i=1}^k q_N(\omega_i)$ represents the first k strongest signals, while the denominator $\sum_{i=1}^k q_A(\omega_i)$ is not guaranteed to have a similar characteristic. We have conducted a further analysis of the condition for achieving a smaller reconstruction error for normal patterns in Corollary 1.

Corollary 1. When $\sum_{i=1}^k q_N(\omega_i) > \frac{k}{n}$, the reconstruction error of normality is smaller than that of anomaly.

Proof. According to assumption 1, the $q_A(\omega_i)$ can be transformed into $\frac{q_N(\omega_i)S + \Delta\mathcal{A}_i}{S + \sum_{i=1}^n \Delta\mathcal{A}_i}$, where $S = \sum_{i=1}^n \mathcal{A}_N(\omega_i)$. Thus, the gap of reconstruction error between anomaly and normality can be transformed to $\log \frac{S + n\Delta A}{S + k\Delta A (\sum_{j=1}^k q_N(\omega_j))^{-1}}$ according to the Law of Large Numbers [48]. When $\sum_{i=1}^k q_N(\omega_i) > \frac{k}{n}$, we can obtain $S + n\Delta A > S + k\Delta A (\sum_{j=1}^k q_N(\omega_j))^{-1}$. As a result, the gap of reconstruction error is bigger than 0.

It is worth noticing that when using trivial DFT and completed spectrum, k is set to n and the reconstruction error gap of anomaly and normality becomes zero, considering that q_A and q_N are normalized values. In contrast, by examining the condition in Corollary 1, it becomes evident that there must be a value of k less than n that yields a reconstruction error gap greater than 0. Thus, compared with using the full spectrum, only using a subset of them widens the reconstruction error gap between the normalities and anomalies. This confirms that anomalies are easier to distinguish when employing the context-aware DFT with a reduced set of bases compared to the standard DFT.

Frequency characterization. The frequency characterization module concatenates the result of context-aware DFT and explicit marked trigonometric bases, as shown in Fig.4(b), where $\mathcal{F}_{i,j}$ denotes the Fourier result of j^{th} base for i^{th} feature dimension, $\omega_{i,j}$ denotes the frequency of j^{th} base for i^{th} feature dimension. Afterward, we use a three-channel convolution to manipulate the concatenated tensors and obtain the frequency representation.

V. EXPERIMENT

We conduct extensive experiments on four real-world anomaly detection datasets and obtain the following conclusions:

- When detecting anomalies for multiple services with kinds of normal patterns by a unified model, MACE achieves

better performance compared with the state-of-the-art methods.

- MACE achieves comparable performance compared with the state-of-the-art methods when the state-of-the-art methods are trained separately for each service and MACE uses a unified model for all the services.
- MACE shows good transferability on unseen datasets.
- MACE consumes obviously less time and memory overhead than the state-of-the-art methods.
- Every module in MACE contributes to its performance.

A. Experiment Setup

The datasets used in this paper contain several subsets, which represent data for different services, servers and detecting sensors. It is assumed that different subsets have different normal patterns. We divide every ten subsets in a dataset as a group. For each group, we train a unified model to detect the anomaly in it.

Datasets. We utilized a selection of datasets, including the widely-recognized Server Machine Dataset (SMD), two cloud service monitoring logs from a top global Internet company (J-D1 and J-D2), the well-established anomaly detection benchmark, Soil Moisture Active Passive (SMAP) and a dataset MC that we collect from one of top ten global cloud providers. As shown in Fig.5, the normal patterns of SMD are the most diverse, while the normal patterns of J-D2 are the most similar. Moreover, SMAP has most point anomalies, while the anomalies in other datasets are lasting.

- Server Machine Dataset (SMD) [3]: SMD spans a 5-week period and originates from a major Internet company, incorporating data from 28 distinct machines. Each machine’s log data, a subset of SMD, is equally divided into training and testing sets. The anomaly ratio in SMD is 4.16%.
- Datasets provided by Jumpstarter (J-D1 and J-D2) [32]: J-D1 and J-D2 are two datasets gathered from a top global Internet company. Each dataset includes logs of 19 metrics from 30 services, with each service’s log data forming a subset in J-D1 and J-D2. The anomaly ratios for J-D1 and J-D2 are 5.25% and 20.26%, respectively.
- Soil Moisture Active Passive (SMAP) [34]: SMAP comprises real spacecraft telemetry data and anomalies from the Soil Moisture Active Passive satellite, featuring an anomaly ratio of 13.13%.
- MC: MC consists of the monitoring data from 25 services for 15 days, whose anomaly ratio is 3.6% and contains substantial point anomalies.

Hyperparameter. The important hyperparameters of MACE are shown in Table IV, where m denotes the number of basis in a subset, γ_f denotes γ for the dualistic convolution in the frequency domain, γ_t denotes the one in the time domain, σ_f denotes the scaling factor of dualistic convolution in the frequency domain and σ_t denotes the one in the time domain.

Baselines. We conducted a comprehensive comparison of MACE with several state-of-the-art methods, including DCdetector [6], AnomalyTransformer [4], DVGCRN [5], JumpStarter

TABLE IV
HYPERPARAMETERS OF MACE.

Hyperparameter	Value	Hyperparameter	Value
m	20	γ_f in SMD	7
γ_t in SMD	11	γ_f in J-D1	11
γ_t in J-D1	11	γ_f in J-D2	13
γ_t in J-D2	13	γ_f in SMAP	13
γ_t in SMAP	13	σ_f in SMD	5
σ_t in SMD	5	σ_f in J-D1	5
σ_t in J-D1	5	σ_f in J-D2	5
σ_t in J-D2	5	σ_f in SMAP	5
σ_t in SMAP	7	kernel length	5
window size	40	learning rate	0.001

[32], OmniAnomaly [3], and MSCRED [7]. To assess its diverse pattern generalization capabilities, we introduced two additional methods: TranAD, a meta-learning-based approach [8], and ProS, a transfer-learning-based method [49]. Furthermore, to evaluate its computational efficiency, we compared MACE with the classical anomaly detection method VAE [50]. Due to space constraints, some figures represent methods using the first two letters of their names as a shorthand.

- DCdetector (DC) [6]: DCdetector is a latest cutting-edge anomaly detection method. It employs a unique dual attention asymmetric design to establish a permuted environment and leverages pure contrastive loss to guide the learning process. This enables the model to learn a permutation-invariant representation with superior discrimination abilities.
- AnomalyTransformer (An) [4]: Anomaly Transformer stands out as one of the most cutting-edge methods, harnessing the formidable capabilities of transformers to model point-wise representation and pair-wise associations through an innovative anomaly-attention mechanism.
- DVGCRN (DV) [5]: DVGCRN stands out as another cutting-edge anomaly detection method, effectively modeling fine-grained spatial and temporal correlations in multivariate time series. It achieves a precise posterior approximation of latent variables, contributing to a robust representation of multivariate time series data.
- JumpStarter (Ju) [32]: JumpStarter stands out as a cutting-edge anomaly detection method, equipped with a shape-based clustering and an outlier-resistant sampling algorithm. This combination ensures a rapid initialization and high F1 score performance.
- OmniAnomaly (Om) [3]: OmniAnomaly is a widely acknowledged anomaly detection method, employing a stochastic variable connection and planar normalizing flow to robustly capture the representation of normal multivariate time series data.
- MSCRED (MS) [7]: MSCRED is a highly acclaimed method known for its ability to detect anomalies across various scales and pinpoint root causes through the utilization of multi-scale signature matrices.
- TranAD (Tr) [8]: TranAD, a meta-learning-based anomaly detection method, represents one of the latest cutting-edge approaches. It excels in learning a robust initialization for anomaly detection models, demonstrating excellent

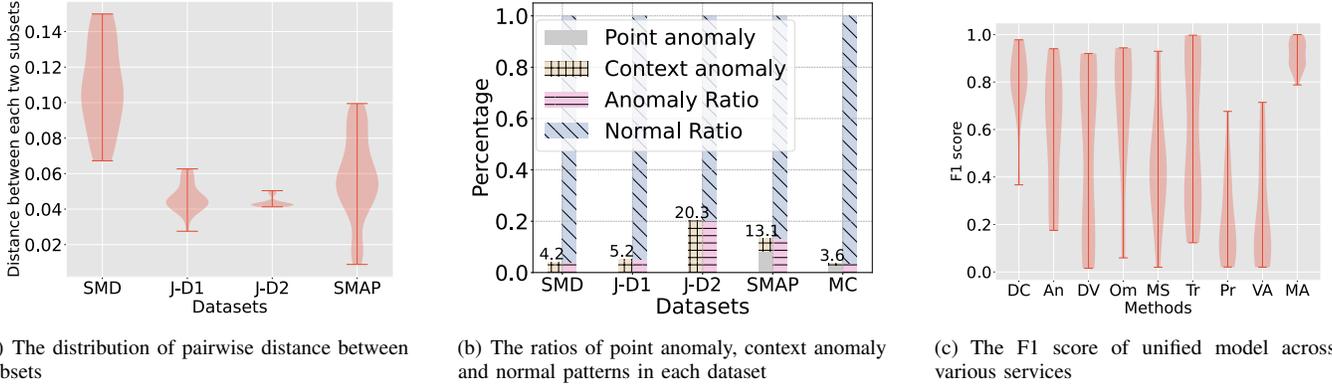


Fig. 5. (a) We first use kernel density estimation to estimate the distribution of each subset. Subsequently, we compute KL divergence between each pair of subsets in a training group. The figure shows the distribution of the KL divergences of different datasets. (b) The figure shows the point anomaly and normal pattern ratios in each dataset. (c) All the methods train a unified model for every ten services. The figure shows their F1 score across different services.

generalization across diverse normal patterns.

- ProS (Pr) [49]: ProS introduces a zero-shot methodology, capable of inferring in the target domains without the need for re-training. This is achieved through the introduction of latent domain vectors, serving as latent representations of the domains.
- VAE (VA) [50]: VAE, a widely recognized and classical anomaly detection method, serves as a foundational framework for numerous state-of-the-art approaches. It introduces low computational and memory overhead, contributing to its popularity.

Metrics. We use three of the most widely-used metrics to evaluate the performance of MACE and baseline methods as many prominent anomaly detection papers [5], [6], [8], [32] do: the precision, recall and F1 score. The definitions of these metrics are given in Eq. 12 - Eq. 14, where TP , FP and FN denote true positive, false positive and false negative respectively.

$$\text{Precision} = \frac{TP}{TP + FP} \quad (12)$$

$$\text{Recall} = \frac{TP}{FN + TP} \quad (13)$$

$$\text{F1} = 2 * \frac{\text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}} \quad (14)$$

B. Prediction Accuracy

In this subsection, we conduct extensive experiments to validate that MACE consistently achieves the best F1 score compared to baselines when distinguishing anomalies from different normal patterns with a unified model. Furthermore, in comparison to baselines that tailor a unique model for each subset, MACE demonstrates competitive performance with a unified model. Additionally, owing to the memory-guided pattern extraction method, MACE exhibits commendable performance on previously unseen normal patterns.

Adaptability to multiple normal patterns. We assume that various subsets in the four datasets contain distinct normal patterns, representing logs for different servers (SMD), services (J-D1 and J-D2), and data for different detector channels

(SMAP). During the training stage, every ten subsets in a dataset are grouped together and utilized to train a unified model for both MACE and baselines. In the testing stage, the trained model is applied to detect anomalies in each corresponding testing subset independently. The average metrics for different subsets are presented in Table V, where the best results are highlighted in bold, and the second-best results are underlined. Since JumpStarter is a signal-based method, multiple normal pattern training is not applicable to it, and thus, JumpStarter is excluded from this analysis. As indicated in Table V, despite occasional deviations, MACE achieves the best performance when detecting multiple normal patterns with a unified model. Moreover, MACE consistently achieves the best F1 score compared with all the baselines across the four datasets. Furthermore, the improvement is substantial: MACE increases the F1 score by an average of 8.7% compared to the best baseline performance. As illustrated in Fig. 5(a), the subsets in SMD exhibit significant differences from each other, where MACE shows a distinct advantage. Intuitively, since the normal patterns in J-D2 are very similar to each other, most methods perform well on this dataset and the advantage of MACE is not as obvious as the one on the former dataset. Additionally, MACE attains a high F1 score on SMAP. Considering that most anomalies in SMAP are point anomalies, this result verifies the effectiveness of dualistic convolution in the time domain, which extends the detection capabilities for short-term anomalies and makes them easier to identify.

Moreover, to verify the unified model can work well on each service, we plot the F1 score for different services when detecting anomalies with a unified model on the Server Machine Dataset in Fig. 5(c). As shown in Fig. 5(c), the performance of MACE centered around a pretty well average value, while performances of other methods vary around a broad range across different services. This verifies MACE can capture diverse normal patterns well with a unified model.

Competitive performance compared to customizing a unique model. In this experiment, MACE employs a unified model for every ten different normal patterns, while baselines customize a unique model for each subset. As depicted in Table

TABLE V
THE PERFORMANCE OF MACE AND BASELINES WHEN DETECTING ANOMALIES FOR MULTIPLE PATTERNS BY A UNIFIED MODEL.

	SMD			J-D1			J-D2			SMAP		
	Precision	Recall	F1									
DCdetector	0.680	0.672	0.669	0.709	0.583	0.626	0.956	0.897	0.923	0.594	0.613	0.597
AnomalyTransformer	0.439	0.947	0.562	0.519	0.945	0.639	0.824	0.981	0.891	0.610	0.947	0.699
DVGCRN	0.481	0.766	0.481	0.344	0.737	0.421	0.695	0.867	0.742	0.475	0.979	0.549
OmniAnomaly	0.674	0.829	0.713	0.957	0.868	0.899	0.948	0.932	0.938	0.789	0.984	0.819
MSCRED	0.444	0.562	0.407	0.880	0.806	0.819	0.927	0.944	0.932	0.838	1.000	0.884
TranAD	0.617	0.616	0.471	0.198	0.631	0.258	0.729	0.952	0.797	0.275	0.577	0.291
ProS	0.153	0.785	0.214	0.505	0.731	0.534	0.796	0.861	0.805	0.412	0.973	0.468
VAE	0.221	0.689	0.246	0.377	0.796	0.425	0.566	0.909	0.665	0.470	0.983	0.557
MACE	0.964	0.870	0.910	0.893	0.984	0.934	0.938	0.989	0.961	0.958	1.000	0.977

TABLE VI
MACE USES A UNIFIED MODEL FOR DIFFERENT PATTERNS, WHILE BASELINES CUSTOMIZE A UNIQUE MODEL FOR EACH NORMAL PATTERN.

	SMD			J-D1			J-D2			SMAP		
	Precision	Recall	F1									
DCdetector	0.836	0.911	0.872	0.766	0.744	0.748	0.956	0.880	0.913	0.956	0.989	0.970
AnomalyTransformer	0.894	0.955	0.923	0.520	0.918	0.645	0.818	0.998	0.896	0.941	0.994	0.967
DVGCRN	0.950	0.883	0.915	0.395	0.806	0.479	0.711	0.852	0.723	0.916	0.920	0.914
JumpStarter	0.904	0.943	0.923	0.921	0.945	0.933	0.941	0.996	0.968	0.471	0.995	0.526
OmniAnomaly	0.695	0.877	0.728	0.891	0.940	0.905	0.945	0.974	0.958	0.713	0.963	0.744
MSCRED	0.746	0.744	0.716	0.975	0.830	0.889	0.949	0.969	0.958	0.872	1.000	0.923
TranAD	0.926	0.997	0.961	0.251	0.918	0.349	0.754	0.965	0.817	0.804	1.000	0.892
ProS	0.146	0.822	0.206	0.422	0.767	0.506	0.763	0.921	0.821	0.447	0.973	0.509
VAE	0.286	0.585	0.255	0.334	0.866	0.385	0.702	0.890	0.763	0.579	0.973	0.648
MACE	0.964	0.870	0.910	0.893	0.984	0.934	0.938	0.989	0.961	0.958	1.000	0.977

TABLE VII
MACE USES A UNIFIED MODEL FOR DIFFERENT PATTERNS, WHILE BASELINES CUSTOMIZE A UNIQUE MODEL FOR EACH NORMAL PATTERN ON MC

	DCdetector	AnomalyTrans	DVGCRN	OmniAnomaly	MSCRED	TranAD	ProS	VAE	JumpStarter	MACE
Precision	0.984	1.000	0.125	0.681	0.841	0.784	0.681	0.583	0.473	0.908
Recall	0.728	0.870	0.304	0.981	0.981	0.989	1.000	0.750	0.393	0.994
F1 Score	0.806	0.923	0.147	0.782	0.878	0.864	0.772	0.639	0.393	0.941

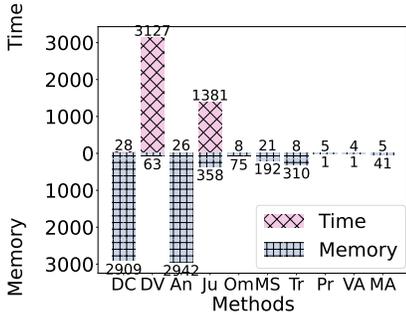
TABLE VIII
THE PERFORMANCE OF MACE AND BASELINES FOR UNSEEN NORMAL PATTERNS.

	SMD			J-D1			J-D2			SMAP		
	Precision	Recall	F1									
DCdetector	0.681	0.685	0.681	0.798	0.771	0.781	0.948	0.857	0.891	0.700	0.760	0.724
AnomalyTransformer	0.490	0.916	0.622	0.555	0.948	0.667	0.838	0.981	0.899	0.586	1.000	0.678
DVGCRN	0.125	0.798	0.173	0.388	0.894	0.478	0.619	0.893	0.664	0.444	1.000	0.525
OmniAnomaly	0.686	0.780	0.701	0.976	0.824	0.880	0.932	0.953	0.941	0.735	0.986	0.794
MSCRED	0.418	0.593	0.409	0.828	0.818	0.806	0.931	0.952	0.939	0.839	1.000	0.896
TranAD	0.255	0.643	0.265	0.127	0.546	0.198	0.516	0.659	0.546	0.205	1.000	0.302
ProS	0.154	0.808	0.215	0.475	0.770	0.564	0.789	0.933	0.855	0.412	0.979	0.469
VAE	0.193	0.789	0.270	0.339	0.875	0.386	0.661	0.884	0.721	0.433	0.979	0.500
MACE	0.915	0.835	0.863	0.972	0.829	0.885	0.963	0.968	0.964	0.954	0.996	0.973

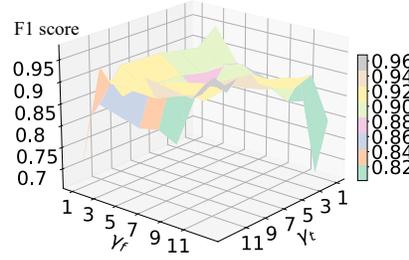
VI and Table VII, MACE achieves comparable performance with the strongest state-of-the-art methods. It is noteworthy that MACE captures ten different patterns simultaneously with a single model, a factor that generally hinders model performance [10], while baselines customize a unique model for each normal pattern. The negative impact of multiple normal patterns is further evident when comparing Table V and Table VI: when normal patterns are diverse (e.g., in SMD), the baselines in Table VI, where they tailor a unique model for each normal pattern, exhibit big strength compared to their performance in

Table V, where they learn a unified model for multiple normal patterns. In contrast, when normal patterns are similar (e.g., in J-D2), the baseline performance gaps between Tables V and VI are narrow. From this comparison, it can be concluded that the diversity of normal patterns hinders model performance when using a unified model for multiple patterns. Thus, it is tolerable for MACE to exhibit a somewhat lower F1 score on SMD, considering that the normal patterns of SMD are the most diverse among the four datasets.

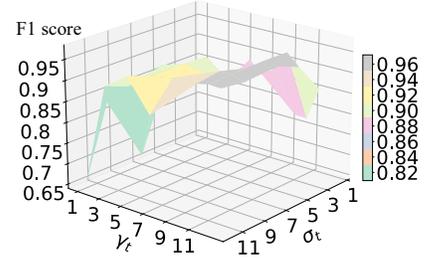
MACE performance on unseen normal patterns. As



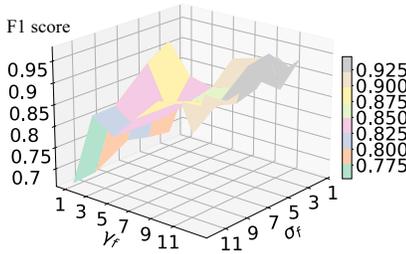
(a) Time (s) and memory overhead (*10KB).



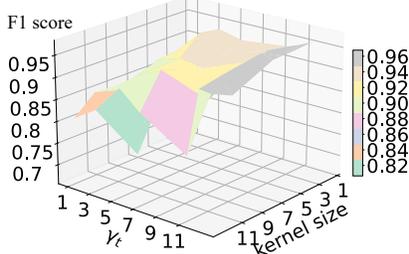
(b) The impact of γ_t and γ_f on MACE performance.



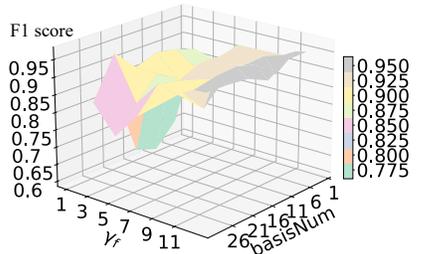
(c) The impact of γ_t and σ_t on MACE performance.



(d) The impact of γ_f and σ_f on MACE performance.



(e) The impact of kernel size and γ_t on MACE performance.



(f) The impact of number of basis in a subset and γ_f on MACE performance.

Fig. 6. Due to the space limitation, we use the first two letters in method names as a shorthand. γ_t and γ_f represent the powers in dualistic convolution for the time domain and frequency domain respectively. σ_t and σ_f represent the scaling factors in dualistic convolution for the time domain and frequency domain respectively. (a) The time and memory overhead of MACE and baselines. (b) The F1 score of MACE for grid search of γ_t and γ_f . (c) The F1 score of MACE for grid search of γ_t and σ_t . (d) The F1 score of MACE for grid search of γ_f and σ_f . (e) The F1 score of MACE for grid search of dualistic convolution kernel size in the time domain and γ_t . (f) The F1 score of MACE for grid search of the number of bases in a subset and γ_f .

mentioned earlier, different subsets are assumed to represent different normal patterns, and every ten subsets in a dataset are divided into a group. MACE and all the baselines are trained on one group and tested on another. The results are presented in Table VIII, where the best performances are bolded, and the second-best performances are underlined. Since JumpStarter is a signal-based method, training on one group while testing on another is not applicable to it, and thus, it is not included in the table. As shown in Table VIII, MACE consistently achieves the highest F1 score on the four datasets. The performance of MACE when the normal patterns are diverse, such as in SMD, is lower than when the normal patterns are similar, such as in J-D2. When the distance between different normal patterns is small, MACE can achieve similar F1 scores to those when MACE is trained and tested on the same group (i.e., the performance on J-D2 and SMAP).

C. Efficiency Analysis

We evaluated both time and memory overhead on a server equipped with a configuration comprising 32 Intel(R) Xeon(R) CPU E5-2620 @ 2.10GHz CPUs and 2 K80 GPUs. For neural network-based methods, we employed a profiling tool to assess their memory overhead. In the case of JumpStarter, a signal-based method, we record its maximum memory consumption during the inference process. The time overhead was calculated based on the training time of each method on a subset group of the SMD dataset. The results, depicted in Fig. 6(a), reveal

that MACE’s time overhead is competitive with some very simple methods, such as VAE and ProS based on VAE, while MACE’s F1 scores significantly surpass the ones of them across all four datasets. Regarding memory overhead, MACE’s value is higher than that of a two-layer VAE and ProS based on a two-layer VAE. However, MACE’s memory overhead is considerably lower than that of other deep neural networks. These findings underscore MACE’s efficiency in terms of both time and memory usage, positioning it favorably among deep neural-network-based methods and showcasing superior performance in terms of anomaly detection as evidenced by its higher F1 scores across diverse datasets.

D. Ablation Study

We conducted experiments to assess the effectiveness of individual modules within MACE by removing them individually. When the context-aware Discrete Fourier Transform (DFT) and Inverse DFT (IDFT) modules were removed, they were replaced with conventional DFT and IDFT. When making ablation experiments for other modules, we compare the completed MACE with the one that has removed them. The results are displayed in Table IX, where “Dualistic Convolution (F)” and “Dualistic Convolution (T)” correspond to dualistic convolution in the frequency and time domains, respectively. As depicted in Table IX, the complete MACE model exhibits considerable superiority over its variants. It is worth noting that when context-aware DFT and IDFT are substituted with vanilla

TABLE IX
THE PERFORMANCES OF MACE WHEN REMOVING DIFFERENT MODULES

Remove module	SMD			J-D1			J-D2			SMAP		
	Precision	Recall	F1									
Context-aware DFT & IDFT	0.762	0.813	0.762	0.624	0.887	0.689	<u>0.958</u>	0.953	0.953	0.775	1.000	0.831
Dualistic Convolution (F)	0.187	0.943	0.184	<u>0.855</u>	0.854	0.820	0.857	0.933	0.886	0.681	0.990	0.713
Dualistic Convolution (T)	0.046	0.842	0.084	0.089	<u>0.946</u>	0.152	0.208	0.440	0.250	0.682	0.989	0.720
Frequency Characterization	<u>0.894</u>	0.860	<u>0.868</u>	0.838	0.940	<u>0.857</u>	0.970	<u>0.980</u>	0.975	<u>0.944</u>	1.000	<u>0.967</u>
Pattern extraction	0.714	0.702	0.696	0.667	0.914	0.740	0.957	0.953	0.954	0.770	0.970	0.797
MACE	0.964	0.870	0.910	0.893	0.984	0.934	0.938	0.989	<u>0.961</u>	0.958	1.000	0.977

counterparts, MACE’s performance takes a sharp nosedive. The computational and memory overhead of vanilla DFT and IDFT increases because they introduce more Fourier bases, yet the performance deteriorates. This observation aligns with our earlier theoretical analysis. Moreover, this experiment underscores the effectiveness of the diverse normal pattern adaptability facilitated by the pattern extraction mechanism. This module significantly enhances performance on SMD, characterized by diverse normal patterns, while showing marginal improvement on J-D2, where the normal patterns are similar. Similarly, the frequency characterization module makes a substantial contribution on SMD but is of limited use on J-D2, given its multi-pattern extraction nature. In summary, the module-by-module evaluation reaffirms the crucial role played by various components in MACE and their impact on anomaly detection performance across different datasets.

E. Hyperparameter Study

We employed grid search to investigate the influence of critical hyperparameters on the performance of MACE. Fig.6(b)-Fig.6(f) present the F1 scores corresponding to different combinations of pairwise hyperparameters. The search ranges for γ in the time domain and frequency domain, σ in the time domain and frequency domain, kernel size, and the number of Fourier bases in Context-aware DFT and IDFT were set to $\{1, 3, 5, 7, 11, 12, 13\}$, $\{3, 5, 7, 10, 12\}$, $\{3, 5, 7, 11, 13\}$, and $\{5, 10, 15, 20, 25, 30\}$ respectively.

Impact of γ_t and γ_f : When γ_t and γ_f are set to 1, dualistic convolution degenerates into a standard convolution, essentially nullifying its contribution. Consequently, MACE’s performances with γ_t and γ_f set to 1 are unsatisfied. In general, the performance of MACE improves as γ_t and γ_f increase, as depicted in Fig.6(c) and Fig.6(d). However, it’s important to note that γ cannot grow infinitely, as excessively large γ values can lead to gradient explosions. Thus, setting γ within the search space mentioned above is a safe approach.

Impact of σ_t and σ_f : These scaling factors are introduced to mitigate gradient explosions. As demonstrated in Fig.6(c)-Fig.6(d), MACE’s performance remains stable across various values of σ .

Impact of dualistic convolution kernel size in the time domain: Intuitively, the performance of MACE initially improves and then declines as the kernel size increases, as shown in Figure 6(e). That is because when the kernel size increases from a small value, it makes anomalies more prominent and

easier to detect. However, when the kernel size becomes excessively large, the dualistic convolution in the time domain distorts the original time series and detrimentally affects model performance.

Impact of the number of Fourier bases in context-aware DFT and IDFT: The performance of MACE generally follows an increasing-then-decreasing pattern as the number of bases grows. As analyzed theoretically in Section.IV-C, when the number of bases increases from a small value, both the reconstruction of normal patterns and anomalies improve, but the enhancement in normality reconstruction is more pronounced. However, when the number of bases becomes relatively large, the improvement in normality reconstruction becomes marginal, while the effect on anomaly reconstruction becomes significant. Therefore, the performance initially increases and then declines with the growth of the number of bases.

VI. CONCLUSION

In this work, we address the challenges of detecting anomalies from diverse normal patterns with a unified and efficient model as well as improve the short-term anomaly sensitivity by proposing MACE. MACE exhibits three key characteristics: (i) a pattern extraction mechanism allowing the model to detect anomalies by the correlation between the data sample and its service normal pattern and adapt to the diverse nature of normal patterns; (ii) a dualistic convolution mechanism that amplifies anomalies in the time domain and hinders the reconstruction in the frequency domain; (iii) the utilization of the inherent sparsity and parallelism of the frequency domain to enhance model efficiency. We substantiate our approach through both mathematical analysis and extensive experiments, demonstrating that selecting a subset of Fourier bases based on normal patterns yields superior performance compared to utilizing the complete spectrum. Comprehensive experiments confirm MACE’s proficiency in effectively handling diverse normal patterns, showcasing optimal performance with high efficiency when benchmarked against state-of-the-art methods.

ACKNOWLEDGMENT

This work was supported by the Key Research Project of Zhejiang Province under Grant 2022C01145, the National Science Foundation of China under Grants 62125206 and U20A20173, the Guangzhou-HKUST(GZ) Joint Funding Program under Grant 2024A03J0620, and in part by Alibaba Group through Alibaba Research Intern Program.

REFERENCES

- [1] G. Zhang, C. Li, K. Zhou, L. Liu, C. Zhang, W. Chen, H. Fang, B. Cheng, J. Yang, and J. Xing, "Dbcatcher: A cloud database online anomaly detection system based on indicator correlation," in *2023 IEEE 39th International Conference on Data Engineering (ICDE)*. IEEE, 2023, pp. 1126–1139.
- [2] F. Chen, Z. Qing, Y. Zhang, S. Deng, Y. Xiao, G. Pang, and Q. Wen, "Lara: A light and anti-overfitting retraining approach for unsupervised time series anomaly detection," *arXiv preprint arXiv:2310.05668*, 2023.
- [3] Y. Su, Y. Zhao, C. Niu, R. Liu, W. Sun, and D. Pei, "Robust anomaly detection for multivariate time series through stochastic recurrent neural network," in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2019, pp. 2828–2837.
- [4] J. Xu, H. Wu, J. Wang, and M. Long, "Anomaly transformer: Time series anomaly detection with association discrepancy," in *The Tenth International Conference on Learning Representations, ICLR 2022*. OpenReview.net, 2022.
- [5] W. Chen, L. Tian, B. Chen, L. Dai, Z. Duan, and M. Zhou, "Deep variational graph convolutional recurrent network for multivariate time series anomaly detection," in *International Conference on Machine Learning, ICML 2022*, ser. Proceedings of Machine Learning Research, vol. 162, 2022, pp. 3621–3633.
- [6] Y. Yang, C. Zhang, T. Zhou, Q. Wen, and L. Sun, "Dcdetector: Dual attention contrastive representation learning for time series anomaly detection," in *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2023*, A. K. Singh, Y. Sun, L. Akoglu, D. Gunopulos, X. Yan, R. Kumar, F. Ozcan, and J. Ye, Eds. ACM, 2023, pp. 3033–3045.
- [7] C. Zhang, D. Song, Y. Chen, X. Feng, C. Lumezanu, W. Cheng, J. Ni, B. Zong, H. Chen, and N. V. Chawla, "A deep neural network for unsupervised anomaly detection and diagnosis in multivariate time series data," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, no. 01, 2019, pp. 1409–1416.
- [8] S. Tuli, G. Casale, and N. R. Jennings, "TranAD: Deep transformer networks for anomaly detection in multivariate time series data," *Proc. VLDB Endow.*, vol. 15, no. 6, pp. 1201–1214, 2022.
- [9] Y. Li, W. Chen, B. Chen, D. Wang, L. Tian, and M. Zhou, "Prototype-oriented unsupervised anomaly detection for multivariate time series," in *International Conference on Machine Learning, ICML 2023*, ser. Proceedings of Machine Learning Research, vol. 202, 2023, pp. 19407–19424.
- [10] H. Park, J. Noh, and B. Ham, "Learning memory-guided normality for anomaly detection," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 14372–14381.
- [11] J. Chung, K. Kastner, L. Dinh, K. Goel, A. C. Courville, and Y. Bengio, "A recurrent latent variable model for sequential data," *Advances in neural information processing systems*, vol. 28, 2015.
- [12] S. Schmidl, P. Wenig, and T. Papenbrock, "Anomaly detection in time series: a comprehensive evaluation," *Proceedings of the VLDB Endowment*, vol. 15, no. 9, pp. 1779–1797, 2022.
- [13] S. Ramaswamy, R. Rastogi, and K. Shim, "Efficient algorithms for mining outliers from large data sets," in *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, 2000, pp. 427–438.
- [14] X. Wang, J. Lin, N. Patel, and M. Braun, "Exact variable-length anomaly detection algorithm for univariate and multivariate time series," *Data Mining and Knowledge Discovery*, vol. 32, pp. 1806–1844, 2018.
- [15] V. Gómez-Verdejo, J. Arenas-García, M. Lazaro-Gredilla, and Á. Navia-Vazquez, "Adaptive one-class support vector machine," *IEEE Transactions on Signal Processing*, vol. 59, no. 6, pp. 2975–2981, 2011.
- [16] P. Boniol, M. Linardi, F. Roncallo, T. Palpanas, M. Meftah, and E. Remy, "Unsupervised and scalable subsequence anomaly detection in large data series," *The VLDB Journal*, pp. 1–23, 2021.
- [17] Y. Zhang, Z. Guan, H. Qian, L. Xu, H. Liu, Q. Wen, L. Sun, J. Jiang, L. Fan, and M. Ke, "CloudRCA: A root cause analysis framework for cloud computing platforms," in *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2021, pp. 4373–4382.
- [18] A. Siffer, P.-A. Fouque, A. Termier, and C. Largouet, "Anomaly detection in streams with extreme value theory," in *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, 2017, pp. 1067–1075.
- [19] S. Subramaniam, T. Palpanas, D. Papadopoulos, V. Kalogeraki, and D. Gunopulos, "Online outlier detection in sensor data using non-parametric models," in *Proceedings of the 32nd international conference on Very large data bases*, 2006, pp. 187–198.
- [20] H. Ren, B. Xu, Y. Wang, C. Yi, C. Huang, X. Kou, T. Xing, M. Yang, J. Tong, and Q. Zhang, "Time-series anomaly detection service at microsoft," in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2019, pp. 3009–3017.
- [21] M. Thill, W. Konen, and T. Bäck, "Online adaptable time series anomaly detection with discrete wavelet transforms and multivariate gaussian distributions," *Archives of Data Science, Series A (Online First)*, vol. 5, no. 1, p. 17, 2018.
- [22] M. Thill, W. Konen, and T. Bäck, "Time series anomaly detection with discrete wavelet transforms and maximum likelihood estimation," in *Intern. Conference on Time Series (ITISE)*, vol. 2, 2017, pp. 11–23.
- [23] J. Gao, X. Song, Q. Wen, P. Wang, L. Sun, and H. Xu, "RobustTAD: Robust time series anomaly detection via decomposition and convolutional neural networks," *arXiv preprint arXiv:2002.09545*, 2020.
- [24] X. Chen, L. Deng, F. Huang, C. Zhang, Z. Zhang, Y. Zhao, and K. Zheng, "Daemon: Unsupervised anomaly detection and interpretation for multivariate time series," in *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 2021, pp. 2225–2230.
- [25] C. Zhang, T. Zhou, Q. Wen, and L. Sun, "TFAD: A decomposition time series anomaly detection architecture with time-frequency analysis," in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 2022, pp. 2497–2507.
- [26] L. Li, J. Yan, Q. Wen, Y. Jin, and X. Yang, "Learning robust deep state space for unsupervised anomaly detection in contaminated time-series," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [27] A. Goodge, B. Hooi, S.-K. Ng, and W. S. Ng, "Lunar: Unifying local outlier detection methods via graph neural networks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 6, 2022, pp. 6737–6745.
- [28] M. Jin, H. Y. Koh, Q. Wen, D. Zambon, C. Alippi, G. I. Webb, I. King, and S. Pan, "A survey on graph neural networks for time series: Forecasting, classification, imputation, and anomaly detection," *arXiv preprint arXiv:2307.03759*, 2023.
- [29] D. R. Choffnes, F. E. Bustamante, and Z. Ge, "Crowdsourcing service-level network event monitoring," in *Proceedings of the ACM SIGCOMM 2010 Conference*, 2010, pp. 387–398.
- [30] S. Guha, N. Mishra, G. Roy, and O. Schrijvers, "Robust random cut forest based anomaly detection on streams," in *International conference on machine learning*. PMLR, 2016, pp. 2712–2721.
- [31] G. Pang, K. M. Ting, and D. Albrecht, "Lesinn: Detecting anomalies by identifying least similar nearest neighbours," in *2015 IEEE international conference on data mining workshop (ICDMW)*. IEEE, 2015, pp. 623–630.
- [32] M. Ma, S. Zhang, J. Chen, J. Xu, H. Li, Y. Lin, X. Nie, B. Zhou, Y. Wang, and D. Pei, "Jump-Starting multivariate time series anomaly detection for online service systems," in *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, 2021, pp. 413–426.
- [33] N. Zhao, J. Zhu, Y. Wang, M. Ma, W. Zhang, D. Liu, M. Zhang, and D. Pei, "Automatic and generic periodicity adaptation for kpi anomaly detection," *IEEE Transactions on Network and Service Management*, vol. 16, no. 3, pp. 1170–1183, 2019.
- [34] K. Hundman, V. Constantinou, C. Laporte, I. Colwell, and T. Soderstrom, "Detecting spacecraft anomalies using lstms and nonparametric dynamic thresholding," in *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 2018, pp. 387–395.
- [35] B. Zong, Q. Song, M. R. Min, W. Cheng, C. Lumezanu, D. Cho, and H. Chen, "Deep autoencoding gaussian mixture model for unsupervised anomaly detection," in *6th International Conference on Learning Representations, ICLR 2018*, 2018.
- [36] W. Grathwohl, K.-C. Wang, J.-H. Jacobsen, D. Duvenaud, M. Norouzi, and K. Swersky, "Your classifier is secretly an energy based model and you should treat it like one," in *8th International Conference on Learning Representations, ICLR 2020*, 2020.
- [37] L. Ruff, N. Görnitz, L. Deecke, S. A. Siddiqui, R. A. Vandermeulen, A. Binder, E. Müller, and M. Kloft, "Deep one-class classification," in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, ser. Proceedings of Machine Learning Research, vol. 80, 2018, pp. 4390–4399.

- [38] L. Shen, Z. Li, and J. Kwok, "Timeseries anomaly detection using temporal hierarchical one-class network," *Advances in Neural Information Processing Systems*, vol. 33, pp. 13 016–13 026, 2020.
- [39] H. Xu, W. Chen, N. Zhao, Z. Li, J. Bu, Z. Li, Y. Liu, Y. Zhao, D. Pei, Y. Feng *et al.*, "Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications," in *Proceedings of the 2018 world wide web conference*, 2018, pp. 187–196.
- [40] J. Audibert, P. Michiardi, F. Guyard, S. Marti, and M. A. Zuluaga, "Usad: Unsupervised anomaly detection on multivariate time series," in *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 2020, pp. 3395–3404.
- [41] A. Deng and B. Hooi, "Graph neural network-based anomaly detection in multivariate time series," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 35, no. 5, 2021, pp. 4027–4035.
- [42] Y. Yang and T. M. Hospedales, "Deep multi-task representation learning: A tensor factorisation approach," in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [43] Y. Liu, X. Yang, D. Xie, X. Wang, L. Shen, H. Huang, and N. Balasubramanian, "Adaptive activation network and functional regularization for efficient and flexible deep multi-task learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 4924–4931.
- [44] J. Ma, Z. Zhao, J. Chen, A. Li, L. Hong, and E. H. Chi, "Snr: Sub-network routing for flexible parameter sharing in multi-task learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 216–223.
- [45] J.-H. Kim, "Further improvement of jensen inequality and application to stability of time-delayed systems," *Automatica*, vol. 64, pp. 121–125, 2016.
- [46] G. Wang, X. Zhang, and Y. Chu, "A power mean inequality involving the complete elliptic integrals," *The Rocky Mountain Journal of Mathematics*, vol. 44, no. 5, pp. 1661–1667, 2014.
- [47] E. Artin, *The gamma function*. Courier Dover Publications, 2015.
- [48] P.-L. Hsu and H. Robbins, "Complete convergence and the law of large numbers," *Proceedings of the national academy of sciences*, vol. 33, no. 2, pp. 25–31, 1947.
- [49] A. Kumagai, T. Iwata, and Y. Fujiwara, "Transfer anomaly detection by inferring latent domain representations," *Advances in neural information processing systems*, vol. 32, 2019.
- [50] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *2nd International Conference on Learning Representations, ICLR 2014*, 2014.