

Graph-Aware Contrasting for Multivariate Time-Series Classification

Yucheng Wang^{1,3}, Yuecong Xu¹, Jianfei Yang³,
Min Wu¹, Xiaoli Li^{1,2,3}, Lihua Xie³, Zhenghua Chen^{1,2*}

¹Institute for Infocomm Research, A*STAR, Singapore

²Centre for Frontier AI Research, A*STAR, Singapore

³Nanyang Technological University, Singapore

{yucheng003, xuyu0014, yang0478, chen0832}@e.ntu.edu.sg, {wumin, xlli}@i2r.a-star.edu.sg, elhxie@ntu.edu.sg

Abstract

Contrastive learning, as a self-supervised learning paradigm, becomes popular for Multivariate Time-Series (MTS) classification. It ensures the consistency across different views of unlabeled samples and then learns effective representations for these samples. Existing contrastive learning methods mainly focus on achieving temporal consistency with temporal augmentation and contrasting techniques, aiming to preserve temporal patterns against perturbations for MTS data. However, they overlook spatial consistency that requires the stability of individual sensors and their correlations. As MTS data typically originate from multiple sensors, ensuring spatial consistency becomes essential for the overall performance of contrastive learning on MTS data. Thus, we propose Graph-Aware Contrasting for spatial consistency across MTS data. Specifically, we propose graph augmentations including node and edge augmentations to preserve the stability of sensors and their correlations, followed by graph contrasting with both node- and graph-level contrasting to extract robust sensor- and global-level features. We further introduce multi-window temporal contrasting to ensure temporal consistency in the data for each sensor. Extensive experiments demonstrate that our proposed method achieves state-of-the-art performance on various MTS classification tasks. The code is available at <https://github.com/Frank-Wang-oss/TS-GAC>.

Introduction

Multivariate Time-Series (MTS) data are widely used in areas such as healthcare and industrial manufacturing for classification tasks, attracting significant research interests. To improve the performance of MTS classification, deep learning has gained popularity for learning effective representations (Craik, He, and Contreras-Vidal 2019; Chen et al. 2021; Deng and Hooi 2021; Chen et al. 2020c; Zhao et al. 2019). However, the need for substantial labeled samples poses challenges as large-scale manual labeling is impractical, limiting their applicability to real-world scenarios. To address this challenge, Contrastive Learning (CL) has emerged as a promising approach (Zhang et al. 2023a; Eldele et al. 2023). By contrasting the different views of unlabeled samples that are commonly generated by augmen-

tations, CL enhances encoder’s robustness to perturbations and learns robust and effective representations.

Researchers have recently begun exploring CL for MTS data (Eldele et al. 2021; Tonekaboni, Eytan, and Goldenberg 2021), with a primary focus on achieving temporal consistency by preserving temporal patterns robustly against perturbations. Specifically, temporal augmentations such as jittering or permutation are commonly used to create different views for MTS data. Encoders are then employed to extract temporal features, based on which CL is performed to make the encoders robust to temporal disturbances, thus preserving temporal patterns within MTS data. To further enhance temporal consistency, temporal contrasting is often achieved with a predictive contrastive loss when predicting the future timestamps with the past information (Choi and Kang 2023; Eldele et al. 2021).

While the current methods have made progress with CL for MTS data, they mainly focused on temporal consistency while ignoring spatial consistency during the CL process. Here, the spatial consistency refers to maintaining the stability of both the individual sensors and the correlations across the different sensors. Specifically, the robustness of MTS data relies on the stability of each individual sensor, i.e., any disturbance in a sensor could have a significant impact on the classification performance of an MTS sample. We take Fig. 1 for illustration. Amplitude disturbances, such as insensitivity, in foot signals can lead to the similar foot amplitude in walking and running actions, potentially causing a classifier to misclassify running as walking. Thus, a robust algorithm should be able to identify disturbances within individual sensors. Moreover, correlations exist between sensors, with certain sensors exhibiting stronger correlations across each other than with others. For example, due to the physical connection between the foot and knee, a foot sensor is more correlated with a knee sensor than a hand sensor. Preserving the robustness of these relative sensor relationships can further help learn robust sensor features (Yu, Yin, and Zhu 2017; Jia et al. 2020). As MTS data typically originate from multiple sensors, it is crucial to ensure the spatial consistency to enhance the overall CL performance on MTS data.

The above discussion motivates us to propose a novel approach called Graph-Aware Contrasting for MTS data (TS-GAC). To achieve spatial consistency, specific augmentation and contrasting methods tailored for MTS data are de-

*Corresponding Author

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Signals from knee and foot for walking and running. The foot sensor is more important for classification than the knee sensor due to its large amplitude. (a) During walking, both knee and foot have low frequency and amplitude. (b) During running, both sensors show increased frequency and amplitude. Disturbances in the foot sensor, like insensitivity, may cause running signals to have a similar amplitude to walking signals, which may mislead a classifier and mis-classify running as walking.

signed. We first design graph augmentations, involving node and edge augmentations, to augment MTS data. For node augmentations, we apply temporal and frequency augmentations (Zhang et al. 2022; Yang and Hong 2022) to fully augment each sensor, while edge augmentations are designed to augment sensor correlations, ensuring robustness in the relationships between sensors. By capturing the augmented sensor correlations, Graph Neural Network (GNN) (Wang et al. 2023; Jia et al. 2020) is utilized to update sensor features.

With updated sensor features, we then design graph contrasting which incorporates both node- and graph-level contrasting to learn robust sensor- and global-level features. For node-level contrasting, we create two views using the proposed augmentations and contrast the sensors in different views within each MTS sample to ensure the robustness of each sensor against perturbations. Additionally, we map the sensor features to global features and introduce graph-level contrasting by contrasting MTS samples in different views within each training batch. Furthermore, we achieve temporal consistency for each sensor through temporal contrasting by following prior works (Choi and Kang 2023; Eldele et al. 2021). Due to the dynamic nature of sensor correlations in MTS data (Wang et al. 2023), we propose segmenting a sample into multiple windows, enabling us to incorporate multi-window temporal contrasting which ensures the consistency of temporal patterns within each sensor.

In summary, our contributions are three folds. First, to promote spatial consistency, we propose novel graph augmentations to enhance the quality of augmented views for MTS data. The graph augmentations involve node and edge augmentations, aiming to augment sensors and their correlations respectively. Second, we design graph contrasting that includes node- and graph-level contrasting, facilitating the learning of robust sensor- and global-level features. We also introduce a multi-window temporal contrasting to achieve temporal consistency for each sensor. Third, we conduct extensive experiments on ten public MTS datasets, showing that our TS-GAC achieves state-of-the-art performance.

Related Work

Contrastive Learning (CL) As a self-supervised learning paradigm, CL has gained popularity due to its ability to learn effective features from unlabeled samples by bringing positive pairs closer while pushing negative pairs farther (Zhang et al. 2023a; Eldele et al. 2023). Augmentations are commonly used to create positive pairs, generating augmented samples from different perspectives. Negative pairs, on the other hand, are created using the remaining samples in the same batch (Chen et al. 2020a) or stored in a memory bank (He et al. 2020). Contrasting these positive and negative pairs helps encoders become robust to perturbations, ensuring consistency in the learned features, and thus learning robust and effective features from unlabeled data.

Researchers have proven the effectiveness of CL in image tasks (Hjelm et al. 2018; He et al. 2020; Caron et al. 2020; Chen et al. 2020a). MoCo (He et al. 2020) designed a momentum encoder with a memory bank to store negative samples, achieving desirable performance with limited computational resources. SimCLR (Chen et al. 2020a) adopted larger batches of negative pairs and achieved comparable results to supervised learning. Inspired by SimCLR, MoCo-v2 (Chen et al. 2020b) improved performance with powerful augmentations without requiring large batches. Besides, negative pairs may occupy computation resources, so BYOL (Grill et al. 2020) and SimSiam (Chen and He 2021) learned representations with only positive pairs. Although these methods have achieved decent performance, they are proposed for images. Different from images, MTS data contain both temporal and spatial information from multiple sensors, making traditional image-based augmentation and contrasting methods unsuitable for MTS data.

CL for MTS Data Pioneering works have successfully utilized CL techniques to learn decent representations from unlabeled MTS data, primarily focusing on achieving temporal consistency (Pöppelbaum, Chadha, and Schwung 2022; Khaertdinov, Ghaleb, and Asteriadis 2021; Hao et al. 2023; Yue et al. 2022; Eldele et al. 2021). Specifically, they augmented MTS data with temporal augmentations such as jittering, cropping, and sub-series, and then conducted CL to ensure encoders robustness to temporal disturbances. Meanwhile, some works (Choi and Kang 2023; Eldele et al. 2021) also introduced temporal contrasting by summarizing past information for contrasting with future timestamps, further enforcing robustness to perturbations within timestamps.

While these works advanced CL for MTS data by ensuring temporal consistency, they overlook spatial consistency for MTS data. Some recent works proposed to incorporate spatial information, e.g., sensor correlations, into CL frameworks. For example, TAGCN (Zhang et al. 2023d) utilized GNN to extract features from sub-series of MTS data and then performed CL. Additionally, TSGCC (Zhang et al. 2023b) designed a graph-based method to compute weights between samples for clustering by instance- and clustering-contrasting. However, these methods only utilized GNN to extract spatial information within MTS data, while still overlooking spatial consistency to achieve better CL for MTS data. Although a few recent studies (Chen et al. 2022; Li

et al. 2022) explored channel-wise signal augmentations, graph-level augmentations and contrasting are still under-explored, limiting their ability to achieve robust spatial consistency for MTS data.

To overcome the limitations, we propose TS-GAC which incorporates both the graph augmentation and graph contrasting techniques to ensure spatial consistency during the CL process for MTS classification.

Methodology

Problem Formulation

Given a dataset with n unlabeled MTS samples $\mathcal{X} = \{X_j\}_{j=1}^n$, each sample $X_j \in \mathbb{R}^{N \times L}$ is collected from N sensors with T timestamps. Our objective is to perform contrastive learning scheme that can achieve spatial consistency for MTS data, enabling the training of an encoder \mathcal{F} without relying on labels. This approach allows us to achieve enhanced CL performance and thus extract effective representations $h_j = \mathcal{F}(X_j) \in \mathbb{R}^d$. With h_j , we employ a simple classifier, e.g., a multi-layer perceptron, to obtain class probabilities $y_j \in \mathbb{R}^c$, where c represents the number of classes in the classification task. For simplicity, the subscript j is removed, and we denote an MTS sample as X .

Overall Structure

Fig. 2 shows the overall structure of TS-GAC, which aims to achieve spatial consistency in CL for MTS classification. Specific augmentation and contrasting techniques are tailored for MTS data. For augmentation, we consider node and edge augmentations to augment individual sensors and their correlations, generating weak and strong views for each sample. Node frequency augmentations are applied first, followed by segmenting augmented samples into multiple windows considering the dynamic local patterns in MTS data. Node temporal augmentations are utilized within each window, followed by a 1-Dimensional Convolutional Neural Network (1D-CNN) to process these windows. Subsequently, graphs are constructed with each sensor as a node and sensor correlations are edges. The constructed graphs are further augmented by edge augmentations, and then processed by a GNN-based encoder to learn representations. Next, to achieve spatial consistency, we design graph contrasting including Node-level Contrasting (NC) and Graph-level Contrasting (GC). NC enables the contrasting of sensors within each sample to learn robust sensor-level features while GC allows the contrasting of samples within each training batch, promoting the learning of robust global-level features. We further introduce Multi-Window Temporal Contrasting (MWTC) to ensure temporal consistency for each sensor, by utilizing past windows in one view to predict the future windows in another view.

Augmentation

CL learns robust representations by contrasting different views of unlabeled data, which are commonly generated by augmentations. Then, the augmented views from the same data are pulled closer and the views from different data

are simultaneously pushed farther for representation learning. Thus, augmentations are critical for CL to learn robust and generalizable representations. To enhance augmentation quality for MTS data, we consider its multi-source nature, i.e., collected from multiple sensors (Zhao et al. 2019). We argue that augmentations for MTS data should be able to ensure the learning of robust sensor features and sensor correlations. For this purpose, we design node and edge augmentations that augment individual sensors and their correlations respectively. Further, following (Eldele et al. 2021), we generate weak and strong views, i.e., weakly and strongly augmented, for each sample with the augmentations for subsequent contrasting.

Node Augmentations We perform both frequency and temporal augmentations for the nodes (i.e., sensors).

Frequency augmentations: We utilize frequency augmentations to augment individual sensors, as the augmentations are widely recognized as effective in augmenting time-series data (Zhang et al. 2022, 2023c). This involves transforming the signals of each sensor into the frequency domain and augmenting the extracted frequency features. The augmented frequency features are then transformed back into the temporal domain to obtain augmented signals.

Particularly, we adopt Discrete Wavelet Transform (DWT) (Boggess et al. 2002) to decompose signals into detail and approximation coefficients using high-pass and low-pass filters, representing detailed and general trends within the signals, respectively. To generate weak and strong views, we add Gaussian noise to the detail and approximation coefficients respectively. The augmented frequency features are then transformed back into the temporal domain using inverse DWT (iDWT) to obtain the augmented signals. Mathematically, frequency augmentations are achieved via Eq. (1), where $\eta_{A,i}$ and $\eta_{D,i}$ denote the approximation and detail coefficients for the i -th sensor, and ξ represents the noise added to coefficients. We denote $\{X^w, X^s\}$ as the augmented signals in weak and strong views.

$$\begin{aligned} \eta_{A,i}, \eta_{D,i} &= DWT(x_i), \\ \eta_{A,i}^s &= \eta_{A,i} + \xi, \eta_{D,i}^w = \eta_{D,i} + \xi, \\ x_i^s &= iDWT(\eta_{A,i}^s, \eta_{D,i}), x_i^w = iDWT(\eta_{A,i}, \eta_{D,i}^w). \end{aligned} \quad (1)$$

Temporal augmentations: We further introduce temporal augmentations to augment each sensor due to its importance in augmenting time-series data (Pöppelbaum, Chadha, and Schwung 2022; Khaertdinov, Ghaleb, and Asteriadis 2021). Before temporal augmentations, we note that MTS data show dynamic properties, i.e., local patterns of MTS data are dynamically changing (Wang et al. 2023). To capture such properties, we segment each MTS sample into mini windows. As displayed in Fig. 3, given the window with length f , we segment an MTS sample into $\bar{L} = \lfloor L/f \rfloor$ windows, where $\lfloor \cdot \rfloor$ represents truncation. Thus, we obtain $X^w = \{\bar{X}_t^w\}_{t=1}^{\bar{L}}$ for the weak view, where t is the index of the window, and $\bar{X}_t^w = \{\bar{x}_{t,i}^w\}_{i=1}^N \in \mathbb{R}^{N \times f}$ contains the local patterns, including local sensor features and correlations. The windows in the strong view $\{\bar{X}_t^s\}_{t=1}^{\bar{L}}$ are obtained in the same way. In this case, if we conduct temporal augmentations before segmentation, it is hard to augment each

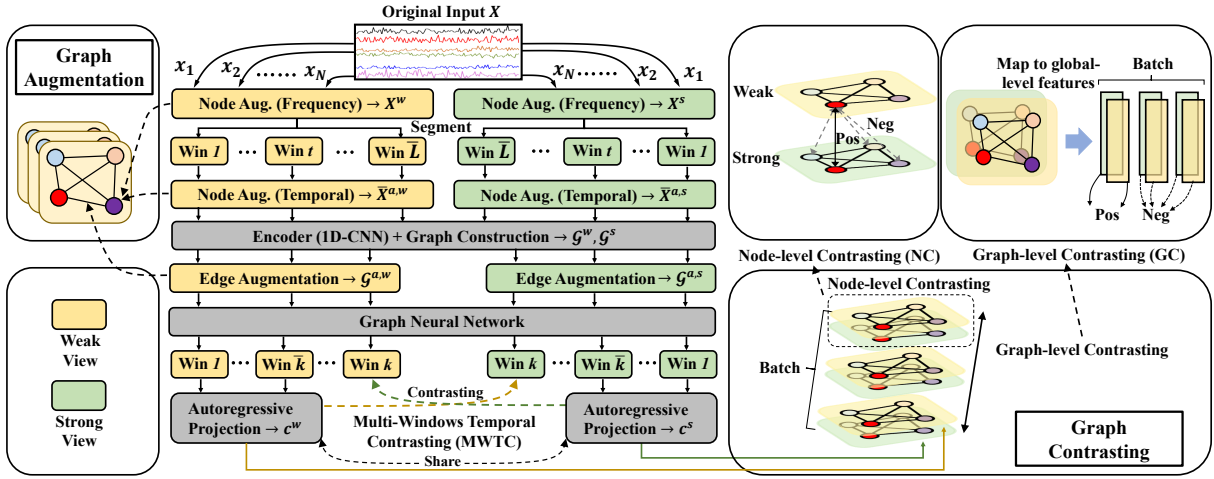


Figure 2: Overall structure of TS-GAC. (1) Graph augmentations to augment MTS data effectively, generating weak and strong views. The graph augmentations involve node and edge augmentations, where node augmentations include both frequency and temporal augmentations to fully augment sensors. Node frequency augmentations are first applied, followed by segmenting augmented samples into multiple windows by considering the dynamic local patterns in MTS data. Node temporal augmentations are utilized within each window, followed by 1D-CNN to process these windows. Subsequently, graphs are constructed and augmented through edge augmentations, and then processed by GNN. (2) Graph contrasting includes NC and GC to achieve spatial consistency. NC ensures robust sensors by pulling closer corresponding sensors in different views and pushing farther different sensors in those views within each sample. GC ensures robust global features by pulling closer corresponding samples in different views and pushing farther different samples in those views within each batch. MWTC further achieves temporal consistency for each sensor by summarizing past windows to contrast with future windows in another view.




Figure 3: The multi-window segmentation to generate multiple windows for one MTS sample.

window averagely, so we propose augmenting each window after segmentation.

We adopt permutation for temporal augmentations due to its wide application (Eldele et al. 2021; Pöppelbaum, Chadha, and Schwung 2022) and augment each sensor of each window. After augmentation, we obtain the augmented windows, e.g., $\{\bar{X}_t^{a,w}\}_{t=1}^L$ in the weak view, where $\bar{X}_t^{a,w} = \{\bar{x}_{t,i}^{a,w}\}_{i=1}^N$. 1D-CNN is then utilized as an encoder to capture the temporal information between windows (Jin et al. 2022), whose details are attached in our supplementary materials. With the encoder, we learn updated windows, e.g., $\{Z_t^w\}_{t=1}^k$ for the weak view, where $Z_t^w = \{z_{t,i}^w\}_{i=1}^N$. Similar notations such as $\bar{X}_t^{a,s}$ and Z_t^s apply to the strong view.

Edge Augmentations The correlations between sensors should remain robust due to their importance for learning

sensor features (Jia et al. 2020; Zhang, Zhang, and Tsung 2022). To ensure robust sensor relationships, we begin by constructing graphs whose nodes and edges represent sensors and the correlations between these sensors respectively. Augmenting the edges allows us to augment the relations effectively. For graph construction, we note that correlated sensors should follow similar properties and their features should be similar in the feature space, so we leverage the features similarities to define the sensor correlations. Given the features $Z_t = \{z_{t,i}\}_{i=1}^N \in \mathbb{R}^{N \times f}$, we compute the correlation between sensors i and j using the dot product of their features, i.e., $e_{t,ij} = z_{t,i}(z_{t,j})^T$. Then, the softmax function is used to restrict the correlations within the range $[0,1]$. Multiple graphs are built based on the windows for two views. For the weak view, the graph for t^{th} window is denoted as $G_t^w = (Z_t^w, E_t^w)$, where $E_t^w = \{e_{t,ij}^w\}_{i,j}^N$. Similar graphs G_t^s are obtained for the strong view.

We then introduce edge augmentations to augment the correlations between sensors. A naive approach would be randomly adding noise, replacing, or dropping certain edges for graph augmentation (You et al. 2020). However, this method may introduce excessive bias and significantly alter the topological structure within MTS data. Note that GNN updates sensor features based on their correlations with other sensors. Thus, strong correlations ensure more information propagation, making them more crucial than weak correlations. Randomly disturbing these strong correlations can introduce excessive bias. To address this issue, it is necessary to add constraints for the edge augmentation. Thus, we propose retaining the s strongest correlations (i.e., top- s

correlations) for each sensor and augmenting the remaining correlations by replacing them with random values within the range [0, 1]. This approach allows us to fully augment sensor correlations while preserving the topological information within MTS data as much as possible. Specifically, we retain more strong correlations for graphs in the weak view and fewer strong correlations for graphs in the strong view. The resulting augmented graph for the t^{th} window in the weak view is denoted as $\mathcal{G}_t^{a,w} = (Z_t^w, E_t^{a,w})$, and $E_t^{a,w}$ are augmented sensor correlations. Similarly, $\mathcal{G}_t^{a,s}$ denotes the augmented graph for the strong view.

With the augmented graphs, we adopt GNN to update sensor features by leveraging the augmented correlations as conventional works did (Jia et al. 2020; Wang et al. 2023). Particularly, the features for sensor i in the weak view are updated by a nonlinear function, i.e., $z_{t,i}^w = \sigma(\sum_j z_{t,j}^w e_{t,i,j}^{a,w} W_g)$, where W_g are learnable weights. The updated sensor features $z_{t,i}^w$ and $z_{t,i}^s$ in weak and strong views are then used for subsequent contrasting.

Contrasting

With the augmentations to generate weak and strong views, we design graph contrasting to achieve spatial consistency and further design MWTC to achieve temporal consistency for each sensor. We begin by presenting MWTC in this section, as it learns high-level sensor features within multi-window for subsequent graph contrasting.

Multi-Window Temporal Contrasting MWTC operates at the sensor-level, ensuring temporal consistency for each sensor. It is noted that the multi-window of each sensor show temporal dependencies where future windows are normally affected and dependent on past windows, which can be incorporated to keep the multi-window robust. Inspired by the idea of predictive coding (Oord, Li, and Vinyals 2018) and temporal contrasting (Choi and Kang 2023; Eldele et al. 2021), we propose to summarize past windows in one view to contrast with the future windows in another view. By doing so, we aim to maintain the temporal dependency robustness against perturbations to the windows, enabling that the temporal patterns within MTS data are preserved.

Specifically, we introduce an auto-regressive model f_a to summarize the sensor features in past \bar{k} windows, i.e., $c_i^w = f_a(z_{1,i}^w, \dots, z_{\bar{k},i}^w | W_a)$, representing the summarized vectors for the i -th sensor in the weak view. c_i^w is then to predict future windows, i.e., $\bar{z}_{\bar{k}+1,i}^w = f_{\bar{k}+1}(c_i^w), \dots, \bar{z}_{k,i}^w = f_k(c_i^w)$, where $f_{\bar{k}+1}(\cdot), \dots, f_k(\cdot)$ are nonlinear functions to predict the $(\bar{k}+1)$ -th, ..., k -th windows. Similar operations are conducted for the strong view. Here, we adopt a transformer model for f_a following (Eldele et al. 2021), the detail of which is attached in our supplementary materials. $\mathcal{L}_{MWTC}^{s \rightarrow w}$ in Eq. (2) is the loss using the past windows in the strong view to predict the future windows in the weak view. Here, the predicted window $\bar{z}_{t,i}^s$ should exhibit similarity with its positive pair $z_{t,i}^w$, while being dissimilar with its negative pairs $z_{v,i}^w, v \in \hat{\mathcal{V}}_{t,i}$, where $\hat{\mathcal{V}}_{t,i}$ denotes the set of windows excluding the t -th window for sensor i .

$$\mathcal{L}_{MWTC}^{s \rightarrow w} = \frac{-1}{N(k-\bar{k})} \sum_i^N \sum_{t=\bar{k}}^k \log \frac{\exp((\bar{z}_{t,i}^s)^T z_{t,i}^w)}{\sum_{v \in \hat{\mathcal{V}}_{t,i}} \exp((\bar{z}_{t,i}^s)^T z_{v,i}^w)}. \quad (2)$$

Similarly, we can obtain $\mathcal{L}_{MWTC}^{w \rightarrow s}$ and thus obtain $\mathcal{L}_{MWTC} = \mathcal{L}_{MWTC}^{s \rightarrow w} + \mathcal{L}_{MWTC}^{w \rightarrow s}$ for sample X .

Graph Contrasting We propose graph contrasting to achieve spatial consistency, including Node-level Contrasting (NC) and Graph-level Contrasting (GC) to learn robust sensor- and global-level features. NC is achieved by contrasting sensors in different views within each MTS sample while GC is achieved by contrasting the samples within each training batch. Notably, we leverage the vectors $\{c_i\}_{i=1}^N$ for graph contrasting, as the vectors represent the high-level features by summarizing the sensor-level features within multi-window. By utilizing the high-level features, we can achieve more effective graph contrasting.

Node-level Contrasting: NC is designed to learn robust sensor-level features. Specifically, it aims to maximize the similarity between the corresponding sensors in two views while minimizing the similarity between different sensors in those views. By doing so, NC encourages the encoder to learn features against perturbations to each sensor. Eq. (3) presents the node-level contrastive loss, where $\hat{\mathcal{V}}_i$ denotes the set of sensors excluding sensor i . The visualization process is shown in NC of Fig. 2.

$$\mathcal{L}_{NC}^{s \rightarrow w} = -\frac{1}{N} \sum_i^N \log \frac{\exp(f_{sim}(c_i^s, c_i^w)/\tau)}{\sum_{v \in \hat{\mathcal{V}}_i} \exp(f_{sim}(c_i^s, c_v^w)/\tau)}. \quad (3)$$

Here $f_{sim}(a, b)$ is a function to measure the similarity of samples implemented as the dot product $a^T b$, and τ is a temperature parameter. $\mathcal{L}_{NC}^{s \rightarrow w}$ denotes that the sensors in the strong view are contrasted with the positive and negative pairs in the weak view. Similarly, we can obtain $\mathcal{L}_{NC}^{w \rightarrow s}$ and thus obtain $\mathcal{L}_{NC} = \mathcal{L}_{NC}^{s \rightarrow w} + \mathcal{L}_{NC}^{w \rightarrow s}$ for sample X .

Graph-level Contrasting: GC aims to learn robust global-level features by contrasting samples within each training batch. For subsequent contrasting, we here obtain the global-level features by stacking all sensor features. For the weak view, $g^w = [c_1^w | \dots | c_N^w]$, where $[\]$ denotes concatenation. Similar operations are conducted for the strong view.

To learn robust global-level features, GC is achieved by maximizing the similarity between the corresponding samples in two views and simultaneously minimizing the similarity between the different samples in those views. Given a batch of B MTS samples, we have $2B$ augmented samples from two augmented views. The corresponding samples in two views are treated as positive pairs, and each view of the sample can form $2B-2$ negative pairs with the remaining augmented samples. We denote the global-level features of the p -th augmented samples in weak and strong views within the batch as $g_p^{\{w,s\}}$. Accordingly, the graph-level contrasting is demonstrated as Eq. (4), which denotes that the samples in the strong view are contrasted with the remaining augmented samples in the batch. Here, $\hat{\mathcal{V}}_p$ denotes the set of samples in

the batch excluding the p -th sample.

$$\mathcal{L}_{GC}^s = -\frac{1}{B} \sum_{p=1}^B \log \frac{\exp(f_{sim}(g_p^s, g_p^w)/\tau)}{\sum_{v \in \hat{V}_p} \exp(f_{sim}(g_p^s, g_v^{\{w,s\}})/\tau)}. \quad (4)$$

Similarly, we can obtain \mathcal{L}_{GC}^w for the weak view and thus obtain $\mathcal{L}_{GC} = \mathcal{L}_{GC}^s + \mathcal{L}_{GC}^w$.

Finally, we combine MWTC, NC, and GC to form the final self-supervised loss as Eq. (5), where λ_{MWTC} , λ_{NC} , and λ_{GC} are hyperparameters that denote relative weights of the losses. Notably, MWTC and NC are both achieved for each MTS sample, so they are denoted as $\mathcal{L}_{p,MWTC}$ and $\mathcal{L}_{p,NC}$ for the p -th sample.

$$\mathcal{L} = \lambda_{MWTC} \sum_p^B \mathcal{L}_{p,MWTC} + \lambda_{NC} \sum_p^B \mathcal{L}_{p,NC} + \lambda_{GC} \mathcal{L}_{GC}. \quad (5)$$

Experimental Results

Datasets We examine our method on ten public MTS datasets for classification, including Human Activity Recognition (HAR) (Anguita et al. 2012), ISRUC (Khalighi et al. 2016), and eight large datasets from UEA archive, i.e., ArticulatoryWordRecognition (AWR), FingerMovements (FM), SpokenArabicDigitsEq (SAD), CharacterTrajectories (CT), FaceDetection (FD), InsectWingbeat (IW), MotorImagery (MI), and SelfRegulationSCP1 (SRSCP1). For HAR and ISRUC, we randomly split them into 80% and 20% for training and testing, while for those from UEA archive, we directly adopt their pre-defined train-test splits. *The statistics of the datasets are in the Appendix.*

Evaluation For evaluation, we follow the standard linear classification scheme as current methods did (Eldele et al. 2021; Yue et al. 2022), i.e., train an encoder with only training data in a self-supervised manner and then train a linear classifier on top of the pre-trained encoder. To evaluate performance, we adopt two metrics, i.e., Accuracy (Accu.) and Macro-averaged F1-Score (MF1) (Eldele et al. 2021; Meng et al. 2022). Besides, to reduce the effect of random initialization, we conduct ten times for all experiments and take the average results for comparisons. The standard variations are reported to show the robustness of the results.

Implementation Details All methods are conducted with NVIDIA GeForce RTX 3080Ti and implemented by PyTorch (Paszke et al. 2019). We set the batch size as 128 and choose ADAM as the optimizer with a learning rate of $3e-4$. We pre-train the model and train the linear classifier 40 epochs. *More implementation details are in the Appendix.*

Comparisons with State-of-the-Arts

We compare our method with SOTA methods, including TNC (Tonekaboni, Eytan, and Goldenberg 2021), TS-TCC (Eldele et al. 2021), TS2Vec (Yue et al. 2022), MHCCCL (Meng et al. 2022), CaSS (Chen et al. 2022), and TAGCN (Zhang et al. 2023d). All methods are re-implemented based on their original settings except for the

encoders, which are replaced by the same encoder as ours for fair comparisons.

Table 1 shows the comparisons with SOTA methods. From the table, we observe that TS-GAC achieves the best performance on eight out of ten datasets. Particularly, TS-GAC gains great improvements on HAR and ISRUC, improving by 1.44% and 3.13% respectively regarding accuracy. In the remaining cases where TS-GAC achieves the second best, the gaps of TS-GAC with the best result are marginal, e.g., only 0.4% lower than the best accuracy for FM. Meanwhile, TS-GAC has smaller variances, indicating that our TS-GAC is more robust and stable.

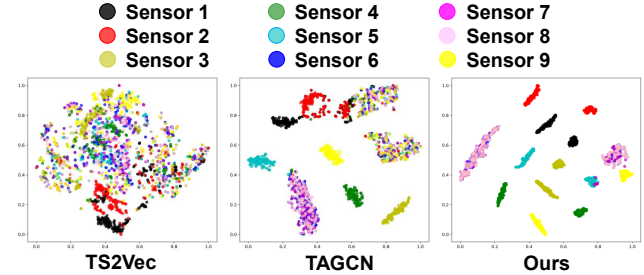


Figure 4: Visualization for sensor features.

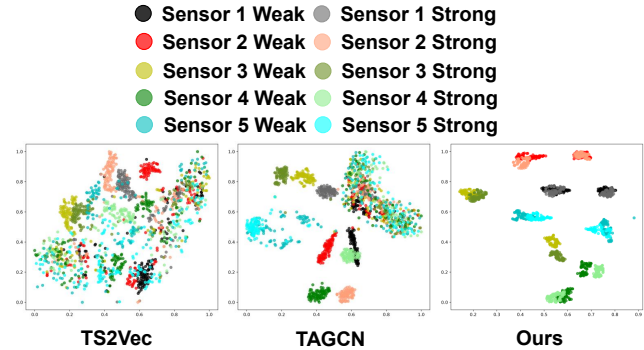


Figure 5: Visualization for spatial consistency.

The superior performance can be attributed to the spatial consistency achieved by TS-GAC. To intuitively demonstrate this spatial consistency, we visualized sensor features from different views, comparing TS-GAC with two competitive methods, TS2Vec and TAGCN. We first visualized the individual sensor features. As shown in Fig. 4, TS-GAC exhibits clearer sensor clusters than TS2Vec and TAGCN, emphasizing its ability to learn robust sensor features. Based on the clear sensor features, the features extracted from weak and strong views are aligned. Specifically, TS-GAC obtains closer feature clusters for the same sensors in weak and strong views, demonstrating its capability to learn consistent sensor features across different perspectives.

Ablation Study

We evaluate designed augmentation and contrasting techniques within TS-GAC, which fall into two categories of

Datasets	Metrics	TNC	TS-TCC	TS2Vec	MHCCL	CaSS	TAGCN	TS-GAC (Ours)
HAR	Accu	81.10±1.88	91.66±0.42	92.78±0.32	82.95±0.55	82.64±0.31	92.83±0.28	94.27±0.12
	MF1	78.24±2.91	91.86±0.40	<u>92.78±0.33</u>	82.70±0.62	82.34±0.31	92.66±0.29	94.07±0.14
ISRUC	Accu	77.69±1.28	80.50±0.42	76.32±0.48	74.71±0.98	81.09±0.19	77.21±0.21	84.22±0.17
	MF1	64.08±1.60	79.12±0.40	74.44±0.59	72.09±1.23	<u>79.73±0.29</u>	76.23±0.27	83.45±0.23
AWR	Accu	82.60±4.21	89.44±0.68	98.30±0.09	93.00±0.56	97.47±0.16	97.87±0.27	98.33±0.08
	MF1	77.42±5.34	89.51±0.73	<u>98.29±0.10</u>	93.14±0.75	97.46±0.16	97.86±0.27	98.33±0.07
FM	Accu	48.90±2.42	47.40±1.63	47.10±4.22	52.40±2.28	50.00±1.79	51.50±1.91	<u>52.00±1.54</u>
	MF1	43.02±5.25	47.36±1.64	47.03±4.18	49.82±3.06	35.10±2.01	<u>49.52±2.04</u>	48.78±0.71
SAD	Accu	90.30±1.36	95.20±0.15	97.31±0.19	95.91±0.56	97.44±0.07	<u>97.50±0.03</u>	97.99±0.05
	MF1	88.83±1.42	95.24±0.15	97.31±0.19	95.92±0.45	97.45±0.07	<u>97.52±0.03</u>	97.99±0.05
CT	Accu	96.23±1.24	98.61±0.17	98.68±0.02	98.21±0.10	97.16±0.07	98.89±0.10	98.82±0.05
	MF1	95.98±1.56	98.49±0.19	98.59±0.02	95.62±0.11	96.92±0.08	98.81±0.09	<u>98.77±0.06</u>
FD	Accu	50.15±0.61	58.00±1.71	<u>59.60±0.61</u>	55.26±1.12	54.38±0.47	58.21±0.74	60.53±0.31
	MF1	41.59±1.03	57.83±2.24	<u>59.20±0.67</u>	53.10±1.92	54.29±0.48	57.68±1.02	60.47±0.51
IW	Accu	30.19±0.27	56.08±1.22	<u>58.60±0.35</u>	29.30±2.34	24.45±2.35	58.07±0.31	65.80±0.36
	MF1	28.86±1.02	55.72±1.22	<u>58.16±0.42</u>	24.29±2.46	22.29±2.40	57.90±0.30	65.49±0.48
MI	Accu	52.40±3.12	51.70±4.63	<u>53.00±0.49</u>	52.45±0.78	51.00±1.67	50.00±2.42	56.00±0.46
	MF1	52.59±3.85	46.53±5.88	48.87±0.58	38.64±1.12	17.54±2.45	46.94±3.21	<u>50.25±0.36</u>
SRSCP1	Accu	76.76±1.27	83.64±0.99	82.94±1.67	82.31±1.01	83.95±1.15	82.18±1.10	84.47±1.19
	MF1	75.90±1.02	83.61±0.99	82.92±1.68	81.81±1.00	<u>83.81±1.15</u>	82.18±1.10	84.44±1.19

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

TS-GAC (Variants)		Augmentations		w/o GC	Contrasting		Complete
		w/o Aug. (N)	w/o Aug. (E)		w/o NC	w/o MWTC	
HAR	Accu	92.97±0.23	93.67±0.11	92.10±0.09	92.29±0.27	93.60±0.19	94.27±0.12
	MF1	92.69±0.27	93.41±0.12	91.76±0.11	92.03±0.30	93.38±0.21	94.07±0.14
ISRUC	Accu	83.86±0.20	83.87±0.18	81.70±0.14	81.29±0.12	81.29±0.34	84.22±0.17
	MF1	82.88±0.28	82.80±0.15	80.62±0.13	80.11±0.11	80.11±0.87	83.45±0.23

Table 2: Ablation study for graph augmentation and graph contrasting (%)

variants. The first category tests augmentations, including w/o Aug. (N) and w/o Aug. (E), representing variants without node and edge augmentations, respectively. The second category assesses the effectiveness of contrastive losses, with variants w/o GC, w/o NC, and w/o MWTC indicating the removal of graph-level contrasting, node-level contrasting, and multi-window temporal contrasting, respectively. Finally, we compare them with the complete TS-GAC.

Table 2 shows the results, where we only present the results on HAR and ISRUC due to limited space. *More results can be found in our supplementary materials.* The experimental results demonstrate the effectiveness of our proposed graph augmentation and contrasting techniques in achieving spatial consistency for MTS data. Specifically, the graph augmentations show significant improvements in learning robust representations. Compared to the variant without node augmentations, our complete TS-GAC achieves improvements of 1.30% and 0.36% on the two datasets. Similarly, compared to the model without edge augmentations, our complete TS-GAC achieves improvements of 0.60% and

0.35% on the two datasets. The improvements indicate the necessity of using graph augmentations for better augmenting MTS data. Meanwhile, the designed contrasting techniques play crucial roles in learning robust representations, and our complete TS-GAC achieves the best performance compared to the variants without any of the contrastive losses. For instance, we see drops of 2.17% and 2.52% by removing GC and drops of 1.98% and 2.93% by removing NC on the two datasets, indicating the effectiveness of graph contrasting in achieving spatial consistency. We further observe drops of 0.67% and 2.90% by removing MWTC on the two datasets, showing the importance of achieving temporal consistency for each sensor. Additionally, we can derive from the results that TS-GAC can still achieve good performance even when only graph contrasting is used, further highlighting the effectiveness of graph contrasting.

Overall, these findings validate the importance of our proposed graph augmentation and contrasting techniques, demonstrating the necessity of achieving spatial consistency when conducting CL for MTS data.

Sensitivity Analysis

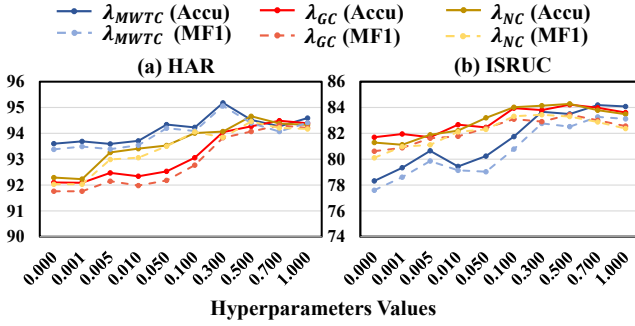


Figure 6: Sensitivity analysis for λ_{MWTC} , λ_{GC} , and λ_{NC} .

Hyperparameter Analysis We analyze λ_{MWTC} , λ_{GC} , and λ_{NC} to test their effects. The hyperparameters are trade-offs between various losses, so we choose the values within [0, 0.001, 0.005, 0.01, 0.05, 0.1, 0.3, 0.5, 0.7, 1.0]. To reduce computation costs, we fixed other hyperparameters as 1 when testing one of them. From Fig. 6, we observe that TS-GAC tends to achieve better performance when the hyperparameters are set as larger values. For example, the accuracy increases by 2% with λ_{GC} increasing from 0 to 0.7 on HAR. The improvements show that the contrastive losses have positive effects on CL performance. However, the improvements become small when the values are large enough. For example, the performance has no obvious improvements with increasing λ_{GC} from 0.7 to 1. Similar trends can also be found in other hyperparameters. Therefore, we can derive that the large hyperparameters have positive effects on the performance; however, too large values are unnecessary.

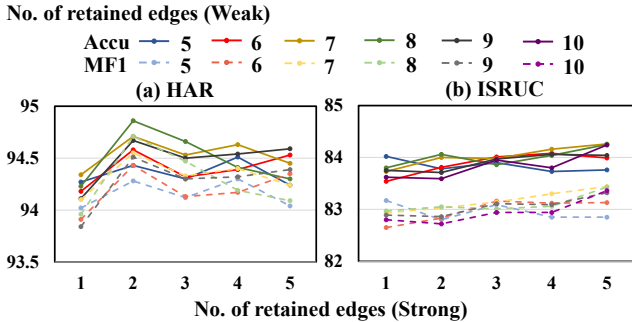


Figure 7: Sensitivity analysis for retained edges in views.

Number of retained edges for edge augmentations To effectively augment sensor correlations, we design edge augmentations by retaining the s strongest correlations, i.e., edges, for each sensor and replacing remaining correlations with random values. The value of s is crucial for augmenting sensor relations and thus requires testing. Here, the weak view should have larger s for weak augmentation while the strong view should have smaller s for strong augmentation. Meanwhile, each sensor in HAR and ISRUC has 9 and 10 edges respectively. Thus, we set s in the weak view within

[5,6,7,8,9] for HAR and add 10 for ISRUC. For the strong view, we set s within [1, 2, 3, 4, 5] for both datasets. Fig. 7 shows the results on HAR and ISRUC, where no. of retained edges represents the value of s . We take the results in HAR for example, and observe that our model shows better performance when s in the strong view is set to 2 while keeping s in the weak view fixed. On the other hand, our model shows better performance when s in the weak view is set to 7 or 8 while keeping s in the strong view fixed. These trends indicate that having fewer retained correlations in the strong view has a positive effect, but the value of s should not be too small so as to avoid overly distorted correlations. Similarly, having more retained correlations in the weak view is beneficial, but the value of s should not be too large.

Conclusion

We propose TS-GAC for MTS classification. To achieve spatial consistency, specific augmentation and contrasting techniques are tailored for MTS data. To better augment MTS data, graph augmentations are proposed, including node and edge augmentations for ensuring the robustness of sensors and their correlations. Besides, graph contrasting is designed, including node- and graph-level contrasting to extract robust sensor- and global-level features. We further introduce multi-window temporal contrasting to ensure temporal consistency for each sensor. Experiments show that TS-GAC achieves SOTA performance in various MTS classification tasks.

Acknowledgements

We thank anonymous reviewers for their constructive comments on this work. This research is supported by the Agency for Science, Technology and Research (A*STAR) under its AME Programmatic Funds (Grant No. A20H6b0151) and Career Development Award (Grant No. C210112046), and the National Research Foundation, Singapore under its AI Singapore Programme (AISG2-RP-2021-027).

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 *International workshop on ambient assisted living*, 216–223. Springer.
- Boggess, A.; Narcowich, F. J.; Donoho, D. L.; and Donoho, P. L. 2002. A first course in wavelets with fourier analysis. *Physics Today*, 55(5): 63.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33: 9912–9924.
- Chen, K.; Zhang, D.; Yao, L.; Guo, B.; Yu, Z.; and Liu, Y. 2021. Deep learning for sensor-based human activity recognition: Overview, challenges, and opportunities. *ACM Computing Surveys (CSUR)*, 54(4): 1–40.

- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020a. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Chen, X.; Fan, H.; Girshick, R.; and He, K. 2020b. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*.
- Chen, X.; and He, K. 2021. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15750–15758.
- Chen, Y.; Zhou, X.; Xing, Z.; Liu, Z.; and Xu, M. 2022. CaSS: A Channel-Aware Self-supervised Representation Learning Framework for Multivariate Time Series Classification. In *International Conference on Database Systems for Advanced Applications*, 375–390. Springer.
- Chen, Z.; Wu, M.; Zhao, R.; Guretno, F.; Yan, R.; and Li, X. 2020c. Machine remaining useful life prediction via an attention-based deep learning approach. *IEEE Transactions on Industrial Electronics*, 68(3): 2521–2531.
- Choi, H.; and Kang, P. 2023. Multi-Task Self-Supervised Time-Series Representation Learning. *arXiv preprint arXiv:2303.01034*.
- Craik, A.; He, Y.; and Contreras-Vidal, J. L. 2019. Deep learning for electroencephalogram (EEG) classification tasks: a review. *Journal of neural engineering*, 16(3): 031001.
- Deng, A.; and Hooi, B. 2021. Graph neural network-based anomaly detection in multivariate time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 4027–4035.
- Eldele, E.; Ragab, M.; Chen, Z.; Wu, M.; Kwoh, C.-K.; and Li, X. 2023. Label-efficient time series representation learning: A review. *arXiv preprint arXiv:2302.06433*.
- 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*.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33: 21271–21284.
- Hao, S.; Wang, Z.; Alexander, A. D.; Yuan, J.; and Zhang, W. 2023. MICOS: Mixed supervised contrastive learning for multivariate time series classification. *Knowledge-Based Systems*, 260: 110158.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.
- Hjelm, R. D.; Fedorov, A.; Lavoie-Marchildon, S.; Grewal, K.; Bachman, P.; Trischler, A.; and Bengio, Y. 2018. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*.
- Jia, Z.; Lin, Y.; Wang, J.; Zhou, R.; Ning, X.; He, Y.; and Zhao, Y. 2020. GraphSleepNet: Adaptive Spatial-Temporal Graph Convolutional Networks for Sleep Stage Classification. In *IJCAI*, 1324–1330.
- Jin, R.; Wu, M.; Wu, K.; Gao, K.; Chen, Z.; and Li, X. 2022. Position encoding based convolutional neural networks for machine remaining useful life prediction. *IEEE/CAA Journal of Automatica Sinica*, 9(8): 1427–1439.
- Khaertdinov, B.; Ghaleb, E.; and Asteriadis, S. 2021. Contrastive self-supervised learning for sensor-based human activity recognition. In *2021 IEEE International Joint Conference on Biometrics (IJCB)*, 1–8. IEEE.
- Khalighi, S.; Sousa, T.; Santos, J. M.; and Nunes, U. 2016. ISRUC-Sleep: A comprehensive public dataset for sleep researchers. *Computer methods and programs in biomedicine*, 124: 180–192.
- Li, R.; Zhong, T.; Jiang, X.; Trajcevski, G.; Wu, J.; and Zhou, F. 2022. Mining Spatio-Temporal Relations via Self-Paced Graph Contrastive Learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '22*, 936–944. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393850.
- Meng, Q.; Qian, H.; Liu, Y.; Xu, Y.; Shen, Z.; and Cui, L. 2022. MHCC: Masked Hierarchical Cluster-wise Contrastive Learning for Multivariate Time Series. *arXiv preprint arXiv:2212.01141*.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Pöppelbaum, J.; Chadha, G. S.; and Schwung, A. 2022. Contrastive learning based self-supervised time-series analysis. *Applied Soft Computing*, 117: 108397.
- Tonekaboni, S.; Eytan, D.; and Goldenberg, A. 2021. Unsupervised representation learning for time series with temporal neighborhood coding. *arXiv preprint arXiv:2106.00750*.
- Wang, Y.; Wu, M.; Li, X.; Xie, L.; and Chen, Z. 2023. Multivariate Time Series Representation Learning via Hierarchical Correlation Pooling Boosted Graph Neural Network. *IEEE Transactions on Artificial Intelligence*.
- Yang, L.; and Hong, S. 2022. Unsupervised time-series representation learning with iterative bilinear temporal-spectral fusion. In *International Conference on Machine Learning*, 25038–25054. PMLR.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph contrastive learning with augmentations. *Advances in neural information processing systems*, 33: 5812–5823.
- Yu, B.; Yin, H.; and Zhu, Z. 2017. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *arXiv preprint arXiv:1709.04875*.

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.

Zhang, K.; Wen, Q.; Zhang, C.; Cai, R.; Jin, M.; Liu, Y.; Zhang, J.; Liang, Y.; Pang, G.; Song, D.; et al. 2023a. Self-Supervised Learning for Time Series Analysis: Taxonomy, Progress, and Prospects. *arXiv preprint arXiv:2306.10125*.

Zhang, Q.; Liang, Z.; NGUEILBAYE, A.; Zhang, P.; Chen, J.; Chen, X.; and Huang, J. Z. 2023b. Graph-Augmented Contrastive Clustering for Time Series. *Available at SSRN 4474418*.

Zhang, W.; Yang, L.; Geng, S.; and Hong, S. 2023c. Self-Supervised Time Series Representation Learning via Cross Reconstruction Transformer. *IEEE Transactions on Neural Networks and Learning Systems*.

Zhang, W.; Zhang, C.; and Tsung, F. 2022. GRELEN: Multivariate Time Series Anomaly Detection from the Perspective of Graph Relational Learning. In Raedt, L. D., ed., *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, 2390–2397. International Joint Conferences on Artificial Intelligence Organization. Main Track.

Zhang, X.; Wang, Y.; Zhang, L.; Jin, B.; and Zhang, H. 2023d. Exploring unsupervised multivariate time series representation learning for chronic disease diagnosis. *International Journal of Data Science and Analytics*, 15(2): 173–186.

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.

Zhao, R.; Yan, R.; Chen, Z.; Mao, K.; Wang, P.; and Gao, R. X. 2019. Deep learning and its applications to machine health monitoring. *Mechanical Systems and Signal Processing*, 115: 213–237.