

Content-Aware Balanced Spectrum Encoding in Masked Modeling for Time Series Classification

Yudong Han^{1,2*}, Haocong Wang^{1*}, Yupeng Hu^{1†}, Yongshun Gong¹, Xuemeng Song³, Weili Guan⁴

¹School of Software, Shandong University,

²Beijing Institute of Technology,

³School of Computer Science and Technology, Shandong University,

⁴Harbin Institute of Technology (Shenzhen)

{hanyudong.sdu, sduwhc, sxmustc, honeyguan}@gmail.com, {huyupeng, ysgong}@sdu.edu.cn

Abstract

Due to the superior ability of global dependency, transformer and its variants have become the primary choice in Masked Time-series Modeling (MTM) towards time-series classification task. In this paper, we experimentally analyze that existing transformer-based MTM methods encounter with two under-explored issues when dealing with time series data: (1) they encode features by performing long-dependency ensemble averaging, which easily results in rank collapse and feature homogenization as the layer goes deeper; (2) they exhibit distinct priorities in fitting different frequency components contained in the time-series, inevitably leading to spectrum energy imbalance of encoded feature. To tackle these issues, we propose an auxiliary content-aware balanced decoder (CBD) to optimize the encoding quality in the spectrum space within masked modeling scheme. Specifically, the CBD iterates on a series of fundamental blocks, and thanks to two tailored units, each block could progressively refine the masked representation via adjusting the interaction pattern based on local content variations of time-series and learning to recalibrate the energy distribution across different frequency components. Moreover, a dual-constraint loss is devised to enhance the mutual optimization of vanilla decoder and our CBD. Extensive experimental results on ten time-series classification datasets show that our method nearly surpasses a bunch of baselines. Meanwhile, a series of explanatory results are showcased to sufficiently demystify the behaviors of our method.

Introduction

Time-series representation learning has emerged as a fundamental and preposed task for time-series analysis. Different from supervised learning (Han et al. 2023b, 2021; Hu et al. 2021a,b, 2024), which heavily relies on human-labeled ground truth, self-supervised learning (Nie et al. 2022; Tonekaboni, Eytan, and Goldenberg 2021; Ozyurt, Feuerriegel, and Zhang 2022) has shown its flexibility using the intrinsic characteristic of data itself towards scalable representation learning. The widely embraced scheme that initially introduced from computer vision domain (He




Figure 1: (a) Comparison of rank of interaction matrix across different layers of the encoder in vanilla MTM and our method; (b) Energy distribution comparison of raw data, reconstructed results of vanilla MTM, and that of our method.

et al. 2022), Masked Time-series Modeling (MTM), randomly masks a portion of input timestamps, and then reconstructs the invisible timestamps based on the visible ones. Recently, transformer and its variants have become the predominant choice in MTM, and achieve new state-of-the-art performance on various classification benchmarks (Dong et al. 2023; Cheng et al. 2023). This greatly attributes to its powerful modeling ability of long-range dependency between different timestamps, which facilitates learning the context-aware feature.

Despite their promising performance towards representation learning, these transformer-based methods overlook two potential problems: **(1) feature homogenization.** Several studies (Dong, Cordonnier, and Loukas 2021; Han et al. 2023a, 2024) point out that the feature encoded by transformer-based backbone easily incurs rank collapse due to the long-dependency *ensemble averaging* (Park and Kim 2022), and we experimentally conclude that this phenomenon also exists when encoding the time-series data. As illustrated in Figure 1 (a), we showcase the learned interaction matrix and calculate their ranks from the encoder of vanilla MTM (Vaswani et al. 2017) and our improved method. We observe that vanilla MTM tends to produce low-rank features, which manifests to encode homogenized information and lacks sufficient semantic richness. **(2) energy imbalance.** Based on *frequency principle* (Xu 2020; Rahaman et al. 2019) and architecture-induced frequency preference (Park and Kim 2022), we observe that vanilla

*These authors contributed equally.

†Corresponding author

transformer-based feature learning is inclined to capture low-frequency energy of time-series. In other words, they will easily memorize the sketchy trend of time-series but need more steps to comprehend variation details conveyed by mid/high-frequency energy. As depicted in Figure 1 (b), we transform the original time-series and reconstructed ones into the Fourier domain respectively, and show their difference of spectrum energy distribution. It can be seen that the reconstruction result from vanilla MTM primarily concentrates on low-frequency energy and lacks adequate attention to other frequency band information, thereby leading to the obvious inconsistency with original time-series.

However, how to tackle above issues within MTM remains an under-explored problem. In this paper, we theoretically analyze that (1) the rank of encoding feature intrinsically hinges on property of its corresponding spectrum space, and (2) the distribution of feature space could also be directly revealed by the energy distribution in the spectrum space. Based on above considerations, we reveal the possibility of resorting to the unified spectrum encoding to address them. Specifically, we propose an auxiliary content-aware balanced spectrum decoding branch to achieve the spectrum reconstruction in MTM, where several cascaded blocks with two tailored units are employed to motivate the encoder to retard homogenized and imbalanced feature learning, respectively. **Content-aware Interaction Modulation (CIM) Unit** is first employed, which maps the intermediate representations from encoder into spectrum space and modulates them with a content-aware complex network (Trabelsi et al. 2017). According to *frequency-domain convolutional theory* (Huang et al. 2023), this process essentially equals to dynamic kernel convolutional operation, which not only restricts unnecessary interaction scope compared with long-dependency tactic but also has the flexibility to adjust the receptive field based on variation of content. Following CIM, **Spectrum Energy Rebalance (SER) Unit**, inspired by Bernstein approximation, is leveraged to adjust the energy distribution of encoded feature across different frequency components, thus compensating for overlooked details in mid/high frequency bands. Moreover, dual-constraint loss is further devised to reduce the gap between temporal space and spectrum space, thereby avoiding the information loss in two decoding branches.

We conduct experiments on ten classification datasets from diverse domains with both univariate and multivariate settings, to demonstrate the superiority of our method in representation learning. More importantly, the related visualization and mechanism exploration results sufficiently indicate the behavior of our model.

In summary, our main contributions are as follows:

- We propose a two-pronged masked time-series reconstruction framework, which seamlessly integrates the content-aware balanced decoder into vanilla temporal decoder, along with devised dual-constraint loss, to establish the connection between temporal domain and spectrum domain.
- Towards content-aware balanced decoder, we endow it with two iterative units, to effectively retard feature ho-

mogenization and achieve spectrum energy rebalance.

- We conduct extensive experiments on numerous real-world datasets to validate the effectiveness of our proposed CBD, the superiority of our overall model compared with state-of-art methods, and convincingly reveal the intriguing behavior of our model.

Related Work

Time-Series Pretraining Recent self-supervised Time-Series pre-training techniques can be roughly categorized into two groups: *Time-Series Contrastive Learning* and *Masked Times-Series Modeling*.

Time-Series Contrastive Learning constructs multiple views to increase variance and align features for robust global representation learning. Existing works (Zhang et al. 2022; Meng et al. 2023) focus on designing different views for contrast to capture different preference on learned representation. For example, TS2Vec (Yue et al. 2022) utilizes temporal contrast to improve discrimination of dynamic semantic variations. TS-TCC (Eldele et al. 2021) formulates a cross-view prediction task that incorporates both temporal and contextual contrast. TimesURL (Liu and Chen 2024) enhances learning process through frequency-temporal augmentation and double universum construction. More flexibly, TS-GAC (Wang et al. 2024) emphasizes the role of spatial consistency in a graph contrastive framework. CSL (Liang et al. 2023) leverages shapelet-based learning to better fit series data.

Concurrent with these strands of research, **Masked Time-Series Modeling** learns representations by reconstructing masked timestamps from partial observations. TST (Zerveas et al. 2021) marks the first attempt to employ a transformer-based encoder for this task. PatchTST(Nie et al. 2023) predicts masked subseries-level patches to capture local semantics. SimMTM (Dong et al. 2023) argues that random masking would disrupt temporal variations, thereby proposing to aggregate the point-wise representations derived from multiple masked variations for robust reconstruction. To mitigate the gap between masked representation in pre-trained stage and unmasked ones in fine-tuning stage, TimeMAE (Cheng et al. 2023) leverages a decoupled scheme to respectively encode the semantic-enhanced representation with high information density. In this paper, we primarily focus on optimizing the masked time-series modeling. Different from above studies, our method reveals two crucial issues in existing transformer-based masked modeling framework: *feature homogenization* and *energy imbalance* from the spectrum perspective.

Learning from Spectrum Perspective More and more investigations delve into learning the robust representation by transforming the temporal domain to frequency domain due to its exclusive properties, which have been actively applied to many fields, such as graph classification, domain generalization, and time-series. For example, FACT (Xu et al. 2021) reckons that model that highlights phase information could better deal with cross-domain setting, thus developing a Fourier-based data augmentation strategy. SFA (Zhang et al. 2023b) introduces a spectrum

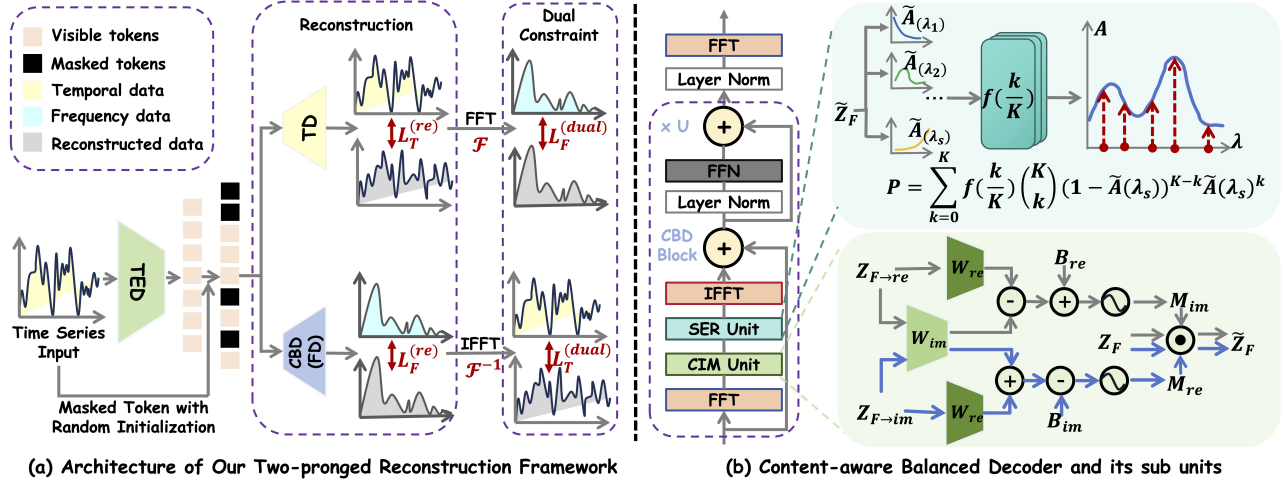


Figure 2: (a) Schematic illustration of our two-pronged framework. TED denotes temporal encoder, and TD and CBD represents temporal decoder and frequency decoder (content-aware balanced decoder). (b) The details of content-aware balanced decoder.

feature alignment method to address the issue of imbalanced feature when performing graph contrast. In the field of time-series analysis, FEDformer (Zhou et al. 2022) launches a frequency-enhanced transformer to emphasize the global view of learned feature for time series forecasting. However, how to effectively retard homogenized and imbalanced feature learning from the spectrum perspective in masked time-series framework still remains an under-explored challenge.

Methodology

In this section, we begin by reviewing the masked modeling framework towards time-series data. Following this, we delve into an in-depth analysis to the key components of our framework, Content-aware Balanced Decoder (CBD), and elaborate on the implementation of its two tailored units. Finally, we provide the design details of loss function that optimizes the overall network.

Review of Masked Time-Series Modeling

Masked Time-Series Modeling (MTM) adopts a transformer-based architecture to reconstruct large portions of masked timestamps via the visible timestamps. Formally, given $\{x_i\}_{i=1}^N$ as a batch of N time-series samples, where $x_i \in \mathbb{R}^{L \times C}$ contains L timestamps and C variables, we randomly mask a certain proportion of timestamps along the temporal dimension. Following the orthodox pipeline of masked modeling framework, MTM encoder takes these visible timestamps as input to generate an intermediate representation $\mathbf{Z} = \{z_i\}_{i=1}^N$, while the following decoder achieves temporal reconstruction of masked timestamps. By aligning the reconstruction results with unmasked timestamps, the encoder could learn the representation that capture context-aware dependencies within time-series data.

Content-aware Balanced Decoder

As suggested by MAE (He et al. 2022), the decoder design is crucial to the MTM model, as it not only models the rela-

tionship between representations of masked tokens and visible tokens, but also determines the semantic preference of learned representation. To address the issues of feature homogenization and spectrum energy imbalance in traditional transformer-based MTM methods, we armed the temporal encoder (TD) with an auxiliary content-aware balanced decoder (CBD), where two tailored units are iteratively deployed. The overall architecture and unfold details are illustrated in Figure 2. In what follows, we elaborate these two units sequentially.

Content-aware Interaction Modulation Unit (CIM) Inspired by the fact that the rank of feature matrix can be determined by its spectrum distribution (i.e., lower bound), $rank(\mathbf{Z}) \geq rank(\mathbf{\Sigma})$, where $\mathbf{\Sigma}$ are the eigenvalue of \mathbf{Z} , we devise a content-aware modulation strategy that operates directly in the spectral space. Specifically, we first transform the intermediate feature \mathbf{Z} to discrete Fourier domain,

$$\mathbf{Z}_F = \mathcal{F}(\mathbf{Z}), \quad (1)$$

where \mathcal{F} denotes the Fourier transformation, and each of components \mathbf{Z}_F can be written as,

$$\mathbf{Z}_F(\lambda_s) = \sum_{t=1}^T \mathbf{Z}(t) \cdot e^{-j \frac{2\pi \lambda_s t}{T}}, \quad (2)$$

and the inverse Fourier transformation \mathcal{F}^{-1} corresponds to

$$\mathbf{Z}(t) = \frac{1}{T} \sum_{\lambda_s=1}^T \mathbf{Z}_F(\lambda_s) \cdot e^{-j \frac{2\pi \lambda_s t}{T}}, \quad (3)$$

where T denotes the total length of intermediate feature of time-series, and λ_s denotes the s -th frequency component in frequency domain. The real and imaginary part of \mathbf{Z}_F can be denoted as $\mathbf{Z}_{F \rightarrow re}$ and $\mathbf{Z}_{F \rightarrow im}$, which exhibits equal contribution for frequency modeling (Trabelsi et al. 2017). In order to achieve more exhaustive spectrum modulation, we encode the real and imaginary parts separately to generate the modulation signals, which is summarized as,

$$\mathbf{M}(\lambda_s) = \sigma(\mathbf{W}_F \mathbf{Z}_F(\lambda_s) + \mathbf{b}_F), \quad (4)$$

where σ is the activation function, $\mathbf{W}_F = \mathbf{W}_{F \rightarrow re} + j \cdot \mathbf{W}_{F \rightarrow im}$ and $\mathbf{b}_F = \mathbf{b}_{F \rightarrow re} + j \cdot \mathbf{b}_{F \rightarrow im}$. More specifically, the generated \mathbf{M} can be further unfolded into real and imaginary parts as follows,

$$\begin{cases} \mathbf{M}_{re}(\lambda_s) = \sigma(\mathbf{W}_{F \rightarrow re} \mathbf{Z}_{F \rightarrow re}(\lambda_s) \\ \quad - \mathbf{W}_{F \rightarrow im} \mathbf{Z}_{F \rightarrow im}(\lambda_s) + \mathbf{B}_{F \rightarrow re}), \\ \mathbf{M}_{im}(\lambda_s) = \sigma(\mathbf{W}_{F \rightarrow im} \mathbf{Z}_{F \rightarrow im}(\lambda_s) \\ \quad + \mathbf{W}_{F \rightarrow re} \mathbf{Z}_{F \rightarrow re}(\lambda_s) + \mathbf{B}_{F \rightarrow im}). \end{cases} \quad (5)$$

Afterwards, the generated complex signal is used to modulate counterparts of original frequency-domain feature, which can be written as,

$$\tilde{\mathbf{Z}}_F(\lambda_s) = \mathbf{M}(\lambda_s) \odot \mathcal{F}(\mathbf{Z}). \quad (6)$$

Based on *theorem 1*, Eq.(6) essentially equals to a dynamic convolutional operation. In a deeper analysis from the perspective of temporal domain, it serves as two purposes: (1) Unlike the long-range dependency mechanism in self-attention block, convolutional operation leads to less freedom of interaction, which naturally restore the rank of features (Han et al. 2023a). (2) its dynamicity indicates that the learned kernel is determined by the content variation of each timestamp, which effectively diversifying the interaction pattern. In other words, for smoothing sub-series, a large interaction scope is possibly unnecessary, while for sub-series that has more obvious fluctuation, enlarging the interaction scope to obtain more context is crucial for understanding its comprehensive semantic. By means of Eq.(6), feature homogenization and rank collapse can be effectively mitigated via direct spectrum modulation. The proof of *theorem 1* is illustrated in supplementary materials.

Theorem 1. (*Frequency-domain convolution theorem*) *The multiplication of two signals in the Fourier domain equals to Fourier transformation of a convolution of these two signals in temporal domain, which can be summarized as,*

$$\mathcal{F}[\mathbf{K}(t) \otimes \mathbf{Z}(t)] = \mathcal{F}(\mathbf{K}(t)) \odot \mathcal{F}(\mathbf{Z}(t)), \quad (7)$$

where \otimes and \odot denote the convolutional operation and element-multiplication operation, respectively, $\mathbf{K}(t)$ and $\mathbf{Z}(t)$ represent two signals with respect with time variable t , and $\mathcal{F}(\cdot)$ denotes the Fourier transformation.

Spectrum Energy Rebalance Unit (SER) For time-series data, timestamps with obvious fluctuations generally contain more high-frequency components (Qin et al. 2021) in spectrum space. However, several previous literatures indicate that networks tend to prioritize learning low-frequency components over high-frequency ones (Xu and Zhou 2021; Xu, Zhang, and Xiao 2019), resulting in challenges in capturing high-frequency details. Furthermore, according to *theorem 2*, overlooking the certain bands of frequency components would greatly impact the temporal information encoding, thereby hindering balanced representation learning. The proof of *theorem 2* is shown in supplementary materials.

Theorem 2. (*Parseval's theorem*) *Assume that \mathbf{Z} denotes the original temporal domain feature, and \mathbf{Z}_F is the corresponding spectrum representation, then the energy of feature in temporal domain equals to the energy in frequency*

domain,

$$\sum_{t=-\infty}^{\infty} |\mathbf{Z}(t)|^2 = \sum_{s=-\infty}^{\infty} |\mathbf{Z}_F(\lambda_s)|^2, \quad (8)$$

where $\mathbf{Z}_F = \mathcal{F}(\mathbf{Z})$, $\mathcal{F}(\cdot)$ denotes the Fourier transformation, λ_s represents the s -th spectrum component, and t denotes the time variable

To tackle it, we design Spectrum Energy Rebalance unit (SER) to flexibly adjust the energy distribution in spectrum space to approximate that of original time-series data. We assume that the ideal response function from spectrum components to energy can be denoted as $\mathcal{G}(\cdot)$, and the initial response is $\mathbf{A} = \sqrt{\tilde{\mathbf{Z}}_{F \rightarrow re}^2 + \tilde{\mathbf{Z}}_{F \rightarrow im}^2} = \|\tilde{\mathbf{Z}}_F\|_2$. SER serves as the energy rebalance mapping function $\mathcal{P}(\cdot)$, which makes the initial response \mathbf{A} head to real response $\mathcal{G}(\tilde{\mathbf{Z}}_F)$, i.e., $\mathcal{G}(\tilde{\mathbf{Z}}_F) = \mathcal{P}(\cdot) \circ \|\cdot\|_2(\tilde{\mathbf{Z}}_F)$, where \circ denotes the compound operation. To better simulate the effect of $\mathcal{P}(\cdot)$, we adopt Bernstein polynomials. Before unfolding our design details of SER, we first give the definition of *Bernstein polynomial approximation* in the following.

Definition 1. (*Bernstein polynomial approximation*) *Given an arbitrary continuous mapping function $f(w)$ from $w \in [0, 1]$, the K -order Bernstein polynomial approximation of $f(w)$ is defined as,*

$$p_K(w) = \sum_{k=0}^K \Theta_k \cdot \mathcal{B}_k^K(w) = \sum_{k=0}^K f\left(\frac{k}{K}\right) \cdot \binom{K}{k} (1-w)^{K-k} w^k, \quad (9)$$

and we have $p_K(w) \rightarrow f(w)$ as $K \rightarrow \infty$, where Θ_k works as the coefficient of $\mathcal{B}_k^K(w)$, and $\mathcal{B}_k^K(w)$ serves as the base of polynomial.

Given that the domain of definition of $p_K(w)$ is restricted to $[0, 1]$, we therefore adopt the softmax function to normalize the input energy distribution $\mathbf{A}(\lambda_s)$, and we substitute the normalized energy distribution $\tilde{\mathbf{A}}(\lambda_s)$ into the approximated polynomial $p_K(w)$,

$$p_K = \sum_{k=0}^K f\left(\frac{k}{K}\right) \cdot \binom{K}{k} (1 - \tilde{\mathbf{A}}(\lambda_s))^{K-k} \tilde{\mathbf{A}}(\lambda_s)^k. \quad (10)$$

Beyond that, in order to adjust \mathcal{P} (i.e., p_K) in a data-driven manner, we leverage an energy-aware gating network to control the filter coefficients,

$$f\left(\frac{k}{K}\right) = (\mathbf{W}_c \mathbf{A}^\dagger + \mathbf{b}_c)_k, \quad (11)$$

where $\mathbf{A}^\dagger = [\tilde{\mathbf{A}}(\lambda_1), \tilde{\mathbf{A}}(\lambda_2), \dots, \tilde{\mathbf{A}}(\lambda_S)] \in \mathbb{R}^S$. $\mathbf{W}_c \in \mathbb{R}^{K \times S}$ and $\mathbf{b}_c \in \mathbb{R}^K$ are the learnable weights and bias of gating network, respectively.

Such mechanism potentially undertakes the role that it allows the model to have elasticity to compensate for the missing spectrum components in the training process. Interestingly, from a graph perspective, different ways of frequency response can be interpreted as different specific filters (Xu, Zhang, and Xiao 2019; Bianchi et al. 2022). When K is sufficiently large, function \mathcal{P} could simulate arbitrary spectrum rebalance functions (i.e., filters). Finally, the learned

spectrum rebalance function is utilized to modulate the frequency component to obtain the balanced spectrum feature,

$$\tilde{\mathbf{Z}}_{F \rightarrow \text{bal}}(\lambda_s) = p_K(\lambda_s) \odot \tilde{\mathbf{Z}}_F(\lambda_s). \quad (12)$$

Iterative Architecture As illustrated in Figure 2(b), it is built upon the vanilla transformer architecture, which modifies the self-attention block by integrating CIM and SER unit. For each layer $u \in [1, U]$, where U is the total number of layer, a CBD block processes and updates the masked representation as follows,

$$\tilde{\mathbf{Z}}_F^{(u)} = \text{CBDBlock}(\tilde{\mathbf{Z}}_F^{(u-1)}; \Omega^{(u)}), \quad (13)$$

where $\Omega^{(u)} = \{\mathbf{W}_F^{(u)}, \mathbf{b}_F^{(u)}, \mathbf{W}_c^{(u)}, \mathbf{b}_c^{(u)}\}$ denotes the set of learned parameters in each block. At each step, CBD block retards rank collapse of interaction matrix and mitigates feature smoothing via CIM unit, and further leverages SER unit to balance spectrum energy. Moreover, each block has independent parameters, allowing blocks at different depths to collaboratively refine the masked representation.

Optimization Strategy

Dual-Constraint Pre-trained Loss During the pre-training stage, as shown in Figure 2(b), temporal encoder (TED) encodes the masked time series and decodes it into two outputs: temporal reconstruction $\tilde{\mathbf{Z}}_T^{(U)}$ and frequency reconstruction $\tilde{\mathbf{Z}}_F^{(U)}$. We denote the ground truth of temporal domain and its frequency domain as $\mathbf{Z}_T^{(gt)}$ and $\mathbf{Z}_F^{(gt)}$, the pre-training loss is formulated as:

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_T^{(re)}(\tilde{\mathbf{Z}}_T, \mathbf{Z}_T^{(gt)}) + \mathcal{L}_F^{(dual)}(\mathcal{F}(\tilde{\mathbf{Z}}_T), \mathbf{Z}_F^{(gt)}) \\ & + \gamma \left[\mathcal{L}_F^{(re)}(\tilde{\mathbf{Z}}_F^{(U)}, \mathbf{Z}_F^{(gt)}) + \mathcal{L}_T^{(dual)}(\mathcal{F}^{-1}(\tilde{\mathbf{Z}}_F^{(U)}), \mathbf{Z}_T^{(gt)}) \right], \end{aligned} \quad (14)$$

where $\mathcal{L}_T^{(re)}$ represents the reconstruction loss in traditional masked modeling, while $\mathcal{L}_F^{(dual)}$ serves as the dual-constraint loss aforementioned. For temporal domain loss $\mathcal{L}_T^{(re)}$, we compute Mean Square Error (MSE) between the reconstructed and raw time-series. As to frequency domain loss $\mathcal{L}_F^{(re)}$, we compute squared Euclidean distance between their real and imaginary parts. γ is the loss weight to control the contribution of each branch. Specially, the dual-constraint loss ensures that the two decoding branches can enhance the information consistency between temporal domain and frequency domain in the training process.

Fine-tuned Loss In the fine-tuning stage, the output \mathbf{Z} from the common encoder is fed into the final classification head f_c . Then the fine-tuned loss is termed as,

$$\mathcal{L} = \mathcal{L}_{CE}(f_c(\mathbf{Z}), y), \quad (15)$$

where the Cross-Entropy Loss is utilized to minimize the distance between label y and the prediction $f_c(\mathbf{Z})$.

Experiments

In this section, we first introduce the experiment setup. Then, we present and analyze the quantitative results of our model and a bunch of baselines on time series classification. Finally, we demonstrate the effectiveness of our model through analytical experiments.

Experiment Setup

Datasets We conduct experiments on ten publicly available datasets to validate the performance of our model, including Human Activity Recognition (HAR)(Anguita et al. 2012) and nine large datasets from the UEA(Bagnall et al. 2018) and UCR archive(Dau et al. 2019), which is referred as PS, SRSCP1, MI, FM, AWR, SAD, ECG5000, FB, Uware. These datasets not only encompass both univariate and multivariate datasets, but also cover diverse domains. More details are provided in supplementary materials.

Baselines We compare our model with recent advanced self-supervised learning methods. To highlight our performance, we select five contrastive learning methods and four masked modeling methods. More details about these baselines are provided in supplementary part:

- **Contrastive Learning Methods:** TS-TCC(Eldele et al. 2021), TS2Vec(Yue et al. 2022), TS-GAC(Wang et al. 2024), TimesURL(Liu and Chen 2024), CSL(Liang et al. 2023).
- **Reconstructed Methods:** PatchTST(Nie et al. 2023), CRT(Zhang et al. 2023a), SimMTM(Dong et al. 2023), TimeMAE(Cheng et al. 2023).

Evaluation In the classification task, we evaluate the performance in two mainstream manner: **linear evaluation** and **fine-tuning evaluation**. The former manner freezes the pre-trained parameters of encoder and trains the classifier layer with labeled data. The latter manner updates both the encoder and classifier layer based on the pre-trained parameters. We utilize accuracy as the metrics for classification tasks and mark the **best** and **second best** values. All experiments are conducted with five different seeds and the average results are taken for comparisons.

Implement Details For the encoder, we use the same backbone as TimeMAE with the default 8 transformer layers, while the decoder is configured with 2 layers for both branches. We also adopt the same masking strategy as TimeMAE. We set the batch size as 128 and choose AdamW optimizer with a learning rate of 1e-4. All methods are conducted with NVIDIA A10 and implemented by PyTorch. More implementation and baseline experiments details are provided in supplementary materials.

Quantitative Experimental Results

Table 1 presents the classification results. Firstly, we observe that our method achieves the highest accuracy across most of datasets, regardless of linear evaluation or fine-tuning task. Particularly, HAR and MI show significant improvements, surpassing the second-best by 2.38% and 1.90% in linear probing accuracy. The superior performance of linear evaluation suggests that our method effectively narrows the gap between pretrained representation and task-specific representation, demonstrating strong feature generalization capabilities. Meanwhile, the performance gain on fine-tuning highlight model’s adaptability to specific tasks. Notably, we exclude TimesURL from fine-tune evaluation due to its implementation constraints to ensure the rigor of our experiments. Secondly, our method supports both univariate and

Methods	Linear Evaluation									
	HAR	PS	SRSCP1	MI	FM	AWR	SAD	ECG5000	FB	UWare
TS-TCC(IJCAI'21)	89.22±0.70	14.27±0.39	83.64±0.99	55.47±0.68	48.00±1.63	93.23±0.68	95.20±0.15	92.67±0.78	50.29±0.31	83.17±0.29
TS2Vec(AAAI'22)	90.36±0.32	10.82±0.38	83.61±0.99	51.00±0.75	47.10±4.22	98.30±0.09	97.31±0.19	93.50±0.65	78.90±0.54	93.40±0.53
TS-GAC(AAAI'24)	<u>91.40±0.16</u>	13.53±0.48	86.04±0.31	56.00±0.46	53.39±0.54	<u>98.50±0.06</u>	97.99±0.05	-	-	-
TimesURL(AAAI'24)	87.18±0.46	16.88±0.42	85.32±0.95	61.20±0.35	56.00±1.84	98.00±0.13	95.86±0.10	93.73±0.58	73.95±0.34	94.22±0.38
CSL(VLDB'24)	83.40±0.33	<u>17.63±0.36</u>	84.60±0.74	61.00±0.45	<u>56.25±2.63</u>	97.67±1.35	94.95±0.26	93.16±0.84	<u>79.50±0.42</u>	<u>95.31±0.20</u>
PatchTST(ICML'22)	77.89±1.98	14.42±0.26	68.36±2.15	61.00±1.36	55.00±2.10	90.63±0.86	91.54±0.41	90.25±0.45	59.11±0.34	90.01±0.30
CRT(TNNLS'23)	89.21±0.56	7.60±1.17	58.03±1.11	51.67±2.86	54.00±1.71	82.50±1.17	85.81±0.29	91.00±0.34	75.73±0.08	80.87±0.24
SimMTM(NeurIPS'23)	87.90±0.35	15.45±0.12	90.78±0.30	62.00±0.36	55.33±3.21	95.00±0.16	95.05±0.13	92.57±0.18	51.98±0.22	91.58±0.30
TimeMAE(Arxiv'23)	91.31±0.84	14.13±0.34	85.53±1.84	<u>62.60±1.67</u>	53.80±1.64	95.07±0.72	95.76±0.51	<u>93.92±0.17</u>	55.31±0.33	86.78±1.64
Ours	93.78±0.70	18.22±0.12	<u>86.72±0.98</u>	64.50±0.20	56.40±2.94	98.52±0.29	<u>97.98±0.21</u>	94.51±0.08	79.57±0.38	95.99±0.12

	Fine-tune Evaluation									
	HAR	PS	SRSCP1	MI	FM	AWR	SAD	ECG5000	FB	UWare
TS-TCC(IJCAI'21)	91.66±0.42	20.72±0.68	83.76±1.07	57.81±0.30	46.00±2.12	96.61±0.33	98.71±0.09	93.42±0.61	63.68±0.52	89.35±0.05
TS2Vec(AAAI'22)	95.10±0.36	16.07±0.66	84.43±0.61	53.00±0.49	49.21±2.33	98.01±0.06	98.36±0.34	93.63±0.52	79.82±0.14	93.45±0.06
TS-GAC(AAAI'24)	94.83±0.10	16.76±0.35	86.44±2.50	56.96±2.54	49.11±4.15	98.06±0.07	99.01±0.20	-	-	-
CSL(VLDB'24)	89.31±0.40	18.13±0.48	83.95±0.68	<u>64.50±0.64</u>	56.00±3.27	<u>98.66±0.16</u>	99.44±0.23	<u>94.42±0.37</u>	79.60±0.30	<u>95.53±0.27</u>
PatchTST(ICML'22)	88.45±0.70	18.42±0.64	81.25±1.78	62.40±2.35	57.81±2.25	98.05±0.06	97.75±0.16	94.40±0.21	75.65±0.23	91.36±0.17
CRT(TNNLS'23)	90.09±0.75	8.38±0.19	67.65±1.36	50.67±2.65	53.00±2.56	87.62±0.84	98.23±0.13	92.53±0.32	79.24±0.13	85.64±0.09
SimMTM(NeurIPS'23)	93.50±0.32	<u>21.95±0.20</u>	92.72±0.93	62.00±0.20	60.00±2.15	98.57±0.10	99.30±0.10	92.76±0.13	82.53±0.18	94.64±0.13
TimeMAE(Arxiv'23)	95.11±0.18	19.49±0.59	87.71±1.27	61.80±2.28	61.00±1.87	97.73±0.43	99.20±0.03	94.25±0.14	67.21±0.44	92.81±0.47
Ours	95.92±0.78	23.27±0.37	86.25±1.14	65.20±0.31	60.60±3.38	98.67±0.19	<u>99.44±0.12</u>	94.84±0.10	80.83±0.45	95.85±0.23

Table 1: Comparisons with State-of-the-Art methods with different evaluations (%)




Figure 3: Learned Interaction Matrix on the HAR dataset.

multivariate setting. It is worth noting that TS-GAC considers variables as nodes in graph network, which struggles to deal with single-variable cases. This demonstrates both flexibility and generalization of our method. Finally, compared with contrastive-based baselines, our improved masked modeling scheme could strike a balance between dynamic local variation and holistic semantic representation. Compared to masked modeling baselines, our content-aware balanced decoder brings diverse and balanced decoding potential from the spectrum perspective.

Analytical Experiment

Analysis of Content-aware Interaction Modulation We pretrain our network and vanilla transformer-based reconstruction model on the HAR dataset for 20, 50 and 80 epochs. Then we visualize the interaction matrix obtained from the last encoder layer in Figure 3. Both methods exhibit similar encoding patterns in the early stages, empha-

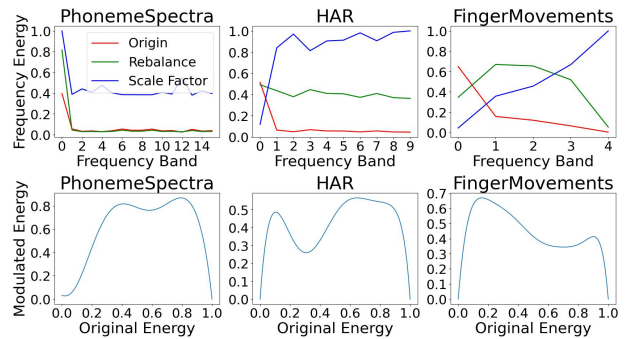


Figure 4: Energy rebalance of SER. The top row depicts the modulation during the energy rebalancing, while the bottom represents the corresponding learned Bernstein polynomials.

sizing long-dependency interactions. However, our method employs a more prudent and distinct tactics for interaction scope control in the initial phase. This ability stems from our Content-aware Interaction Modulation Unit (CIM), enabling more accurate and flexible interaction scope compared to the pure self-attention decoding. As training progresses, the encoding receptive fields of both methods gradually shrink, but our approach maintains a more semantically compact receptive field. More visualizations on other datasets are provided in supplementary materials.

Analysis of Spectrum Energy Rebalance We investigate an in-depth investigation into the modulation mechanism of the Spectrum Energy Rebalance Unit (SER). Figure 4 visualizes the original and modulated spectrum energy distributions for PhonemeSpectra, HAR and FingerMovements. SER adaptively adjusts the energy distribution by learning

Method	LP. Acc.(%)	FT. Acc.(%)
$\mathcal{L}_T^{(re)}$	89.47	93.34
$\mathcal{L}_F^{(re)}$	91.61	94.02
$\mathcal{L}_T^{(re)} + \mathcal{L}_F^{(dual)}$	91.68	94.16
$\mathcal{L}_F^{(re)} + \mathcal{L}_T^{(dual)}$	91.78	94.12
$\mathcal{L}_T^{(re)} + \mathcal{L}_F^{(re)}$	93.37	95.61
w/o SER	91.92	94.06
w/o CIM	93.45	95.78
Ours	93.78	95.92

Table 2: Ablation study of designed losses and components in CBD unit.

K Value	1	2	4	8	12	16
LP Acc. (%)	92.39	92.76	93.61	93.04	93.78	93.48
LP F1 (%)	91.86	92.36	93.13	92.61	93.32	93.05

Table 3: Ablation study of K value in SER.

data-specific Bernstein coefficients. For HAR (top center), the red line shows the original network concentrating energy in the low-frequency range, indicating a bias towards low-frequency information. As frequency increases, energy decreases, leading to insufficient fitting of mid-to-high frequencies. With SER, the green line shows a marked increase in mid-to-high-frequency energy. The blue line also illustrates SER’s modulation role, which serves as the normalized scaling factor from original energy distribution to balanced one. This flexibility allows SER to compensate for overlooked information and prevent representation degradation. Notably, the model does not excessively prioritize mid-to-high frequencies, as seen in PhonemeSpectra. To further clarify SER’s effect, we visualize the learned Bernstein polynomials in the bottom row, summarizing the mapping from normalized energy to the final energy response and highlighting the rebalancing preferences for each dataset.

Ablation Study To take a closer look on each part in our framework, we devise several variants and measure their performance using both linear probing (LP) and fine-tuning (FT) on HAR. Results are reported in Table 2 and Table 3. **(1) Effect of Designed Losses.** Key observations from designed losses include: (1) Both the vanilla temporal decoder and our spectrum-aware decoder benefit from the dual-constraint loss. Notably, the frequency-constraint loss $\mathcal{L}_F^{(dual)}$, without the frequency decoding branch, improves linear evaluation accuracy by 2.21%. (2) Our frequency decoder outperforms the traditional temporal decoder in feature encoding, achieving improvements of 2.14% in LP and 0.68% in FT by addressing feature homogenization and spectrum imbalance. (3) The two-pronged decoders could collaboratively improve encoding capability more than either alone, yielding improvements of 3% and 1.76% in LP accuracy compared to the standalone temporal and frequency decoders, respectively. Combining them with dual-

Model	Avg. LP Acc.	Avg. FT Acc.
PatchTST + CBD	69.82 71.32 (+1.50)	76.55 77.56 (+1.01)
CRT + CBD	67.64 72.95 (+5.31)	71.31 75.58 (+4.27)
SimMTM + CBD	73.81 74.74 (+0.93)	79.80 80.39 (+0.59)

Table 4: Performance by integrating CBD (frequency decoder) to three advanced mask modeling models.

constraint loss $\mathcal{L}_F^{(dual)}$ and $\mathcal{L}_T^{(dual)}$ results in the best representation learning performance. **(2) Effectiveness of Each Component in CBD.** Table 2 also examines the roles of the CIM and SER units in the frequency decoder. The SER can rebalance the spectrum energy without the CIM. This mechanism can encourage the model to emphasize high-frequency overlooked details. Meanwhile, the CIM dynamically controls the interaction scope based on the local signal variation via flexible large kernel convolutions. With the CIM, the model reduces unnecessary information aggregation and facilitate feature encoding. **(3) Performance Variation with Different Order K .** To explore the theoretical boundary of Bernstein approximation, we examined performance with varying order K of the Bernstein polynomial in Table 3. Theoretically, as K approaches infinity, the polynomial can fit any function. However, we found that performance improves with increasing K up to 12, after which it declines. This suggests that $K = 12$ is sufficient for SER to learn the scaling factor towards spectrum rebalancing, while higher value will introduce unnecessary noise. Thus, we select $K = 12$ as the optimal value.

CBD Generality From Table 4, we can find that the CBD, as an additional decoder, improves the classification performance of diverse advanced masked modeling models. Specifically, for CRT, our module increases the LP accuracy by an average of 5.31% and FT accuracy by 4.27% across ten datasets. It is worth mentioning that since TimeMAE does not use the origin data for reconstruction, the CBD may receive inaccurate supervisory signals, leading to performance decline. Detailed experimental results are available in supplementary materials.

Conclusion

In this paper, we propose a two-pronged reconstruction framework to improve the quality of time-series representation. Specifically, we harness the content-aware balanced decoder with two built-in units to form our powerful frequency decoder, which works collaboratively with the vanilla temporal decoder to effectively address the issues of feature homogenization and spectral energy imbalance. The extensive experiments convincingly demonstrate the superiority of our method based on the observation on quantitative results, behavior analysis, and representation visualization.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China, No.:62276155, No.:62376137, No.:U24A20328, No.:62476071, No.:62206156, and No.:62206157; in part by the National Natural Science Foundation of Shandong Province, No.:ZR2021MF040, No.:ZR2024QF104 and No.:ZR2022QF047.

References

- Anguita, D.; Ghio, A.; Oneto, L.; Parra, X.; and Reyes-Ortiz, J. L. 2012. Human activity recognition on smartphones using a multiclass hardware-friendly support vector machine. In *Ambient Assisted Living and Home Care: 4th International Workshop, IWAAL 2012, Vitoria-Gasteiz, Spain, December 3-5, 2012. Proceedings 4*, 216–223. Springer.
- Bagnall, A.; Dau, H. A.; Lines, J.; Flynn, M.; Large, J.; Bostrom, A.; Southam, P.; and Keogh, E. 2018. The UEA multivariate time series classification archive, 2018. *arXiv preprint arXiv:1811.00075*.
- Bianchi, F. M.; Grattarola, D.; Livi, L.; and Alippi, C. 2022. Graph Neural Networks With Convolutional ARMA Filters. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(7): 3496–3507.
- Cheng, M.; Liu, Q.; Liu, Z.; Zhang, H.; Zhang, R.; and Chen, E. 2023. TimeMAE: Self-Supervised Representations of Time Series with Decoupled Masked Autoencoders. *arXiv preprint arXiv:2303.00320*.
- Dau, H. A.; Bagnall, A.; Kamgar, K.; Yeh, C.-C. M.; Zhu, Y.; Gharghabi, S.; Ratanamahatana, C. A.; and Keogh, E. 2019. The UCR time series archive. *IEEE/CAA Journal of Automatica Sinica*, 6(6): 1293–1305.
- Dong, J.; Wu, H.; Zhang, H.; Zhang, L.; Wang, J.; and Long, M. 2023. SimMTM: A Simple Pre-Training Framework for Masked Time-Series Modeling. *arXiv preprint arXiv:2302.00861*.
- Dong, Y.; Cordonnier, J.; and Loukas, A. 2021. Attention is not all you need: pure attention loses rank doubly exponentially with depth. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, 2793–2803. PMLR.
- Eldele, E.; Ragab, M.; Chen, Z.; Wu, M.; Kwoh, C. K.; Li, X.; and Guan, C. 2021. Time-series representation learning via temporal and contextual contrasting. *arXiv preprint arXiv:2106.14112*.
- Han, D.; Pan, X.; Han, Y.; Song, S.; and Huang, G. 2023a. FLatten Transformer: Vision Transformer using Focused Linear Attention. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, 5938–5948. IEEE.
- Han, Y.; Guo, Y.; Yin, J.; Liu, M.; Hu, Y.; and Nie, L. 2021. Focal and Composed Vision-semantic Modeling for Visual Question Answering. In Shen, H. T.; Zhuang, Y.; Smith, J. R.; Yang, Y.; César, P.; Metz, F.; and Prabhakaran, B., eds., *MM '21: ACM Multimedia Conference, Virtual Event, China, October 20 - 24, 2021*, 4528–4536. ACM.
- Han, Y.; Hu, Y.; Song, X.; Tang, H.; Xu, M.; and Nie, L. 2024. Exploiting the Social-Like Prior in Transformer for Visual Reasoning. In *AAAI, 2058–2066*. AAAI Press.
- Han, Y.; Yin, J.; Wu, J.; Wei, Y.; and Nie, L. 2023b. Semantic-Aware Modular Capsule Routing for Visual Question Answering. *TIP*, 32: 5537–5549.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16000–16009.
- Hu, Y.; Liu, M.; Su, X.; Gao, Z.; and Nie, L. 2021a. Video Moment Localization via Deep Cross-Modal Hashing. *IEEE Trans. Image Process.*, 30: 4667–4677.
- Hu, Y.; Nie, L.; Liu, M.; Wang, K.; Wang, Y.; and Hua, X. 2021b. Coarse-to-Fine Semantic Alignment for Cross-Modal Moment Localization. *IEEE Trans. Image Process.*, 30: 5933–5943.
- Hu, Y.; Wang, K.; Liu, M.; Tang, H.; and Nie, L. 2024. Semantic Collaborative Learning for Cross-Modal Moment Localization. *ACM Trans. Inf. Syst.*, 42(2): 50:1–50:26.
- Huang, Z.; Zhang, Z.; Lan, C.; Zha, Z.; Lu, Y.; and Guo, B. 2023. Adaptive Frequency Filters As Efficient Global Token Mixers. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, 6026–6036. IEEE.
- Liang, Z.; Zhang, J.; Liang, C.; Wang, H.; Liang, Z.; and Pan, L. 2023. Contrastive Shapelet Learning for Unsupervised Multivariate Time Series Representation Learning. *CoRR*, abs/2305.18888.
- Liu, J.; and Chen, S. 2024. TimesURL: Self-Supervised Contrastive Learning for Universal Time Series Representation Learning. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, 13918–13926. AAAI Press.
- Meng, Q.; Qian, H.; Liu, Y.; Cui, L.; Xu, Y.; and Shen, Z. 2023. MHCCL: masked hierarchical cluster-wise contrastive learning for multivariate time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 9153–9161.
- Nie, Y.; Nguyen, N. H.; Sinthong, P.; and Kalagnanam, J. 2022. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*.
- Nie, Y.; Nguyen, N. H.; Sinthong, P.; and Kalagnanam, J. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Ozyurt, Y.; Feuerriegel, S.; and Zhang, C. 2022. Contrastive learning for unsupervised domain adaptation of time series. *arXiv preprint arXiv:2206.06243*.

- Park, N.; and Kim, S. 2022. How Do Vision Transformers Work? In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Qin, Z.; Zhang, P.; Wu, F.; and Li, X. 2021. FcaNet: Frequency Channel Attention Networks. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, 763–772. IEEE.
- Rahaman, N.; Baratin, A.; Arpit, D.; Draxler, F.; Lin, M.; Hamprecht, F. A.; Bengio, Y.; and Courville, A. C. 2019. On the Spectral Bias of Neural Networks. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, 5301–5310. PMLR.
- Tonekaboni, S.; Eytan, D.; and Goldenberg, A. 2021. Unsupervised representation learning for time series with temporal neighborhood coding. *arXiv preprint arXiv:2106.00750*.
- Trabelsi, C.; Bilaniuk, O.; Serdyuk, D.; Subramanian, S.; Santos, J. F.; Mehri, S.; Rostamzadeh, N.; Bengio, Y.; and Pal, C. J. 2017. Deep Complex Networks. *CoRR*, abs/1705.09792.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, Y.; Xu, Y.; Yang, J.; Wu, M.; Li, X.; Xie, L.; and Chen, Z. 2024. Graph-Aware Contrasting for Multivariate Time-Series Classification. *arXiv:2309.05202*.
- Xu, Q.; Zhang, R.; Zhang, Y.; Wang, Y.; and Tian, Q. 2021. A fourier-based framework for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14383–14392.
- Xu, Z. J.; Zhang, Y.; and Xiao, Y. 2019. Training Behavior of Deep Neural Network in Frequency Domain. In Gedeon, T.; Wong, K. W.; and Lee, M., eds., *Neural Information Processing - 26th International Conference, ICONIP 2019, Sydney, NSW, Australia, December 12-15, 2019, Proceedings, Part I*, volume 11953 of *Lecture Notes in Computer Science*, 264–274. Springer.
- Xu, Z. J.; and Zhou, H. 2021. Deep Frequency Principle Towards Understanding Why Deeper Learning Is Faster. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, 10541–10550. AAAI Press.
- Xu, Z.-Q. J. 2020. Frequency Principle: Fourier Analysis Sheds Light on Deep Neural Networks. *Communications in Computational Physics*, 28(5): 1746–1767.
- Yue, Z.; Wang, Y.; Duan, J.; Yang, T.; Huang, C.; Tong, Y.; and Xu, B. 2022. Ts2vec: Towards universal representation of time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 8980–8987.
- Zerveas, G.; Jayaraman, S.; Patel, D.; Bhamidipaty, A.; and Eickhoff, C. 2021. A transformer-based framework for multivariate time series representation learning. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, 2114–2124.
- Zhang, W.; Yang, L.; Geng, S.; and Hong, S. 2023a. Self-Supervised Time Series Representation Learning via Cross Reconstruction Transformer. *IEEE Transactions on Neural Networks and Learning Systems*.
- Zhang, X.; Zhao, Z.; Tsiligkaridis, T.; and Zitnik, M. 2022. Self-supervised contrastive pre-training for time series via time-frequency consistency. *Advances in Neural Information Processing Systems*, 35: 3988–4003.
- Zhang, Y.; Zhu, H.; Song, Z.; Koniusz, P.; and King, I. 2023b. Spectral feature augmentation for graph contrastive learning and beyond. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11289–11297.
- Zhou, T.; Ma, Z.; Wen, Q.; Wang, X.; Sun, L.; and Jin, R. 2022. FEDformer: Frequency Enhanced Decomposed Transformer for Long-term Series Forecasting. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, 27268–27286. PMLR.