

Sub-Adjacent Transformer: Improving Time Series Anomaly Detection with Reconstruction Error from Sub-Adjacent Neighborhoods

Wenzhen Yue¹, Xianghua Ying^{1*}, Ruohao Guo¹, DongDong Chen², Ji Shi¹, Bowei Xing¹, Yuqing Zhu³ and Taiyan Chen¹

¹National Key Laboratory of General Artificial Intelligence, School of Intelligence Science and Technology, Peking University

²Microsoft Cloud + AI

³Tsinghua University

yuewenzhen@stu.pku.edu.cn, xhying@pku.edu.cn, dochen@microsoft.com

Abstract

In this paper, we present the Sub-Adjacent Transformer with a novel attention mechanism for unsupervised time series anomaly detection. Unlike previous approaches that rely on all the points within some neighborhood for time point reconstruction, our method restricts the attention to regions not immediately adjacent to the target points, termed *sub-adjacent neighborhoods*. Our key observation is that owing to the rarity of anomalies, they typically exhibit more pronounced differences from their sub-adjacent neighborhoods than from their immediate vicinities. By focusing the attention on the sub-adjacent areas, we make the reconstruction of anomalies more challenging, thereby enhancing their detectability. Technically, our approach concentrates attention on the non-diagonal areas of the attention matrix by enlarging the corresponding elements in the training stage. To facilitate the implementation of the desired attention matrix pattern, we adopt linear attention because of its flexibility and adaptability. Moreover, a learnable mapping function is proposed to improve the performance of linear attention. Empirically, the Sub-Adjacent Transformer achieves state-of-the-art performance across six real-world anomaly detection benchmarks, covering diverse fields such as server monitoring, space exploration, and water treatment.

1 Introduction

In modern industrial systems such as data centers and smart factories, numerous sensors consistently produce significant volumes of measurements [Zhou *et al.*, 2019; Shen *et al.*, 2020]. To effectively monitor the systems’ real-time conditions and avoid potential losses, it is crucial to identify anomalies in the multivariate time series [Pereira and Silveira, 2019; Chen, 2018]. This problem is known as time series anomaly detection [Blázquez-García *et al.*, 2021]. Comprehensive surveys can be found in the literature [Blázquez-García *et al.*, 2021; Schmidl *et al.*, 2022; Gupta *et al.*, 2013].

*Corresponding Author




Figure 1: Illustration of our method in time domain. The point marked with a red circle, along with its neighbors, represents an anomaly on the sinusoidal signal. (a) Previous works typically utilize the attention across all the points within the window, while (b) our method encourages the use of attention of sub-adjacent neighborhoods (the highlighted area). Such an imposed constraint enlarges the reconstruction challenge for anomalies and thus improves the anomaly detection performance.

Effective and robust time series anomaly detection remains a challenging and open problem [Zhao *et al.*, 2022]. Real-world time series usually exhibit nonlinear dependencies and intricate interactions among time points. Besides, the vast scale of data makes the labeling of these anomalies time-consuming and costly in practice [Lai *et al.*, 2023]. Therefore, time series anomaly detection is usually conducted in an unsupervised manner [Xu *et al.*, 2022], which is also the focus of this paper.

Before the era of deep learning, many classic anomaly detection methods were proposed. These include the density-estimation method proposed in [Breunig *et al.*, 2000], the clustering-based method presented in one-class SVM [Schölkopf *et al.*, 2001], and a series of SVDD works [Tax and Duin, 2004; Liu *et al.*, 2013], along with graph-based methods [Cheng *et al.*, 2008]. Traditional methods, which use handcrafted features, struggle to accurately express the relationship among time points and often suffer from poor generalization. Deep learning-based techniques have witnessed a diversity of methods during the recent decade. Broadly, they can be categorized into three main categories: Reconstruction-based methods reconstruct the time series and compare it with the original data [Park *et al.*, 2018;

Li *et al.*, 2021; Su *et al.*, 2019; Pintilie *et al.*, 2023], whereas prediction-based methods predict future points for comparison with actual data [Deng and Hooi, 2021; Hundman *et al.*, 2018a; Ding *et al.*, 2019; Lai *et al.*, 2023]. Dissimilarity-based methods emphasize the discrepancy representation between normal and abnormal points in various domains [Shen *et al.*, 2020; Sun *et al.*, 2021; Xu *et al.*, 2022]. It is noteworthy that these methods often combine multiple approaches to enhance performance. For example, TranAD [Tuli *et al.*, 2022] integrates an integrated reconstruction error and discriminator loss; while [Lai *et al.*, 2023] uses prediction results as the nominal data and combines them with reconstruction results.

In this paper, we focus on the attention matrix within the Transformer module. Transformers [Vaswani *et al.*, 2017] have achieved great success in natural language processing [Brown *et al.*, 2020], computer vision [Liu *et al.*, 2021] and time series [Xu *et al.*, 2022] in recent years. Previous work in [Xu *et al.*, 2022] primarily concentrated on the dissimilarities in attention distributions. Our paper takes a different approach and introduces a simple yet effective attention learning paradigm. Our fundamental assumption is that anomalies are less related to their non-immediate neighborhoods when compared with normal points. Therefore, focusing exclusively on these non-immediate neighborhoods is likely to result in larger reconstruction errors for anomalies. Based on this assumption, our study introduces two key concepts: *sub-adjacent neighborhoods* and *sub-adjacent attention contribution*. As shown in Figure 1, sub-adjacent neighborhoods indicate the areas not immediately adjacent to the target point, and the sub-adjacent attention contribution is defined as the sum of particular non-diagonal elements in the corresponding column of the attention matrix. These concepts are integrated with reconstruction loss, forming the cornerstone of our anomaly detection strategy. Furthermore, we observe that the traditional `Softmax` operation used in standard self-attention impedes the formation of the desired attention matrix, where the predefined non-diagonal stripes are dominant. In response, we adopt linear attention [Shen *et al.*, 2021; Han *et al.*, 2023] for its greater flexibility in attention matrix configurations. We also tailor the mapping function within this framework, using learnable parameters to enhance performance. The main contributions of this paper are summarized as follows.

- We propose a novel attention learning regime based on the sub-adjacent neighborhoods and attention contribution. Specifically, a new attention matrix pattern is designed to enhance discrimination between anomalies and normal points.
- Furthermore, we leverage the linear attention to achieve the desired attention pattern. To the best of our knowledge, this is the first introduction of linear attention with a learnable mapping function to time series anomaly detection.
- Extensive experiments show that the proposed Sub-Adjacent Transformer delivers state-of-the-art (SOTA) performance across six real-world benchmarks and one synthetic benchmark.

2 Related Works

Time Series Anomaly detection. In recent year, graph neural networks and self-attention have been explored in time series anomaly detection [Deng and Hooi, 2021; Zhao *et al.*, 2020; Boniol and Palpanas, 2022; Xu *et al.*, 2022; Lai *et al.*, 2023; Wu *et al.*, 2023]. The most related work to ours is Anomaly Transformer [Xu *et al.*, 2022], which also exploits the attention information. The discrepancy between the learned Gaussian distribution and the actual distribution is used to distinguish abnormal points from normal ones. A somewhat complicated min-max training strategy is adopted to train the model. Generally, Anomaly Transformer is complex and somewhat indirect. Our proposed method is more straightforward and exhibits improved performance.

Linear Attention. Compared with vanilla self-attention, linear attention enjoys more flexibility and lighter computation burden. Existing linear attention methods can be divided into three categories: pattern based methods, kernel based methods and mapping based methods. We mainly focus on mapping based methods in this paper. Efficient attention [Shen *et al.*, 2021] applies `Softmax` function to \mathbf{Q} (in a row-wise manner) and \mathbf{K} (in a column-wise manner) to ensure each row of $\mathbf{Q}\mathbf{K}^T$ sums up to 1. Hydra attention [Bolya *et al.*, 2022] studies the special case where attention heads are as many as feature dimension and \mathbf{Q} and \mathbf{K} are normalized so that $\|\mathbf{Q}_i\|$ and $\|\mathbf{K}_i\|$ share the same value. EfficientVit [Cai *et al.*, 2023] uses the simple ReLU function as the mapping function. FLatten Transformer [Han *et al.*, 2023] proposes the power function as the mapping function to improve the focus capability. In this paper, we apply `Softmax` to matrices \mathbf{Q} and \mathbf{K} both in a row-wise manner to approximate the focus property of the vanilla self-attention. Empirical experiments verify its performance superiority to other mapping methods in time series anomaly detection.

3 Methods

3.1 Problem Formulation

Let $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_T\} \in \mathbb{R}^{T \times D}$ denote a set of time series with $\mathbf{x}_t \in \mathbb{R}^D$, where T is the time steps and D is the number of channels. The label vector $\mathbf{y} = \{y_1, \dots, y_T\}$ indicates whether the corresponding time stamp is normal ($y_t = 0$) or abnormal ($y_t = 1$). Our task is to determine anomaly labels for all time points $\hat{\mathbf{y}} = \{\hat{y}_1, \dots, \hat{y}_T\}$, where $\hat{y}_t \in \{0, 1\}$, to match the ground truth \mathbf{y} as much as possible. Following the practice of [Xu *et al.*, 2022; Schmidl *et al.*, 2022; Tuli *et al.*, 2022; Lai *et al.*, 2023], we mainly focus on the F1 score with point adjustment.

3.2 Sub-Adjacent Neighborhoods

Time points usually have stronger connections with their neighbors and fewer connections with distant points. This characteristic is more pronounced for anomalies [Xu *et al.*, 2022]. As shown in Figure 1, if we rely solely on sub-adjacent neighborhoods to reconstruct time points, the reconstruction errors of anomalies will become more pronounced, thereby enhancing their distinguishability. This is the core idea of the Sub-Adjacent Transformer.




Figure 2: Illustration of attention contribution and the desired attention matrix. For clearness, only the main stripes are depicted.

In this paper, we define the *sub-adjacent neighborhoods* as the region where the distance to the target point is between K_1 and K_2 . $K_2 \geq K_1 > 0$. The highlighted red area in Figure 1 (b) show the sub-adjacent neighborhoods of the point marked with the red circle.

We now apply our thought to the attention matrix. The sub-adjacent neighborhoods are represented by the highlighted stripes in the attention matrix, as depicted in Figure 2. To focus attention on these stripes, we introduce the concept of *attention contribution*. This concept involves viewing the columns of the attention matrix as each point’s contribution to others within the same window. Let \mathbf{A}_{ij} denote the element in the i -th row and j -th column of the attention matrix \mathbf{A} . The value of \mathbf{A}_{ij} reflects the extent of point i ’s contribution to point j ; the larger \mathbf{A}_{ij} is, the more significant the contribution. We define the sub-adjacent attention contribution of each point as the aggregate of values within the pre-defined sub-adjacent span in the corresponding column, which is

$$\begin{aligned} \text{SACon}(\mathbf{A}) &= [\text{SACon}(\mathbf{A}_{:,i})]_{i=0, \dots, \text{win_size}-1} \\ \text{SACon}(\mathbf{A}_{:,i}) &= \sum_{\substack{|j-i| \leq K_2 \\ |j-i| \geq K_1}} \mathbf{A}_{ji}, 0 \leq j < \text{win_size} \end{aligned} \quad (1)$$

where subscript $(:, i)$ denotes the i th column of the corresponding matrix, and win_size is the window size.

The sub-adjacent attention contribution plays a pivotal role in two aspects. Firstly, it steers the focus of attention towards the sub-adjacent neighborhoods by being integrated into the loss function (Eq. 5). This is achieved by increasing $\text{SACon}(\mathbf{A})$ for all points during the training stage. Secondly, it assists in anomaly detection. This is due to its incorporation into the anomaly score calculation, where anomalies typically show a lower attention contribution than normal points.

Moreover, the number of the highlighted cells in Figure 2 is lower for the marginal points (where $i < K_2$ or $i > \text{win_size} - K_2$), leading to imbalance among points. It is mainly because the j in Eq. 1 is bounded by $0 \leq j < \text{win_size}$. To break through this limitation, one plausible way is to use a circular shift function to calculate \mathbf{A}_{ji} in Eq. 1:




Figure 3: Attention matrices obtained using (a) vanilla self-attention and (b) the linear attention with the proposed mapping function. The SMAP dataset [Hundman *et al.*, 2018b] and the proposed sub-adjacent neighborhoods are used.

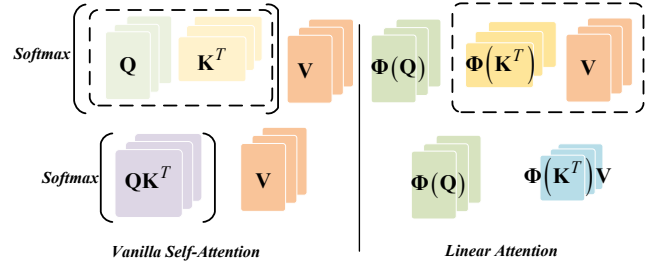


Figure 4: Illustration of vanilla self-attention and linear attention. Without the direct application of Softmax , the attention matrix $\Phi(\mathbf{Q})\Phi(\mathbf{K}^T)$ of linear attention usually exhibits more flexibility.

$$\mathbf{A}_{ji} = \langle [\text{Roll}(\mathbf{Q}, i-j)]_{i,:}, \mathbf{K}_{i,:} \rangle \quad (2)$$

where $\mathbf{Q}, \mathbf{K} \in \mathbb{R}^{\text{win_size} \times \text{d_model}}$ are the query and key matrix, respectively, $\langle \cdot, \cdot \rangle$ represents the inner product of two vectors, $\text{Roll}(\mathbf{Q}, i-j)$ cyclically shift matrix \mathbf{Q} by $i-j$ along the first dimension. The j values in \mathbf{A}_{ji} of Eq. 2 can satisfy the conditions $j < 0$ or $j \geq \text{win_size}$ and ensure that the number of j s for each i is the same. Actually, the cyclic operation in Eq. 2 is equivalent to setting $\text{SACon}(\mathbf{A}_{:,i})$ in Eq. 1 as:

$$\text{SACon}(\mathbf{A}_{:,i}) = \sum_{\substack{|j-i| \leq K_2 \\ |j-i| \geq K_1}} \mathbf{A}_{[j]i}, |j| < \text{win_size} \quad (3)$$

where $[j] = j + n \cdot \text{win_size}$, $n \in \{0, \pm 1\}$ such that $0 \leq [j] < \text{win_size}$. In implementation, we use Eq. 3 instead of Eq. 2 for efficiency. As shown in Figure 3(b), the use of Eq. 3 results in two extra side stripes in the attention matrix.

Linear Attention. Vanilla self-attention employs the Softmax function on a row-wise basis within the attention matrix, leading to competition among values in the same row, as depicted in Figure 3(a). In contrast, linear attention does not face such constraints. As shown in Figure 4, linear attention can be expressed as $\Phi(\mathbf{Q})\Phi(\mathbf{K}^T)\mathbf{V}$, where $\Phi(\cdot)$ is the mapping function. As shown in Figure 3, linear attention demonstrates better attention matrix shaping capabilities. Quantitative results can be found in Table 5.

The mapping function directly impacts the focusing capability and overall performance of linear attention. Herein we introduce a novel learnable mapping function:

$$\Phi(\cdot) = \text{Softmax}_{\text{row}}(\cdot/\tau) \quad (4)$$

where Softmax is applied row-wise to the input matrix, and τ is a learnable parameter to adjust dynamically the Softmax temperature. Note that in Eq. 4, the Softmax function is applied to the matrices \mathbf{Q} and \mathbf{K} , rather than directly to the attention matrix, as is the case with vanilla self-attention. This distinction is fundamental to the increased flexibility of the attention matrix. Moreover, we set all negative values in the matrices \mathbf{Q} and \mathbf{K} to a large negative number, such as -100 , to ensure that these values are close to 0 after applying the mapping function. Various mapping functions have been explored in prior research, including the power function [Han *et al.*, 2023], column-wise Softmax [Shen *et al.*, 2021], ReLU [Cai *et al.*, 2023], and elu function [Katharopoulos *et al.*, 2020]. Our empirical experiments demonstrate the effectiveness of our mapping function in time series anomaly detection, as detailed in Table 5.

3.3 Loss Function and Anomaly Score

Loss Function. Reconstruction loss is fundamental in unsupervised time series anomaly detection. As aforementioned, we also introduce the sub-adjacent attention contribution into the loss function, which guides the model to focus on the sub-adjacent neighborhoods. By integrating these two losses, the loss function for the input series $\mathbf{X} \in \mathbb{R}^{T \times D}$ is formulated as follows:

$$\begin{aligned} \mathcal{L}_{\text{Total}}(\mathbf{X}, \hat{\mathbf{X}}, \mathbf{A}) &= \mathcal{L}_{\text{rec}} + \lambda \cdot \mathcal{L}_{\text{attn}} \\ &= \|\mathbf{X} - \hat{\mathbf{X}}\|_F^2 - \lambda \cdot \|\text{SACon}(\mathbf{A})\|_1 \end{aligned} \quad (5)$$

where \mathcal{L}_{rec} is the reconstruction loss and $\mathcal{L}_{\text{attn}}$ is the attention loss, $\lambda > 0$ is the weight to trade off the two terms. $\|\cdot\|_F$ and $\|\cdot\|_1$ in Eq. 5 are the Frobenius and k-norm, respectively. $\hat{\mathbf{X}} \in \mathbb{R}^{T \times D}$ denotes the reconstructed \mathbf{X} .

Anomaly Score. To identify anomalies, we combine the sub-adjacent attention contribution with the reconstruction errors. Consistent with the approach in [Xu *et al.*, 2022], the Softmax function is applied to $-\text{SACon}(\mathbf{A})$ to highlight the anomalies with less attention contribution. Subsequently, we perform an element-wise multiplication of the attention results and reconstruction errors. Finally, the anomaly score for each point can be expressed as

$$\begin{aligned} \text{AnomalyScore}(\mathbf{X}) &= \text{Softmax}(-\text{SACon}(\mathbf{A})) \\ &\odot \left[\|\mathbf{X}_{:,i} - \hat{\mathbf{X}}_{:,i}\|_F^2 \right]_{i=1, \dots, T} \end{aligned} \quad (6)$$

where \odot is the element-wise multiplication. Typically, the anomalies exhibit less attention contribution and higher reconstruction errors, leading to larger anomaly scores.

Dynamic Gaussian Scoring. Following the practice of [Schmidl *et al.*, 2022], we fit a dynamic Gaussian distribution to anomaly scores obtained by Eq. 6 and design a score based on the fitted distribution. Let μ_t and σ_t denote the dynamic mean and standard variance, respectively, which can

Dataset	Dims	Entities	#Train (K)	#Test (K)	AR (%)
SWaT	51	1	495	449	12.14
WADI	123	1	1209	172	5.71
PSM	25	1	132	87	27.76
MSL	55	27	58	73	10.48
SMAP	25	55	140	444	12.83
SMD	38	28	708	708	4.16
NeurIPS-TS	1	1	20	20	22.44

Table 1: Datasets used in this study. AR is short for the anomaly rate. #Train and #Test denote the number of the training and test time points, respectively.

be computed in the way of sliding windows. Then the final score with dynamic Gaussian fitting can be computed via

$$\text{DyAnoSco}_t = -\log \left(1 - \text{cdf} \left(\frac{\text{AnomalyScore}_t - \mu_t}{\sigma_t} \right) \right) \quad (7)$$

where cdf is the cumulative distribution function of the standard Gaussian distribution $N(0, 1)$.

4 Experiments

4.1 Datasets

We evaluate the Sub-Adjacent Transformer on the following datasets, whose statistics are summarized in Table 1.

SWaT (Secure Water Treatment) [Mathur and Tippenhauer, 2016] is collected continuously over 11 days from 51 sensors located at a water treatment plant. **WADI** (Water DIstribution) [Ahmed *et al.*, 2017] is acquired from 123 sensors of a reduced water distribution system for 16 days. **PSM** (Pooled Server Metrics) [Abdulaal *et al.*, 2021] is collected from multiple servers at eBay with 26 dimensions. **MSL** (Mars Science Laboratory rover) and **SMAP** (Soil Moisture Active Passive satellite) are datasets released by NASA with 55 and 25 dimensions respectively. **SMD** (Server Machine Dataset) [Su *et al.*, 2019] is collected from a large compute cluster, consisting of 5 weeks of data from 28 server machines with 38 sensors. **NeurIPS-TS** (NeurIPS 2021 Time Series Benchmark) is a synthetic dataset proposed by [Lai *et al.*, 2021] and includes 5 kinds of anomalies that cover point- and pattern-wise behaviors: global (point), contextual (point), shapelet (pattern), seasonal (pattern), and trend (pattern).

4.2 Implementation Details

Following the common practice, we adopt a non-overlapping sliding window mechanism to obtain a series of sub-series. The sliding window size is set as 100 without particular statements. K_1 and K_2 are set as 20 and 30, respectively. The choices of these parameters will be discussed later in ablation studies. The points are judged to be anomalies if their anomaly scores (Eq. 7) are larger than a certain threshold δ . In this study, following the practice of paper [Lai *et al.*, 2023; Xu *et al.*, 2022], the thresholds are chosen to output the best F1 scores. The widely-used point adjustment strategy [Xu *et al.*, 2022; Tuli *et al.*, 2022; Garg *et al.*, 2022;

Category	Single-Entity						Multi-Entity†					
Datset	SWaT		WADI		PSM		MSL		SMAP		SMD	
Metric	AUC	F1	AUC	F1	AUC	F1	AUC	F1	AUC	F1	AUC	F1
DAGMM [2018]	91.6	85.3	29.2	20.9	85.2	76.1	80.6	70.1	81.5	71.2	82.3	72.3
LSTM-VAE [2018]	88.3	80.5	49.8	38.0	88.6	80.9	91.6	85.4	84.8	75.6	88.6	80.8
MSCRED [2019]	88.5	80.7	49.1	37.4	74.3	62.6	96.6	93.6	92.4	86.6	90.8	84.1
OmniAnomaly [2019]	92.4	86.6	53.9	41.7	77.6	66.4	94.6	90.1	91.6	85.4	98.0	96.2
MAD-GAN [2019]	89.0	81.5	67.9	55.6	77.1	65.8	95.5	91.7	92.3	86.5	95.4	91.5
MTAD-GAT [2020]	92.0	86.0	72.2	60.2	86.6	78.0	95.0	90.8	94.6	90.1	95.0	90.8
USAD [2020]	91.1	84.6	55.3	43.0	82.5	72.5	95.2	91.1	89.3	81.9	97.2	94.6
THOC [2020]	93.3	88.1	63.1	50.6	94.2	89.5	96.7	93.7	97.5	95.2	66.5	54.1
UAE [2022]	92.6	86.9	97.8	95.7	96.6	93.6	95.7	92.0	94.3	89.6	<u>98.6</u>	<u>97.2</u>
GDN [2021]	96.5	93.5	91.7	85.5	95.9	92.3	94.7	90.3	81.1	70.8	81.8	71.6
GTA [2021]	95.1	91.0	90.7	84.0	91.7	85.5	95.2	91.1	94.7	90.4	95.6	91.9
TranAD [2022]	89.0	81.5	62.0	49.5	93.4	88.2	97.3	94.9	94.0	89.2	98.0	96.1
Heuristics [2022]	98.4	96.9	98.2	96.5	<u>99.2</u>	<u>98.5</u>	<u>98.2</u>	<u>96.5</u>	98.0	96.1	96.5	93.4
Anomaly Transformer [2022]	96.9	94.1	98.3	96.6	98.9	97.9	96.6	93.6	98.3	96.7	95.9	92.3
NPSR [2023]	97.5	95.3	<u>96.7</u>	93.8	97.8	95.7	97.9	96.0	<u>98.9</u>	<u>97.8</u>	91.4	85.0
TimesNet [2023]	-	92.1	-	-	-	97.5	-	85.2	-	71.5	-	85.8
Ours	99.5	99.0	99.7	99.3	99.4	98.9	98.3	96.7	99.1	98.2	98.7	97.7

Table 2: Quantitative results for various anomaly detection methods in the six real-world datasets. AUC means area under the ROC curve. The largest and second-largest values are highlighted with bold text and underlined text, respectively. The values in this table are as % for ease of display. †: For multi-entity datasets, we use a single model to train and test all entities together, posing additional challenges.

Methods	DAGMM [2018]	LSTM-VAE [2018]	OmniAnomaly [2019]	THOC [2020]	Heuristics [2022]	Anomaly Transformer [2022]	NPSR [2023]	TimesNet [2023]	Ours
AUC	64.3	70.5	72.2	74.2	85.5	<u>86.2</u>	76.7	79.6	92.4
F1	51.8	58.4	60.2	62.5	76.3	<u>77.5</u>	65.4	69.5	86.6

Table 3: Anomaly detection performance in the synthetic dataset NeurIPS-TS. AUC means area under the ROC curve. The largest and second-largest values are highlighted with bold text and underlined text, respectively. The values in this table are presented in percentages.

Wu *et al.*, 2023] is adopted. Note that point adjustment has its practicality: one detected anomaly point will guide system administrators to identify the entire anomalous segment. The Sub-Adjacent Transformer contains 3 layers. Specifically, we set the hidden dimension d_{model} as 512, and the head number as 8. The hyper-parameter λ as 10 in Eq. 5 to balance recognition loss and attention loss. The Adam optimizer [Kingma and Ba, 2015] is used with an initial learning rate of 10^{-4} . Following [Xu *et al.*, 2022], the training is early stopped within 10 epochs with the batch size of 128. Experiments are conducted using PyTorch and one NVIDIA RTX A6000 GPU.

4.3 Main Results

Real-World Datasets. In Table 2, we report AUC and F1 metrics for various methods on six real-world datasets with 15 competitive baselines. The Sub-Adjacent Transformer consistently achieves state-of-the-art performance across all benchmarks, with an average improvement of 1.7 and 0.4 percentage points in F1 score over the current SOTA for single-entity and multi-entity datasets, respectively. Note that for

multi-entity datasets, we combine all entities and train them together using a single model, whereas some methods, such as simple heuristics [Garg *et al.*, 2022] and NPSR [Lai *et al.*, 2023], train each entity separately and average the results. Even though entity-to-entity variations are large (especially MSL and SMD datasets), the Sub-Adjacent Transformer still achieves better anomaly detection results using one model than the multi-model counterpart, showing the advantage of the proposed attention mechanism. We have retained all dimensions of the original datasets, even though some of them are constant and do not play a role. Additionally, we also report F1 results without using point-adjustment in the supplementary material, where the Sub-Adjacent Transformer also exhibits competitive performance.

Synthetic Dataset. NeurIPS-TS is a challenging synthetic dataset. We generate the NeurIPS-TS dataset using the source codes released by [Lai *et al.*, 2021], and it includes 5 anomalous types covering both the point-wise (global and contextual) and pattern-wise (shapelet, seasonal and trend) anomalies. As shown in Table 3, the Sub-Adjacent Trans-

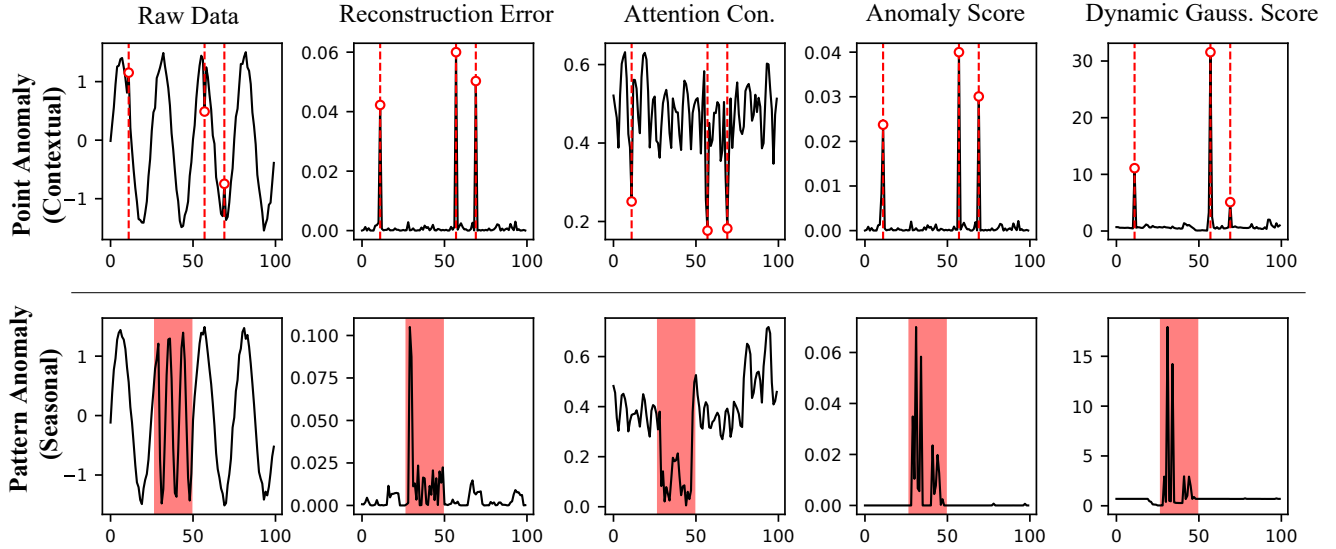


Figure 5: Visualization of detection results for different anomaly categories in NeurIPS-TS benchmark. The anomalous area are highlighted with red lines/areas. The first and second row represent point and pattern anomalies, respectively. From left to right, the columns indicate raw data, recognition error (Eq. 5), attention contribution (Eq. 1), anomaly score (Eq. 6) and dynamic Gaussian score (Eq. 7).

$K_1 : K_2$	SWaT	WADI	PSM	MSL	SMAP	SMD
0 : 0	94.8	93.5	95.3	91.5	92.4	93.8
0 : 5	96.7	95.9	97.6	93.6	93.7	94.1
0 : 10	<u>98.5</u>	<u>98.2</u>	<u>98.7</u>	94.4	96.9	95.9
10 : 20	<u>97.6</u>	<u>98.1</u>	98.1	<u>96.5</u>	<u>97.8</u>	94.4
20 : 30	99.0	99.3	98.9	<u>96.7</u>	98.2	<u>97.7</u>
30 : 40	<u>97.6</u>	94.8	97.8	95.6	97.5	94.0
10 : 30	<u>97.8</u>	<u>97.8</u>	<u>98.8</u>	95.0	97.4	94.5
20 : 40	<u>97.9</u>	<u>98.0</u>	<u>93.5</u>	95.7	95.6	94.4
10 : 40	<u>98.0</u>	<u>97.7</u>	<u>98.7</u>	96.1	97.5	97.8

Table 4: Performance comparison with different K_1 and K_2 settings for real-world datasets. The largest value for each dataset is emphasized in bold, while the second largest value is underlined. The values are marked in red if the value is not less than SOTA.

former achieves state-of-the-art performance with an astonishing 9.1 percentage points improvement of F1 score relative to Anomaly Transformer [Xu *et al.*, 2022], demonstrating the effectiveness of the proposed model on various anomalies. Visualization results can be found in Figure 5. We can see that anomalies do have less sub-adjacent contribution (the third column), which is consistent with our assumption. A full version of Figure 5 can be found in the supplementary material.

4.4 Ablation Studies

Vanilla and Linear Transformer. We compare vanilla self-attention and linear attention with different mapping functions using our method. As shown in Table 5, the proposed softmax-based mapping function with learnable pa-

rameters outperforms other mapping functions and vanilla self-attention, with a maximum F1 score increase of 1.5 percentage points.

Choices of K_1 and K_2 . Choices of K_1 and K_2 directly affect the performance of the Sub-Adjacent Transformer. We evaluate the effect in Table 4 with the window size being 100. The configuration of $K_1 = 20$ and $K_2 = 30$ typically yields the best performance (with the exception of the SMD dataset), while the performance of other settings is also satisfactory, with nearly half of the values exceeding the current SOTA. When the sub-adjacent span is getting close to the diagonal (the second and third rows in Table 4), the performance worsens. We also explore the extreme case where $K_1 = K_2 = 0$. In that case, the performance drops significantly, highlighting the importance of sub-adjacent attention.

Module Ablation. We evaluate the contribution of linear attention module and dynamic scoring module in Table 6. For the Sub-Adjacent Transformer, the baseline means the proposed sub-adjacent attention mechanism accompanied with vanilla self-attention and anomaly score (Eq. 6). Then linear attention and Gaussian dynamic scoring (Eq. 7) are introduced in turn. As shown in Table 6, compared with the baseline, linear attention brings +1.0 improvement on average over six datasets and Gaussian dynamic scoring brings another +0.2 improvement, verifying the effectiveness of the proposed modules.

Furthermore, we also apply the two modules to Anomaly Transformer [Xu *et al.*, 2022]. As we can see from Table 6, the proposed linear attention model provides +1.2 percentage points performance gain averagely, and Gaussian dynamic scoring provides another +0.1 gain. Moreover, as shown in Table 6, without the support of linear attention and dynamic Gaussian score, our method (baseline) still outper-

Dataset		SWaT	WADI	PSM	MSL	SMAP	SMD
Lin. Attn.	Vanilla self-attention [Vaswani <i>et al.</i> , 2017]	98.6	98.2	<u>98.8</u>	94.1	97.4	<u>96.5</u>
	Softmax_column [Shen <i>et al.</i> , 2021]	98.0	97.5	96.9	<u>95.2</u>	<u>97.8</u>	95.8
	Power function [Han <i>et al.</i> , 2023]	98.4	<u>98.3</u>	98.4	93.7	97.5	95.9
	ReLU [Cai <i>et al.</i> , 2023]	<u>98.8</u>	98.1	97.6	94.3	97.5	94.8
	ELU+1 [Katharopoulos <i>et al.</i> , 2020]	95.4	95.6	94.1	94.7	96.3	94.5
	Ours (Eq. 4)	99.0	99.3	98.9	96.7	98.2	97.7

Table 5: F1 values of the proposed method with vanilla self-attention and linear attention with various mapping functions. Bold and underlined text denote the largest and second-largest values per dataset. We mark the value in red if it is not less than SOTA.

	Baseline	Linear Attention	Dynamic Scoring	SWaT	WADI	PSM	MSL	SMAP	SMD
Ours	✓			98.2	98.0	98.7	93.9	97.0	96.1
	✓	✓		98.8 _(+0.6)	99.1 _(+1.1)	98.8 _(+0.1)	96.4 _(+2.5)	97.9 _(+0.9)	97.0 _(+0.9)
	✓	✓	✓	99.0_(+0.8)	99.3_(+1.3)	98.9_(+0.2)	96.7_(+2.8)	98.2_(+1.2)	97.7_(+1.6)
Anomaly Transformer [2022]	✓			94.1	96.6	97.9	93.6	96.7	92.3
Transformer [2022]	✓	✓		98.1 _(+4.0)	97.6 _(+1.0)	98.3_(+0.4)	93.6 _(+0.0)	97.6 _(+0.9)	93.4 _(+1.1)
	✓	✓	✓	98.3_(+4.2)	97.9_(+1.3)	98.2 _(+0.3)	93.7_(+0.1)	97.7_(+1.0)	93.6_(+1.3)

Table 6: Ablation on modules for the proposed Sub-Adjacent Transformer and Anomaly Transformer. The largest values for each dataset are highlighted with bold text. The gain with respect to baseline is given in the subscript.

W_S	SWaT	WADI	PSM	MSL	SMAP	SMD
50	97.6	97.8	98.4	94.7	96.9	92.1
100	<u>99.0</u>	99.3	98.9	96.7	98.2	97.7
150	<u>98.8</u>	98.8	98.7	94.9	97.3	96.4
200	99.1	<u>98.7</u>	<u>98.5</u>	94.9	<u>97.5</u>	<u>94.2</u>
250	97.2	98.1	98.3	<u>95.3</u>	<u>97.4</u>	93.9
300	98.9	98.2	98.1	94.3	96.9	94.3

Table 7: F1 values for different window sizes. The largest and second-largest values are made bold and underlined, respectively.

forms the Anomaly Transformer by an average of +1.8 percentage points, thereby validating the efficacy of our sub-adjacent attention design.

Choices of Window Size. Table 7 gives F1 values with different window sizes for six datasets. One can see that our model exhibits good robustness to the window size. The optimal performance is typically achieved with a window size of 100, the only exception being the SWaT dataset.

Choices of Parameter λ . Table 8 compares the performance for different values of λ . The best results are achieved when λ is set to 10, while the second-best results occur at different values of λ for different datasets. It is noteworthy that without the inclusion of attention loss ($\lambda = 0$), performance drops drastically, affirming the effectiveness of our sub-adjacent attention design.

λ	SWaT	WADI	PSM	MSL	SMAP	SMD
0	95.8	94.2	95.6	92.7	95.3	94.8
4	<u>98.8</u>	97.8	<u>98.8</u>	94.7	97.6	96.7
6	<u>98.6</u>	<u>98.3</u>	<u>98.6</u>	<u>96.1</u>	97.4	<u>97.3</u>
8	98.3	<u>97.8</u>	98.7	<u>95.8</u>	98.0	<u>96.9</u>
10	99.0	99.3	98.9	96.7	98.2	97.7
12	98.5	<u>98.3</u>	98.6	95.8	97.6	97.1
14	98.7	98.2	98.6	95.2	97.3	96.3

Table 8: F1 values for different λ settings. Bold and underlined texts denote the largest and second largest values, respectively.

5 Conclusion

In this paper, we present the Sub-Adjacent Transformer, a novel paradigm for utilizing attention in time series anomaly detection. Our method distinctively combines sub-adjacent attention contribution, linear attention and reconstruction error to effectively detect anomalies, thereby enhancing the efficacy of anomaly detection. It offers a novel perspective on the utilization of attention in this domain. Without bells and whistles, our model demonstrates superior performance across common benchmarks. We hope that the Sub-Adjacent Transformer could act as a baseline framework for the future works in time series anomaly detection.

Acknowledgments

This work was supported by the National Science Foundation of China (NSFC) (No. 62371009 and No. 61971008).

References

- [Abdulaal *et al.*, 2021] Ahmed Abdulaal, Zhuanghua Liu, and Tomer Lancewicki. Practical approach to asynchronous multivariate time series anomaly detection and localization. In *SIGKDD*, pages 2485–2494, 2021.
- [Ahmed *et al.*, 2017] Chuadhry Mujeeb Ahmed, Venkata Reddy Palleti, and Aditya P. Mathur. Wadi: a water distribution testbed for research in the design of secure cyber physical systems. In *Proceedings of the 3rd international workshop on cyber-physical systems for smart water networks*, pages 25–28, 2017.
- [Audibert *et al.*, 2020] Julien Audibert, Pietro Michiardi, Frédéric Guyard, Sébastien Marti, and Maria A. Zuluaga. Usad: Unsupervised anomaly detection on multivariate time series. In *SIGKDD*, pages 3395–3404, 2020.
- [Blázquez-García *et al.*, 2021] Ane Blázquez-García, Angel Conde, Usue Mori, and Jose A Lozano. A review on outlier/anomaly detection in time series data. *ACM Computing Surveys*, 54(3):1–33, April 2021.
- [Bolya *et al.*, 2022] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, and Judy Hoffman. Hydra attention: Efficient attention with many heads. In *ECCV*, pages 35–49. Springer, 2022.
- [Boniol and Palpanas, 2022] Paul Boniol and Themis Palpanas. Series2graph: Graph-based subsequence anomaly detection for time series. *arXiv preprint arXiv:2207.12208*, 2022.
- [Breunig *et al.*, 2000] Markus M Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Jörg Sander. Lof: identifying density-based local outliers. In *SIGMOD*, pages 93–104, 2000.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NIPS*, 33:1877–1901, 2020.
- [Cai *et al.*, 2023] Han Cai, Chuang Gan, and Song Han. Efficientvit: Lightweight multi-scale attention for high-resolution dense prediction. In *ICCV*, pages 17302–17313, 2023.
- [Chen *et al.*, 2021] Zekai Chen, Dingshuo Chen, Xiao Zhang, Zixuan Yuan, and Xiuzhen Cheng. Learning graph structures with transformer for multivariate time-series anomaly detection in iot. *IEEE Internet of Things Journal*, 9(12):9179–9189, 2021.
- [Chen, 2018] Tiankai Chen. *Anomaly detection in semiconductor manufacturing through time series forecasting using neural networks*. PhD thesis, Massachusetts Institute of Technology, 2018.
- [Cheng *et al.*, 2008] Haibin Cheng, Pang-Ning Tan, Christopher Potter, and Steven Klooster. A robust graph-based algorithm for detection and characterization of anomalies in noisy multivariate time series. In *ICDM workshops*, pages 349–358. IEEE, 2008.
- [Deng and Hooi, 2021] Ailin Deng and Bryan Hooi. Graph neural network-based anomaly detection in multivariate time series. In *AAAI*, volume 35, pages 4027–4035, 2021.
- [Ding *et al.*, 2019] Nan Ding, HaoXuan Ma, Huanbo Gao, YanHua Ma, and GuoZhen Tan. Real-time anomaly detection based on long short-term memory and gaussian mixture model. *Computers & Electrical Engineering*, 79:106458, 2019.
- [Fox, 1972] Anthony J Fox. Outliers in time series. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 34(3):350–363, 1972.
- [Garg *et al.*, 2022] Astha Garg, Wenyu Zhang, Jules Samaran, Ramasamy Savitha, and Chuan-Sheng Foo. An evaluation of anomaly detection and diagnosis in multivariate time series. *IEEE Transactions on Neural Networks and Learning Systems*, 33(6):2508–2517, 2022.
- [Gupta *et al.*, 2013] Manish Gupta, Jing Gao, Charu C. Aggarwal, and Jiawei Han. Outlier detection for temporal data: A survey. *IEEE Transactions on Knowledge and data Engineering*, 26(9):2250–2267, 2013.
- [Han *et al.*, 2023] Dongchen Han, Xuran Pan, Yizeng Han, Shiji Song, and Gao Huang. Flatten transformer: Vision transformer using focused linear attention. In *ICCV*, pages 5961–5971, 2023.
- [Ho *et al.*, 2020] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NIPS*, 33:6840–6851, 2020.
- [Hundman *et al.*, 2018a] Kyle Hundman, Valentino Constantinou, Christopher Laporte, Ian Colwell, and Tom Soderstrom. Detecting spacecraft anomalies using lstms and nonparametric dynamic thresholding. In *SIGKDD*, pages 387–395, 2018.
- [Hundman *et al.*, 2018b] Kyle Hundman, Valentino Constantinou, Christopher Laporte, Ian Colwell, and Tom Soderstrom. Detecting spacecraft anomalies using lstms and nonparametric dynamic thresholding. In *SIGKDD*, pages 387–395, 2018.
- [Katharopoulos *et al.*, 2020] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *ICML*, pages 5156–5165. PMLR, 2020.
- [Kieu *et al.*, 2019] Tung Kieu, Bin Yang, Chenjuan Guo, and Christian S. Jensen. Outlier detection for time series with recurrent autoencoder ensembles. In *IJCAI*, pages 2725–2732, 2019.
- [Kingma and Ba, 2015] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [Lai *et al.*, 2021] Kwei-Herng Lai, Daochen Zha, Junjie Xu, Yue Zhao, Guanchu Wang, and Xia Hu. Revisiting time series outlier detection: Definitions and benchmarks. In *NeurIPS Dataset and Benchmark Track*, 2021.

- [Lai *et al.*, 2023] Chih-Yu Lai, Fan-Keng Sun, Zhengqi Gao, Jeffrey H Lang, and Duane S Boning. Nominality score conditioned time series anomaly detection by point/sequential reconstruction. *arXiv preprint arXiv:2310.15416*, 2023.
- [Li *et al.*, 2019] Dan Li, Dacheng Chen, Baihong Jin, Lei Shi, Jonathan Goh, and See-Kiong Ng. Mad-gan: Multivariate anomaly detection for time series data with generative adversarial networks. In *International conference on artificial neural networks*, pages 703–716. Springer, 2019.
- [Li *et al.*, 2021] Zhihan Li, Youjian Zhao, Jiaqi Han, Ya Su, Rui Jiao, Xidao Wen, and Dan Pei. Multivariate time series anomaly detection and interpretation using hierarchical inter-metric and temporal embedding. In *SIGKDD*, pages 3220–3230, 2021.
- [Liu *et al.*, 2013] Bo Liu, Yanshan Xiao, Longbing Cao, Zhifeng Hao, and Feiqi Deng. Svdd-based outlier detection on uncertain data. *Knowledge and information systems*, 34:597–618, 2013.
- [Liu *et al.*, 2021] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Ching-Feng Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021.
- [Mathur and Tippenhauer, 2016] Aditya P. Mathur and Nils Ole Tippenhauer. Swat: a water treatment testbed for research and training on ics security. In *CySWATER*, pages 31–36. IEEE, 2016.
- [Park *et al.*, 2018] Daehyung Park, Yuuna Hoshi, and Charles C Kemp. A multimodal anomaly detector for robot-assisted feeding using an lstm-based variational autoencoder. *IEEE Robotics and Automation Letters*, 3(3):1544–1551, 2018.
- [Pereira and Silveira, 2019] João Pereira and Margarida Silveira. Learning representations from healthcare time series data for unsupervised anomaly detection. In *2019 IEEE international conference on big data and smart computing (BigComp)*, pages 1–7. IEEE, 2019.
- [Pintilie *et al.*, 2023] Ioana Pintilie, Andrei Manolache, and Florin Brad. Time series anomaly detection using diffusion-based models. *arXiv preprint arXiv:2311.01452*, 2023.
- [Schmidl *et al.*, 2022] Sebastian Schmidl, Phillip Wenig, and Thorsten Papenbrock. Anomaly detection in time series: a comprehensive evaluation. *Proceedings of the VLDB Endowment*, 15(9):1779–1797, 2022.
- [Schölkopf *et al.*, 2001] Bernhard Schölkopf, John C. Platt, J. Shawe-Taylor, Alex Smola, and R. C. Williamson. Estimating the support of a high-dimensional distribution. *Neural computation*, 13(7):1443–1471, 2001.
- [Shen *et al.*, 2020] Lifeng Shen, Zhuocong Li, and James T. Kwok. Timeseries anomaly detection using temporal hierarchical one-class network. *NIPS*, 33:13016–13026, 2020.
- [Shen *et al.*, 2021] Zhuoran Shen, Mingyuan Zhang, Haiyu Zhao, Shuai Yi, , and Hongsheng Li. Efficient attention: Attention with linear complexities. In *WACV*, pages 3531–3539, 2021.
- [Su *et al.*, 2019] Ya Su, Youjian Zhao, Chenhao Niu, Rong Liu, Wei Sun, and Dan Pei. Robust anomaly detection for multivariate time series through stochastic recurrent neural network. In *SIGKDD*, pages 2828–2837, 2019.
- [Sun *et al.*, 2021] Fan-Keng Sun, Chris Lang, and Duane Boning. Adjusting for autocorrelated errors in neural networks for time series. *NIPS*, 34:29806–29819, 2021.
- [Tax and Duin, 2004] David M.J. Tax and Robert P.W. Duin. Support vector data description. *Machine learning*, 54:45–66, 2004.
- [Tuli *et al.*, 2022] Shreshth Tuli, Giuliano Casale, and Nicholas R Jennings. Tranad: Deep transformer networks for anomaly detection in multivariate time series data. *arXiv preprint arXiv:2201.07284*, 2022.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NIPS*, 30, 2017.
- [Wu *et al.*, 2023] Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. Timesnet: Temporal 2d-variation modeling for general time series analysis. *ICLR*, 2023.
- [Xu *et al.*, 2022] Jiehui Xu, Haixu Wu, Jianmin Wang, and Mingsheng Long. Anomaly transformer: Time series anomaly detection with association discrepancy. In *ICLR*, 2022.
- [Zhang *et al.*, 2019] Chuxu Zhang, Dongjin Song, Yuncong Chen, Xinyang Feng, Cristian Lumezanu, Wei Cheng, Jingchao Ni, Bo Zong, Haifeng Chen, and Nitesh V Chawla. A deep neural network for unsupervised anomaly detection and diagnosis in multivariate time series data. In *AAAI*, volume 33, pages 1409–1416, 2019.
- [Zhao *et al.*, 2020] Hang Zhao, Yujing Wang, Juanyong Duan, Congrui Huang, Defu Cao, Yunhai Tong, Bixiong Xu, Jing Bai, Jie Tong, and Qi Zhang. Multivariate time-series anomaly detection via graph attention network. In *ICDM*, pages 841–850. IEEE, 2020.
- [Zhao *et al.*, 2022] Yan Zhao, Liwei Deng, Xuanhao Chen, Chenjuan Guo, Bin Yang, Tung Kieu, Feiteng Huang, Torben Bach Pedersen, Kai Zheng, and Christian S Jensen. A comparative study on unsupervised anomaly detection for time series: Experiments and analysis. *arXiv preprint arXiv:2209.04635*, 2022.
- [Zhou *et al.*, 2019] Bin Zhou, Shenghua Liu, Bryan Hooi, Xueqi Cheng, and Jing Ye. Beatgan: Anomalous rhythm detection using adversarially generated time series. In *IJCAI*, volume 2019, pages 4433–4439, 2019.
- [Zong *et al.*, 2018] Bo Zong, Qi Song, Martin Renqiang Min, Wei Cheng, et al. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In *ICLR*, 2018.