

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 LOW RANK TRANSFORMER FOR MULTIVARIATE TIME SERIES ANOMALY DETECTION AND LOCALIZATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Multivariate time series (MTS) anomaly diagnosis, which encompasses both anomaly detection and localization, is critical for the safety and reliability of complex, large-scale real-world systems. The vast majority of existing anomaly diagnosis methods offer limited theoretical insights, especially for anomaly localization, which is a vital but largely unexplored area. The aim of this contribution is to study the learning process of a Transformer when applied to MTS by revealing connections to statistical time series methods. Based on these theoretical insights, we propose the Attention Low-Rank Transformer (ALoRa-T) model, which applies low-rank regularization to self-attention, and we introduce the Attention Low-Rank score, effectively capturing the temporal characteristics of anomalies. Finally, to enable anomaly localization, we propose the ALoRa-Loc method, a novel approach that associates anomalies to specific variables by quantifying interrelationships among time series. Extensive experiments and real data analysis, show that the proposed methodology significantly outperforms state-of-the-art methods in both detection and localization tasks.

1 INTRODUCTION

Driven by the rapid growth of the Internet of Things (IoT), real-world systems have become increasingly more complex and vulnerable to faults. These anomalies frequently result in abnormal patterns for a stream of MTS. In this context, the diagnosis of anomalies in MTS is of great importance to ensure the reliability, safety, and efficiency of critical systems. Due to the scarcity of labeled data, anomaly diagnosis is commonly formulated as an unsupervised learning problem. It typically involves two key tasks: anomaly detection, which determines which timestamps are anomalous, and anomaly localization, which identifies the specific time series responsible for the detected anomalies.

MTS data exhibit complex dynamics, including temporal dependencies (relationships over time) and spatial dependencies (relationships across series). Effective anomaly diagnosis depends on reliably estimating these spatio-temporal dynamics. Deep learning models are widely applied to this task for their strong representation learning capabilities, with Transformer-based architectures shown to be especially effective in modeling the complex dynamics of MTS (Zerveas et al., 2021). However, the limited available theoretical insights into their decision process undermine reliability and trust in safety-critical settings, while also complicating anomaly localization. Practitioners often raise important questions regarding the learning process, the interrelationships within the data learned by the model, and the need for a localization method, since detection alone provides limited practical value in complex, large-scale systems. As illustrated in Fig. 1, two of the main contributions of this work are to address these open questions, which have remained largely unanswered.

In addition, several existing detection methods are built around the point-adjustment evaluation strategy, whereby detecting any point within an anomalous segment suffices to consider the entire segment as an observed anomaly (Xu et al., 2022; Song et al., 2023; Shen et al., 2020). This approach often inflates performance metrics, since even random scoring methods can appear competitive under such evaluation (Kim et al., 2021; Huet et al., 2022). To ensure reliable evaluation, trustworthy detection metrics have to accurately reflect the temporal characteristics of anomalies.

To address the challenges of reliable detection scoring, interpretability, and effective localization, this study makes the following contributions:

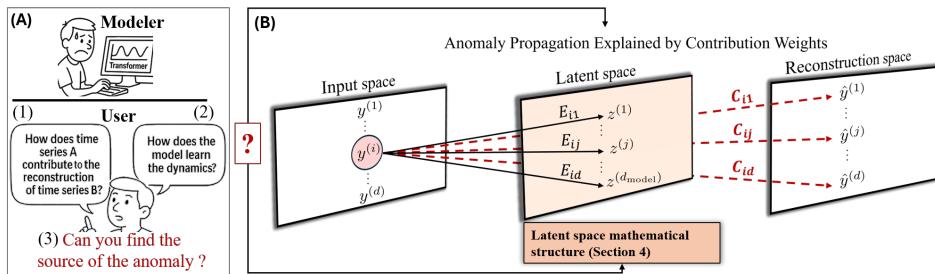


Figure 1: (A): Unresolved key questions commonly raised by practitioners. (B): The proposed localization method (Sec. 5.2), and the theoretical analysis (Sec. 4), addresses these challenges. The Figure illustrates anomaly propagation in MTS data. An anomaly originating in an input series $i \in \{1, \dots, d\}$ can propagate to the latent space and subsequently to the reconstructed outputs. This propagation is quantified by contribution weights: $E_{ij} \in \mathbb{R}$, capturing the influence of input series i on latent feature $j \in \{1, \dots, d_{\text{model}}\}$, and $C_{ij} \in \mathbb{R}$, capturing its influence on the reconstructed output series $\{j \in 1, \dots, d\}$. Together, these weights enable tracing of anomalies across different stages of the model, thereby supporting effective localization and interpretability.

- **Theoretical insights of Transformer encoders on MTS:** We study how Transformer encoders learn from MTS data, shedding light on how the learned representations relate to classical time series. These insights advance the design of anomaly diagnosis methods and deepen the understanding of Transformers in sequential modeling. (Section 4)
- **Attention Low-Rank Transformer (ALoRa-T) for anomaly detection:** We propose the ALoRa-T method, which enhances anomaly detection by applying low-rank regularization to self-attention, and introduce a novel detection metric, the Self-Attention Low-Rank score (ALoRa-T score). This method outperforms state-of-the-art methods across multiple benchmark datasets, and captures effectively the temporal characteristics of anomalies. (Section 5.1)
- **ALoRa-Loc method for anomaly localization:** We introduce ALoRa-Loc, a localization method that first derives contribution weights that quantify the learned interrelationships between time series, offering key insights into the model’s decision-making process. More critically, these weights capture how anomalies propagate across time series during modeling, enabling ALoRa-Loc to trace anomalies back to their origin and attribute them to the most relevant input time series. (Section 5.2)

2 RELATED WORK

Anomaly detection: Due to its importance for the safety and reliability of many critical infrastructure systems, such as water distribution networks (Eliades et al., 2025), power systems (Kyriakides & Polycarpou, 2015) and healthcare (Tang et al., 2022), unsupervised anomaly detection in MTS has attracted significant attention. Traditional unsupervised methods (Liu et al., 2008; Breunig et al., 2000; Schölkopf et al., 2001) have been applied to this task, but may not perform well since they do not learn complex dependencies. To address these limitations, deep learning models that capture temporal and inter-variable relationships have shown improvements over existing methods. A major class of deep learning approaches is reconstruction-based methods, typically using an encoder-decoder architecture trained in a self-supervised manner. The main idea is to learn meaningful latent representations that mimic the normal behavior. The models are expected to reconstruct normal sequences accurately. An unsuccessful reconstruction, in the sense of inflated errors, indicates potential anomalies in the data. Representative examples include RNN-based models such as LSTM-VAE (Park et al., 2018), OmniAnomaly (Su et al., 2019), and InterFusion (Li et al., 2021), as well as the CNN-based model MSCRED (Zhang et al., 2019). More recently, Transformer models (Vaswani et al., 2017) have been proven highly effective for sequential modeling and well-suited to the complex dynamics of MTS (Nie et al., 2023). Anomaly-Transformer (Xu et al., 2022) introduces the AssDis metric, which compares each row of the self-attention matrix to a prior Gaussian distribution. MEMTO (Song et al., 2023) enhances a Transformer encoder with a gated memory module that

108 learns memory items representing normal behavior. As new data arrives, it selectively updates these
 109 items. For anomaly detection, they introduce the LSD score, the distance between input queries and
 110 their closest memory items. Although the aforementioned models have shown strong performance
 111 under point-adjustment metrics, they may often not produce meaningful and reliable anomaly scores,
 112 as illustrated later in Figure 4. Lai et al. (2023) use the Performer model (Choromanski et al., 2021)
 113 to perform both point-wise and segment-wise reconstruction, defining the anomaly score as the ratio
 114 of their reconstruction errors. SARAD (Dai et al., 2024) uses a Transformer on transposed input
 115 windows, and thus the attention weights express feature-wise associations. To guide training and
 116 detect anomalies, SARAD relies on SAR metric, which tracks changes in attention weights, based
 117 on the intuition that anomalies disrupt feature associations.

118 **Anomaly Localization:** Anomaly localization has remained to a large extent an unexplored re-
 119 search area in deep learning. Recently it begun to attract more attention, with some of the afore-
 120 mentioned approaches also addressing this critical issue. Localization effectiveness depends on two
 121 key factors: first, the model’s ability to reconstruct normal patterns accurately. Second, the ability
 122 to interpret the reconstruction errors in a way that isolates the true sources of anomalies without
 123 being distorted by their effects. Some methods rely directly on the reconstruction error to inform
 124 the localization. For example, SARAD and OmniAnomaly (Su et al., 2019) perform localization by
 125 ranking the reconstruction errors of individual time series. While this provides a baseline approach,
 126 anomaly propagation can obscure true anomaly sources. More advanced techniques refine the recon-
 127 struction signal to improve localization. InterFusion (Li et al., 2021) applies Markov Chain Monte
 128 Carlo to adjust reconstruction errors. However, due to its computational complexity, it is limited
 129 to interpreting anomaly segments instead of individual time steps. DAEMON (Chen et al., 2023)
 130 leverages integrated gradients (Sundararajan et al., 2017) to attribute anomalies back to individual
 131 dimensions. Finally, AERCA (Han et al., 2025) focuses exclusively on the localization task, without
 132 addressing detection and localization simultaneously, while providing an interpretable approach and
 133 representing one of the most recent and effective methods in this area.

134 **Research gaps:** Although the aforementioned methods propose innovative architectures and scoring
 135 metrics, they often overlook key aspects: understanding what the model learns, how each time series
 136 is represented, and how these representations contribute to the final decisions. Such understanding
 137 is crucial for making detection applicable and trustworthy, and for enabling effective localization.

139 3 MULTIVARIATE TIME SERIES (MTS) NOTATION

140
 141 A Multivariate Time Series (MTS) is a sequence of observations recorded over time across multiple
 142 variables. It is denoted as $\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N]^\top$, where each row vector $\mathbf{y}_t \in \mathbb{R}^d$, $t = 1, \dots, N$,
 143 denotes the multivariate observation at time step t . Each column vector of \mathbf{Y} , denoted as $\mathbf{y}^{(i)} \in \mathbb{R}^N$
 144 represents the i -th univariate time series of length N , and $y_t^{(i)} \in \mathbb{R}$ is the value of the i -th series
 145 at time t . Given the streaming nature of time-series data, modeling the MTS at an arbitrary time
 146 t relies on its history $\mathbf{Y}_{(t-T:t]}$, defined as the most recent T observations up to t : $\mathbf{Y}_{(t-T:t)} =$
 147 $[\mathbf{y}_{t-T+1}, \dots, \mathbf{y}_t]^\top \in \mathbb{R}^{T \times d}$. Applying this construction to the full MTS $\mathbf{Y} \in \mathbb{R}^{N \times d}$ produces a
 148 sequence of overlapping windows, each serving as the model input. Since the subsequent sections
 149 concern an arbitrary time t , for simplicity we denote the window $\mathbf{Y}_{(t-T:t)}$ by $\mathbf{Y}_{[t]}$.
 150

152 4 THEORETICAL ANALYSIS OF TRANSFORMER ENCODER

153 This section presents a theoretical analysis of the learning dynamics of the Transformer encoder
 154 in the context of MTS. In particular, it relates the components of a standard Transformer model to
 155 classical statistical models.

156 **Embedding:** Given a MTS window at time t , denoted as $\mathbf{Y}_{[t]} \in \mathbb{R}^{T \times d}$, a Transformer encoder
 157 typically begins by embedding the input into a higher-dimensional space using a 1D-convolutional
 158 layer, i.e it maps $\mathbf{Y}_{[t]}$ to $\tilde{\mathbf{Y}}_{[t]} \in \mathbb{R}^{T \times d_{\text{model}}}$, where d_{model} is the embedding (model) dimension. Ap-
 159 pllying the 1D-convolution, the k -th embedded time series value at time step t , $\tilde{y}_t^{(k)}$, is computed

162 as:

$$\tilde{y}_t^{(k)} = \sum_{i=1}^d \underbrace{\left(\sum_{j=-\frac{m-1}{2}}^{\frac{m-1}{2}} w_{i,j}^{(k)} \cdot y_{t+j}^{(i)} \right)}_{\text{Weighted average of the } i\text{'th raw time series}} \quad (1)$$

168 This operation is mathematically equivalent to a learnable **Vector Moving Average (VMA) filtering**
 169 (Brockwell & Davis, 2002), where the output at time t is a weighted sum of multiple time series
 170 values in a local window around t . The filter weights $w_{i,j}^{(k)}$ determine the influence of time series i
 171 on the output feature k , at lag j . The constant m is the kernel size.

172 **Self-Attention Latent Space:** The output of the embedding layer, $\tilde{\mathbf{Y}}_{[t]} \in \mathbb{R}^{T \times d_{\text{model}}}$, is passed as
 173 input to the attention mechanism. At each layer l , the self-attention mechanism projects the previous
 174 layer's output $\mathbf{Z}^{(l-1)} \in \mathbb{R}^{T \times d_{\text{model}}}$ into the Query (\mathbf{Q}^l), Key (\mathbf{K}^l), and Value (\mathbf{V}^l) matrices using
 175 learnable projection matrices $\mathbf{W}^{(Q,l)} = \{w_{ij}^{(Q,l)}\}$, $\mathbf{W}^{(K,l)} = \{w_{ij}^{(K,l)}\}$, and $\mathbf{W}^{(V,l)} = \{w_{ij}^{(V,l)}\}$,
 176 each of dimension $\mathbb{R}^{d_{\text{model}} \times d_{\text{model}}}$. These projections are computed as:
 177

$$\mathbf{Q}^l = \mathbf{Z}^{(l-1)} \mathbf{W}^{(Q,l)}, \quad \mathbf{K}^l = \mathbf{Z}^{(l-1)} \mathbf{W}^{(K,l)}, \quad \mathbf{V}^l = \mathbf{Z}^{(l-1)} \mathbf{W}^{(V,l)}.$$

180 The self-attention scores at layer l are computed as the matrix $\mathbf{S}^{(l)} \in \mathbb{R}^{T \times T}$:

$$\mathbf{S}^{(l)} = \text{softmax} \left(\frac{\mathbf{Q}^l (\mathbf{K}^l)^\top}{\sqrt{d_{\text{model}}}} + \mathbf{M} \right), \quad (2)$$

184 where $\mathbf{M} \in \mathbb{R}^{T \times T}$ is an optional masking matrix, and the softmax is applied row-wise. The latent
 185 representation at layer l is updated through the residual attention mechanism, defined as:

$$\mathbf{Z}^{(l)} = \tilde{\mathbf{Z}}^{(l)} + \mathbf{Z}^{(l-1)}, \quad \text{where } \tilde{\mathbf{Z}}^{(l)} = \mathbf{S}^{(l)} \mathbf{Z}^{(l-1)} \mathbf{W}^{(V,l)}. \quad (3)$$

188 If skip connections are omitted, the update simplifies to $\mathbf{Z}^{(l)} = \tilde{\mathbf{Z}}^{(l)}$. By unrolling Eq. (3), we
 189 can derive formulations that relate the final latent representation, to useful statistical models. In the
 190 absence of skip connections, the latent representation at time step t can be written as:

$$\mathbf{Z}_t = \mathbf{A}_t \tilde{\mathbf{Y}}_{[t]} \mathbf{B}, \quad (4)$$

193 where $\mathbf{A}_t = \mathbf{S}_{t,:}^{(L)} \mathbf{S}_{t,:}^{(L-1)} \dots \mathbf{S}_{t,:}^{(1)} \in \mathbb{R}^{1 \times T}$, $\mathbf{B} = \mathbf{W}^{(V,1)} \dots \mathbf{W}^{(V,L)} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{model}}}$ and $\mathbf{S}_{t,:}^{(L)} \in$
 194 $\mathbb{R}^{1 \times T}$ denotes the t -th row of the final self-attention (SA) matrix $\mathbf{S}^{(L)}$. When skip connections are
 195 included (the standard case), the final representation becomes
 196

$$\mathbf{Z}^{(L)} = \tilde{\mathbf{Y}}_{[t]} + \sum_{\emptyset \neq I \subseteq \{1, \dots, L\}} \left(\prod_{i \in I^\downarrow} \mathbf{S}^{(i)} \right) \tilde{\mathbf{Y}}_{[t]} \left(\prod_{i \in I^\uparrow} \mathbf{W}^{(V,i)} \right), \quad (5)$$

200 where I^\downarrow and I^\uparrow denote the indices in descending and ascending order, respectively. Based on
 201 Eqs (4) and (5), the transformation of the input data in the Transformer is expressed as a left multi-
 202 pllication of the self-attention matrices across layers and a right multiplication of the Value projection
 203 matrices. The former is input-dependent, while the latter are learned parameters independent of the
 204 data. These two equations form the basis for the following proposition, which defines the struc-
 205 ture of the self-attention latent space. An analytical proof is provided in Appendix G, along with
 206 additional details on the derived statements.

208 **Proposition 1 (Space-Time Autoregressive (STAR) structure of the self-attention latent space)**

210 **1. Without skip connections:** Each time series in the Transformer's latent space follows a STAR-like
 211 structure. In particular, it can be expressed as $z_t^{(j)} = \sum_{k=1}^{d_{\text{model}}} b_{kj} \left(\sum_{q=1}^t a_{tq} \tilde{y}_q^{(k)} \right)$, which has the
 212 exact same form as the classical STAR model (Cressie & Wikle, 2011). The key distinction is that,
 213 while traditional STAR models use fixed lag weights a_{tq} estimated by minimizing a loss function
 214 (e.g., mean squared error), the Transformer computes these weights dynamically from the input
 215 through the attention mechanism, with Q and K guiding their estimation in real time.

216 2. **With skip connections:** Each time series in the final representation, Eq. (5), can be interpreted
 217 as a **linear combination of multiple STAR-like processes**, where each component captures distinct
 218 temporal and feature-level dependencies of the original MTS.
 219

220 3. **With feed-forward layers:** Adding feed-forward layers does not alter the STAR-like structure of the
 221 latent space. The only difference is that the spatial weights take a more complex form, as explained
 222 analytically in Proof G.

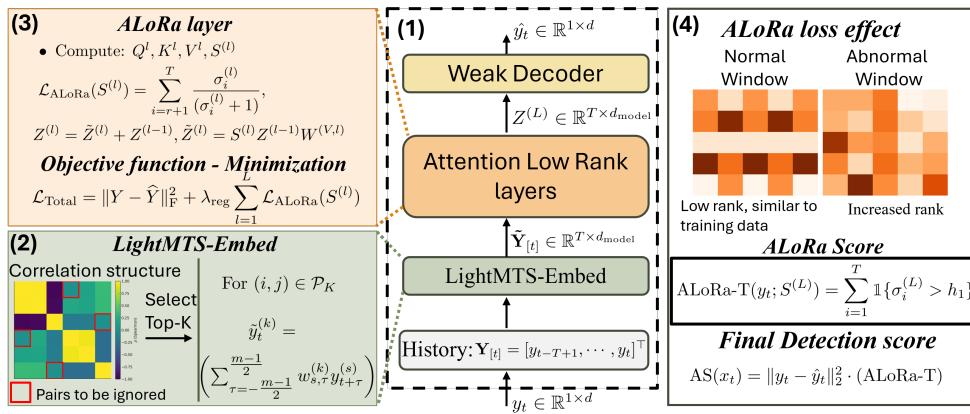
223 **Linear Projection for the Reconstruction:** The final reconstructed MTS at time step t is com-
 224 puted by applying a linear projection to the latent representation at the same time step:
 225

$$226 \hat{Y}_t = \mathbf{Z}_t^{(L)} \mathbf{W}^{\text{out}} \implies \hat{y}_t^{(k)} = \sum_{j=1}^{d_{\text{model}}} w_{jk}^{\text{out}} z_t^{(L,j)}, \quad (6)$$

229 where $\mathbf{W}^{\text{out}} \in \mathbb{R}^{d_{\text{model}} \times d}$ is the output projection matrix that maps the final latent representation to
 230 the original input space. This reveals that the model learns to reconstruct each time series using a
 231 linear combination of multiple STAR processes (Proposition 1).
 232

233 5 PROPOSED ANOMALY DIAGNOSIS METHOD

236 This section presents the proposed anomaly diagnosis method. Section 5.1 introduces the ALoRa-
 237 T model and its associated anomaly detection score, which together form the ALoRa-Det module,
 238 with an overview provided in Fig. 2. Section 5.2 then presents the ALoRa-Loc method for anomaly
 239 localization, illustrated in Fig. 1. The corresponding pseudocodes are provided in Appendix F and a
 240 computational analysis of the model is provided in Appendix C.



255 Figure 2: Overview of ALoRa-T and ALoRa-Det. (1) The architecture comprises the LightMTS-
 256 Embed module, ALoRa layers, and a decoder. (2) The embedding module exploits correlation
 257 structures to retain only the most significant time series pairs, significantly reducing complexity
 258 by avoiding unnecessary information. (3) ALoRa layers impose a low-rank constraint on the self-
 259 attention matrix through a novel loss and a regularization term in the objective function, producing
 260 a signal for abnormality. (4) During inference, anomalous windows yield higher attention ranks,
 261 which are captured by the ALoRa score, providing a clear indicator of anomalies.

263 5.1 ATTENTION LOW-RANK TRANSFORMER FOR ANOMALY DETECTION (ALoRA-DET)

265 ALoRa-Det consists of the ALoRa-T architecture together with its detection scoring method. We
 266 begin by describing the ALoRa-T architecture. The model begins with a lightweight MTS em-
 267 bedding module, followed by multi-head low-rank self-attention layers (ALoRa layers) with skip
 268 connections. As stated in Proposition 1 (see the proof in Section G), feedforward layers do not alter
 269 the structure of the latent space and are therefore omitted to avoid unnecessary complexity. The
 reconstruction step is implemented through a linear projection layer.

270 **LightMTS-Embed:** As shown in Section 4, embedding with 1D convolutions is equivalent to applying VMA filtering to MTS. However, using fully dense filters that mix all input series across embedding dimensions is both computationally expensive and less interpretable. In practice, not all time series are correlated with one another, and ignoring this fact can obscure the underlying dynamics in subsequent layers while further reducing interpretability. To address this, **each convolutional kernel is restricted to aggregate information from exactly two input series**. Specifically, for each output time series k , the (sparse) kernel contains only two non-zero weights $w_{i,j}^{(k)}$, corresponding to a particular pair of input series. To avoid including pairs without strong dependence, only the top- K pairs are retained, ranked by Spearman correlation on the training data (commonly $K = 512$). When the total number of possible pairs does not exceed K , all $\binom{d}{2}$ pairs are included. This design improves efficiency, promotes sparsity, and enhances interpretability while preserving performance. An ablation study on the parameter efficiency and performance of the proposed embedding module, compared with the standard Transformer embedding, is provided in Appendix D.1.

283 **Low-Rank Regularization in Self-Attention:** Based on Eq. (4), since the SA-matrices are the only
284 input-dependent learnable components, their spectral properties can serve as informative signals for
285 anomaly detection. In particular, the *rank* of a matrix, defined as the number of non-zero singular
286 values, provides a useful indicator of abnormal behavior. Empirically, we observe that the rank of
287 SA-matrices increases in the presence of anomalies, as illustrated in Fig. 3 (right). Motivated by
288 this observation, we propose **ALoRa layers**, which extend the standard self-attention mechanism
289 by explicitly promoting a low-rank structure in the attention matrices. This is achieved through
290 the *ALoRa loss*, a regularization term applied to each attention matrix $\mathbf{S}^{(l)}$, defined as the truncated
291 Geman nuclear norm (Geman & Yang, 1995), which enforces singular values close to zero.

$$\mathcal{L}_{\text{ALoRa}}(\mathbf{S}^{(l)}) = \sum_{i=r+1}^T \frac{\sigma_i^{(l)}}{(\sigma_i^{(l)} + 1)}, \quad (7)$$

295 where $\sigma_1^{(l)} \geq \dots \geq \sigma_T^{(l)} \geq 0$ are the singular values of $\mathbf{S}^{(l)}$. The parameter r specifies the number of
296 leading singular values that are preserved without penalty. Since $\mathbf{S}^{(l)}$ is row-stochastic (Eq. 2) with
297 fixed largest singular value $\sigma_1 = 1$, we set $r = 1$, since penalizing it is unnecessary. The ALoRa
298 layer is extended to the multi-head attention (MHA) setting by defining $\mathbf{S}^{(l)} = \frac{1}{H} \sum_{h=1}^H \mathbf{S}_h^{(l)}$, where
299 $\mathbf{S}_h^{(l)}$ denotes the self-attention matrix of head $h = 1, \dots, H$, and H is the number of heads.

301 **Training:** The objective function, $\mathcal{L}_{\text{Total}}$, used to train the model includes two key components:

$$\mathcal{L}_{\text{Total}} = \|\mathbf{Y} - \hat{\mathbf{Y}}\|_{\text{F}}^2 + \lambda_{\text{reg}} \sum_{l=1}^{l=L} \mathcal{L}_{\text{ALoRa}}(\mathbf{S}^{(l)}). \quad (8)$$

306 The first term penalizes the Frobenius norm of the reconstruction error, while the second term en-
307 courages a low-rank structure by summing the ALoRa losses of the SA-matrices $\mathbf{S}^{(l)}$ across all lay-
308 ers. The regularization strength is controlled by the parameter λ_{reg} . The effectiveness of the proposed
309 low-rank regularization is illustrated in Fig. 3, which compares the rank of self-attention matrices
310 with and without regularization. In both cases, rank differences between normal and anomalous
311 inputs are evident but become substantially more pronounced with regularization. Consequently,
312 anomalous patterns are more easily distinguished, providing a new abnormality signal.

313 **Model inference - Anomaly detection:** As in training, the MTS is processed in overlapping
314 windows to ensure fair reconstruction across all time steps, enabling unbiased anomaly detection
315 (Toscano & Recchioni, 2022). The detection anomaly score (AS) for the time steps t is defined as:

$$\text{AS}(y_t) = \|y_t - \hat{y}_t\|_2^2 \cdot (\text{ALoRa-T}) \in \mathbb{R}, \quad (9)$$

$$\text{ALoRa-T}(y_t; \mathbf{S}^{(L)}) = \sum_{i=1}^T \mathbf{1}_{\{\sigma_i^{(L)} > h_1\}} \in \mathbb{R}, \quad (10)$$

321 where $\{\sigma_1^{(L)}, \dots, \sigma_k^{(L)}\}$ are the singular values of the SA-matrix $\mathbf{S}^{(L)}$, used during the encoding-
322 decoding of the MTS at time step t . The indicator function $\mathbf{1}_{\{\cdot\}}$ returns 1 if the condition is true
323 and 0 otherwise. The threshold h_1 is introduced because, while the low-rank regularization encour-
ages many singular values to be close to zero, they are often not exactly zero. Moreover, when the

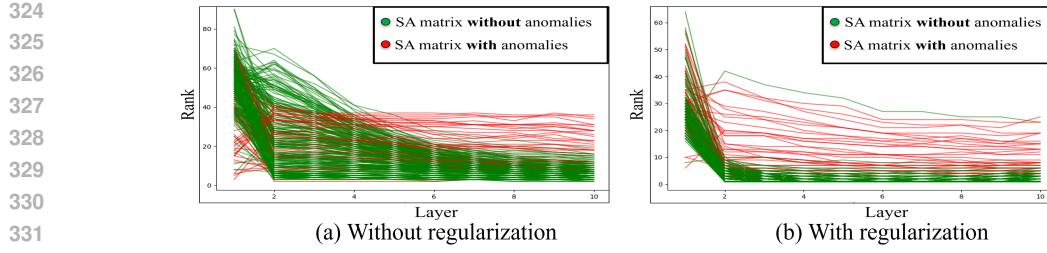


Figure 3: Layer-wise rank of SA-matrices across transformer layers, without (left) and with (right) the low-rank regularization term, on the SMD dataset. The discrepancy in rank between normal and anomalous inputs becomes more pronounced with regularization, improving anomaly sensitivity.

anomaly score exceeds a threshold, $AS(y_t) > h_2$, an anomaly alarm is triggered. Details about h_1 and h_2 selection is provided in Appendix A. The ALoRa-T score captures the temporal characteristics of anomalies, providing earlier and more reliable indications than other scoring functions, as shown in Figure 4 (see also Figs. 12 and 11 in Appendix F).

5.2 ALoRA-LOC FOR MTS ANOMALY LOCALIZATION AND INTEPRETATIONS

Effective anomaly localization in MTS requires a deep understanding of the model’s learning process. Practitioners often ask: *How does the model make decisions, and which time series influence each output?* However, due to a limited understanding of the learning dynamics, such questions remain largely unanswered in the context of deep learning. In contrast, linear regression remains popular for its simplicity and interpretability. Each reconstructed value follows $\hat{x}_i = \sum_j c_{ij} x_j + b_i$, where c_{ij} quantifies the influence of input x_j on output \hat{x}_i .

The proposed localization method (ALoRa-Loc) first addresses this gap by deriving weights that capture how each input time series contributes to both the learned latent representation and the reconstruction of each output time series. This is particularly important, since during model inference one can quantify how strongly the model relies on a given time series i , as measured by its relative contribution with magnitude $C_{ij} / \sum_k C_{ik}, j \in \{1, \dots, d\}$. For the anomaly localization task, such interpretability of the model’s decision process is even more critical. As demonstrated in Appendix E, anomalies in one series can propagate to others during the modeling process, making it essential to capture these propagation effects. The contribution weights E_{ij} and C_{ij} provide exactly this capability: E_{ij} traces the influence of input series i on the latent features, while C_{ij} extends this influence to the reconstructed outputs. By leveraging these weights, ALoRa-Loc enables tracing anomalies backward from the reconstructions to the latent space and ultimately to the originating input series, thereby supporting accurate anomaly localization. The derivation of the contribution weights is based on the theoretical analysis in Section 4, with analytical details in Appendix G.1.

Contribution of each input time series to the latent space: From Eq. (4) and Eq. (5), the latent representation is governed by the product of two matrices. The left product captures the covariance contributions, whereas the right product corresponds to shared weights across all time series and therefore does not affect the variability of individual contributions (see Appendix G.1 for details).

Accounting for the embedding module, the final contribution weights from the input space to the latent space are given by:

$$E_{ij} = \sum_{k=1}^{d_{\text{model}}} \left(\sum_{l=-\frac{m-1}{2}}^{\frac{m-1}{2}} w_{i,l}^{(k)} \right) b_{kj}, \quad \mathbf{B} = \begin{cases} \prod_{i=1}^L (\mathbf{W}^{(V,i)} + I), & \text{with skip connections (Eq. (5)),} \\ \prod_{i=1}^L \mathbf{W}^{(V,i)}, & \text{without skip connections (Eq. (4)).} \end{cases} \quad (11)$$

Here, b_{ij} denotes the (i, j) -th entry of $\mathbf{B} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{model}}}$, and $w_{i,l}^{(k)}$ are the embedding filter weights for input series i at lag l . In the absence of embedding, the weights reduce to $\mathbf{E} = \mathbf{B}$.

378 **Contribution of each input time series to the reconstruction space:** From Eq. (6), it follows that
 379 the overall contribution of the i -th input time series to the j -th reconstructed time series is given by:
 380

$$381 \quad 382 \quad 383 \quad C_{ij} = \sum_{k=1}^{d_{\text{model}}} w_{kj}^{\text{out}} \cdot E_{ik} \quad (12)$$

384 Based on these insights, we define the localization anomaly score (LAS) for each time series i at
 385 time t as:
 386

$$387 \quad 388 \quad LAS_t^{(i)} = \sum_{j=1}^d C_{ij} \|y_t^{(j)} - \hat{y}_t^{(j)}\|_2^2 \quad (13)$$

389 The intuition behind this score is that each term $(C_{ij} \|y_t^{(j)} - \hat{y}_t^{(j)}\|_2^2)$ represents the magnitude of the
 390 anomaly in the i -th time series that has propagated to the reconstruction of the j -th one. Summing
 391 over all j captures the total influence of the anomaly in time series i across the system. In practice,
 392 it is often more effective to sum only over the top- k dimensions with the largest C_{ij} , focusing on the
 393 most influential components. We refer to this variant as ALoRa-Loc (top- k).
 394

395 6 EXPERIMENTS

396 6.1 DATASETS, BASELINES AND EVALUATION METRICS

400 **Datasets & baselines:** We evaluated ALoRa on six widely used datasets from diverse domains:
 401 SWaT (Mathur & Tippenhauer, 2016) and HAI (Shin et al., 2021) from industrial control, SMD (Su
 402 et al., 2019) and PSM (Abdulaal et al., 2021) from IT monitoring, and MSL (Hundman et al.,
 403 2018) from spacecraft telemetry. For localization evaluation, we use SMD, SWaT, and the MSDS
 404 dataset (Nedelkoski et al., 2020). We compare ALoRa against a wide range of baselines, includ-
 405 ing classical methods (PCA (Shyu et al., 2003), KNN (Ramaswamy et al., 2000), IForest (Liu
 406 et al., 2008), LOF (Breunig et al., 2000), OC-SVM (Schölkopf et al., 2001)), notable deep-learning
 407 models (Omni-Anomaly (Su et al., 2019), Interfusion (Li et al., 2021), DAEMON (Chen et al.,
 408 2021), contrastive learning-based methods (Yang et al., 2023), and recent SOTA approaches such
 409 as Anomaly Transformer (A.T) (Xu et al., 2022), MEMTO (Song et al., 2023), NPSR (Lai et al.,
 410 2023), D^3R (Wang et al., 2023), and SARAD (Dai et al., 2024). For localizaiton we also compared
 411 with AERCA (Han et al., 2025). All results are from our own runs using official or public code with
 412 recommended settings. See Appendix B for details.

413 **Detection metrics:** Recent studies on MTS anomaly detection evaluation have shown that range-
 414 based metrics are the most appropriate, as they address the limitations of point-adjusted and point-
 415 wise evaluation methods (Liu & Paparrizos, 2024). Moreover, since MTS anomaly detection in-
 416 volves highly imbalanced datasets, F_1 -score-based metrics are considered the most reliable. Ac-
 417 cordingly, our evaluation focuses on the affiliation-based F_1 -score, precision, and recall (Huet et al.,
 418 2022). For fair comparison, we report the best F_1 -scores and the corresponding precision and re-
 419 call, avoiding method-specific thresholds. Additional results using the range-based F_1 -score (RF_1)
 420 (Hwang et al., 2019), VUS-AUC (VA), and VUS-PR (VPR) (Paparrizos et al., 2022) are provided in
 421 Appendix F. Additional details on the selection of evaluation metrics are provided in Appendix B.3.

422 **Localization metrics:** We evaluate localization performance using standard metrics such as Hit Rate
 423 (Su et al., 2019) and Normalized Discounted Cumulative Gain (NDCG) (Järvelin & Kekäläinen,
 424 2002). Additionally, we use the Interpretation Score (IPS), which measures how accurately anom-
 425 alyses are localized at the segment level. Further details are provided in Appendix B.3.

426 6.2 RESULTS

427 All experiments were conducted over five runs, and we report the average performance values. Due
 428 to space restrictions, the standard deviations are presented in Appendix F.
 429

430 **Detection Results:** As shown in Table 1, ALoRa-Det outperforms all baselines on four out of five
 431 datasets and ranks second on SWaT, it still significantly outperforms most other methods. Compared

to the second-best performing models, ALoRa-Det achieves absolute improvements in affiliation-based F_1 -score of 11.5% on SMD, 7.8% on PSM, 5.9% on MSL, and 8.9% on the HAI dataset. These results highlight the effectiveness and generalizability of our detection approach. Notably, our method is not only highly accurate, but its ALoRa-T score also provides a highly informative anomaly signal (see Figure 4 and Appendix F) that captures the characteristics of anomalies and enables much earlier detection than competing methods in most cases.

Table 1: **Detection performance.** P, R and F_1 denotes Precision, Recall, and F_1 -score respectively. The best F_1 -scores are highlighted in bold, and the second-best are underlined.

Method	SMD			PSM			MSL			SWaT			HAI		
	P	R	F_1	P	R	F_1	P	R	F_1	P	R	F_1	P	R	F_1
KNN	0.70	0.34	0.46	0.53	0.98	0.68	0.50	0.25	0.33	0.45	0.41	0.43	0.48	0.36	0.41
PCA	0.84	0.40	0.54	0.92	0.38	0.54	0.55	0.32	0.41	0.49	0.43	0.46	0.51	0.40	0.45
LOF	0.56	0.35	0.43	0.60	0.41	0.49	0.52	0.30	0.38	0.50	0.38	0.43	0.47	0.34	0.39
OC-SVM	0.44	0.28	0.34	0.88	0.47	0.62	0.53	0.29	0.38	0.46	0.36	0.40	0.42	0.30	0.35
IsolationForest	1.00	0.09	0.17	1.00	0.03	0.06	0.62	0.08	0.14	0.50	0.11	0.18	0.68	0.07	0.13
OmniAnomaly	0.68	0.66	0.67	0.70	0.64	0.67	0.5	0.93	0.64	0.49	0.52	0.50	0.51	0.39	0.44
InterFusion	0.67	0.66	0.67	0.69	0.65	0.67	0.51	0.92	0.64	0.47	0.40	0.43	0.50	0.93	0.64
A.T.	0.58	0.88	0.70	0.55	0.83	0.66	0.51	0.96	0.67	0.57	0.37	0.45	0.61	0.52	0.56
DAEMON	0.70	0.69	0.70	0.72	0.68	0.70	0.52	0.93	0.65	0.60	0.59	0.60	0.62	0.84	0.71
DCdetector	0.54	0.94	0.69	0.53	0.82	0.67	0.51	0.9	0.65	0.58	0.76	0.66	0.72	0.86	0.78
MEMTO	0.78	0.86	0.79	0.63	0.75	0.68	0.53	0.95	0.67	0.60	0.61	0.60	0.60	0.66	0.64
NPSR	0.77	0.98	0.87	0.67	0.89	0.76	0.52	0.98	0.68	0.98	0.16	0.28	0.98	0.65	0.79
D^3R	0.77	0.99	0.87	0.63	0.96	0.76	0.65	0.63	0.64	0.65	0.77	0.71	0.74	0.87	0.79
SARAD	0.88	0.67	0.78	0.73	0.505	0.56	0.56	0.83	0.67	0.54	0.33	0.41	0.51	0.70	0.66
ALoRa-Det	0.97	0.98	0.97	0.82	0.83	0.82	0.57	0.98	0.72	0.68	0.67	0.68	0.98	0.76	0.86

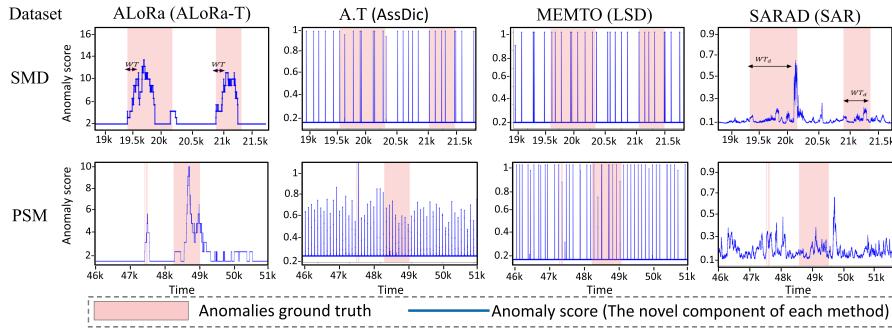


Figure 4: Anomaly scores for the SMD and PSM, datasets. The red segment indicates the ground truth of anomalies. MEMTO and Anomaly Transformer behave close to random guessing, making detection unreliable. While SARAD provides more informative scores does not fully capture anomaly temporal patterns, often resulting in increased waiting time (WT) until detection. In contrast, ALoRa-T score captures these patterns effectively, enabling faster and more precise anomaly detection. Appendix F presents additional visualizations, including more datasets and methods.

Computational Analysis: To evaluate the practical efficiency of ALoRa-Det, we compare its model size (number of learnable parameters), training time, and inference time per sample with other SOTA methods. Due to page limitations, the detailed results are provided in Appendix C. As shown in Table 6, the proposed method is computationally efficient and suitable for real-world applications.

Localization Results: We use ALoRa-Loc (top-2) for SMD and the standard version for the other two datasets. Table 2 shows that ALoRa-Loc consistently outperforms the compared methods. Moreover, ALoRa-Loc provides meaningful interpretations of the learning process, which are essential for practitioners, as discussed in Section 5.2.

486 Table 2: **Localization performance.** We report HR@P%, NDCG@P%, and range-based IPS@P%
 487 for $P \in \{100, 150\}$. The best scores are shown in bold, with the second-best underlined. The
 488 Interfusion method applies only to segment-based localization, so only the IPS metric is reported.
 489

490 Method	SMD				MSDS				SWaT									
	HR@P		NDCG@P		IPS@P		HR@P		NDCG@P		IPS@P		HR@P		NDCG@P		IPS@P	
492 MEMTO	0.32	0.48	0.26	0.36	0.19	0.28	0.14	0.33	0.10	0.22	0.02	0.04	0.01	0.01	0.01	0.02	0.05	0.05
493 OMNI	0.29	0.46	0.24	0.34	0.17	0.26	0.13	0.35	0.09	0.21	0.01	0.03	0.01	0.01	0.01	0.01	0.04	0.04
494 Interfusion	-	-	-	-	0.59	0.75	-	-	-	-	0.03	0.05	-	-	-	-	0.10	0.12
495 SARAD	<u>0.44</u>	<u>0.56</u>	<u>0.47</u>	<u>0.55</u>	0.61	<u>0.74</u>	0.25	0.40	0.31	0.39	0.04	<u>0.06</u>	0.03	<u>0.03</u>	0.03	<u>0.04</u>	<u>0.11</u>	<u>0.12</u>
496 DAEMON	0.26	0.39	0.31	0.40	0.24	0.26	0.23	0.36	0.27	0.34	0.02	0.03	0.02	0.02	0.02	0.03	0.06	0.07
497 AERCA	0.21	0.25	0.36	0.45	0.13	0.17	0.33	<u>0.56</u>	<u>0.31</u>	<u>0.40</u>	0.08	0.12	0.013	0.014	<u>0.031</u>	<u>0.04</u>	0.028	0.029
498 ALoRa-Loc	0.56	0.76	0.60	0.70	<u>0.60</u>	0.81	<u>0.30</u>	<u>0.57</u>	0.32	0.45	0.03	0.05	0.042	0.068	0.041	0.056	0.16	0.20

7 ABLATION STUDIES

500 To verify the effectiveness of the individual components of the proposed method, we conduct a se-
 501 ries of ablation studies. First, we evaluate the contribution of the proposed ALoRA loss, Eq. (7),
 502 and the corresponding ALoRA-T score, Eq. (10). The results of this analysis, presented in Table 3,
 503 demonstrate the effectiveness of both the ALoRA loss and the ALoRA-T score. Additional abla-
 504 tion studies exploring alternative implementations of the ALoRA loss are reported in Table 10 of
 505 Appendix D.3.

506 Table 3: **Ablation Study of ALoRA-Loss and ALoRA-Score Effectiveness:** We report the F_1
 507 score for each case.

509 Study Focus	510 ALoRA-Loss	511 ALoRA-Score	512 SMD	513 PSM	514 MSL	515 SWaT	516 HAI
517 Effectiveness of ALoRA-Loss	✓ (Yes)	✓ (Yes)	0.97	0.82	0.72	0.68	0.86
518 Effectiveness of ALoRA-Score	✓ (Yes)	✗ (No)	0.95	0.74	0.69	0.61	0.82

519 Next, for the Lightweight-MTS embedding module, we evaluate how specific design choices affect
 520 both performance and computational efficiency. In particular, we compare the model’s performance
 521 when all time series pairs are used versus when only the Top- K pairs are selected. The results
 522 presented in Table 4, illustrate that selecting only the top- K pairs, is more effective in terms of both
 523 performance and computational efficiency. Additional ablation studies for the Lightweight-MTS
 524 embedding module are provided in Appendix D.1. Finally, Appendix D.2 presents an ablation study
 525 supporting our decision to omit feed-forward layers. This choice is justified theoretically, since such
 526 layers do not alter the latent representation (Proposition 1), and is further confirmed experimentally,
 527 as shown in Table 9.

528 Table 4: **Ablation Study of LightMTS-Embed Design:** **Cost** : Computation Cost (Model Size
 529 in Millions (M) / Training Time (s) / Inference Time per Sample (ms)). **Det** : Anomaly detection
 530 performance measured using the F_1 score.

532 Study Focus	SMD				SWaT				HAI	
	533 Top-K	534 Correlation	535 Cost	536 Det.	537 Cost	538 Det.	539 Cost	540 Det.		
541 Our Model (Baseline)	✓ (Top-K)	Spearman	3.2M / 0.13 / 75	0.97	3.2M / 0.12 / 45	0.68	3.2M / 0.13 / 36	0.86		
542 (A) All Pairs Instead of Top-K	✗ (All Pairs)	Spearman	5.9M / 0.19 / 90	0.96	19.5 / 0.25 / 62	0.68	108M / 0.27 / 80	0.85		
543 (B) Pearson Instead of Spearman	✓ (Top-K)	Pearson	3.2M / 0.13 / 75	0.948	3.2M / 0.12 / 45	0.666	3.2M / 0.13 / 36	0.85		

8 CONCLUSION

544 We study Transformer encoders on MTS and contribute theoretical insights into their learning behav-
 545 ior. Based on this, we propose ALoRa-Det for anomaly detection, and ALoRa-Loc for localization.
 546 Our methods achieve SOTA performance on both tasks. Future work includes extending our theore-
 547 tical analysis from single-head to multi-head attention, and adapting our methods to settings where
 548 concept drift occurs alongside anomalous events.

540 **9 REPRODUCIBILITY STATEMENT**
 541

542 To ensure the reproducibility of our results, we provide the source code in the supplementary material,
 543 along with a README file that outlines the steps required to reproduce the experiments
 544 and results presented in this work. Appendix A provides details on parameter selection, while Ap-
 545 pendix B describes the datasets used and the baseline methods. For ease of reproducibility the
 546 supplementary material also contains data from two datasets, while the remaining datasets can be
 547 obtained from the sources listed in Appendix B.1. For each compared method, we provide source
 548 code links (Appendix B.2) and follow the recommended settings from the original studies.
 549

550 **REFERENCES**
 551

552 Ahmed Abdulaal, Zhuanghua Liu, and Tomer Lancewicki. Practical approach to asynchronous
 553 multivariate time series anomaly detection and localization. In *Proceedings of the 27th ACM*
 554 *SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 2485–2494. ACM, 2021.
 555 doi: 10.1145/3447548.3467364. URL <https://dl.acm.org/doi/10.1145/3447548.3467364>.

556 Markus M. Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Jörg Sander. Lof: identifying density-
 557 based local outliers. In *Proceedings of the 2000 ACM SIGMOD International Conference on*
 558 *Management of Data*, SIGMOD ’00, pp. 93–104, New York, NY, USA, 2000. Association for
 559 Computing Machinery. ISBN 1581132174. doi: 10.1145/342009.335388. URL <https://doi.org/10.1145/342009.335388>.

560 Peter J. Brockwell and Richard A. Davis. *Time Series: Theory and Methods*. Springer Series in
 561 Statistics. Springer, 2nd edition, 2002. ISBN 978-0387955411.

562 Xuanhao Chen, Liwei Deng, Feiteng Huang, Chengwei Zhang, Zongquan Zhang, Yan Zhao, and Kai
 563 Zheng. Daemon: Unsupervised anomaly detection and interpretation for multivariate time series.
 564 In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, pp. 2225–2230, 2021.
 565 doi: 10.1109/ICDE51399.2021.00228.

566 Xuanhao Chen, Liwei Deng, Yan Zhao, and Kai Zheng. Adversarial autoencoder for unsupervised
 567 time series anomaly detection and interpretation. In *Proceedings of the Sixteenth ACM Interna-*
 568 *tional Conference on Web Search and Data Mining*, WSDM ’23, pp. 267–275. Association for
 569 Computing Machinery, 2023. ISBN 9781450394079.

570 Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea
 571 Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser,
 572 David Benjamin Belanger, Lucy J Colwell, and Adrian Weller. Rethinking attention with per-
 573 formers. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Ua6zuk0WRH>.

574 N. Cressie and C.K. Wikle. *Statistics for Spatio-Temporal Data*. CourseSmart. Wiley, 2011. ISBN
 575 9780471692744.

576 Zhihao Dai, Ligang He, Shuanghua Yang, and Matthew Leeke. Sarad: Spatial association-aware
 577 anomaly detection and diagnosis for multivariate time series. In *Advances in Neural Infor-*
 578 *mation Processing Systems (NeurIPS)*, 2024. URL <https://openreview.net/forum?id=gmf5Aj01Hz>.

579 Demetrios G. Eliades, Kleanthis Malialis, Stelios Vrachimis, and Marios M. Polycarpou. Smart
 580 water networks as cyber-physical-socio-environmental systems. *IEEE Transactions on Industrial*
 581 *Cyber-Physical Systems*, 3:95–105, 2025. doi: 10.1109/TICPS.2024.3520558.

582 D. Geman and Chengda Yang. Nonlinear image recovery with half-quadratic regularization. *IEEE*
 583 *Transactions on Image Processing*, 4(7):932–946, 1995. doi: 10.1109/83.392335.

584 Borjan Geshkovski, Cyril Letrouit, Yury Polyanskiy, and Philippe Rigollet. The emergence of clus-
 585 ters in self-attention dynamics. In *Advances in Neural Information Processing Systems (NeurIPS)*,
 586 volume 36, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/b2b3e1d9840eba17ad9bbf073e009afe-Abstract-Conference.html.

594 Xiao Han, Saima Absar, Lu Zhang, and Shuhan Yuan. Root cause analysis of anomalies in mul-
 595 tivariate time series through granger causal discovery. In *The Thirteenth International Confer-
 596 ence on Learning Representations*, 2025. URL <https://openreview.net/forum?id=k38Th3x4d9>.

598 Alexis Huet, Jose Manuel Navarro, and Dario Rossi. Local evaluation of time series anomaly de-
 599 tection algorithms. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Dis-
 600 covery and Data Mining*, pp. 635–645, 2022. doi: 10.1145/3534678.3539339. URL <https://dl.acm.org/doi/10.1145/3534678.3539339>.

602 Kyle Hundman, Valentino Constantinou, Christopher Laporte, Ian Colwell, and Tom Soderstrom.
 603 Detecting spacecraft anomalies using lstms and nonparametric dynamic thresholding. In *Pro-
 604 ceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data
 605 Mining*, pp. 387–395. ACM, 2018. doi: 10.1145/3219819.3219845. URL <https://dl.acm.org/doi/10.1145/3219819.3219845>.

607 Won-Seok Hwang, Jeong-Han Yun, Jonguk Kim, and Hyoung Chun Kim. Time-series aware pre-
 608 cision and recall for anomaly detection: Considering variety of detection result and addressing
 609 ambiguous labeling. In *Proceedings of the 28th ACM International Conference on Information
 610 and Knowledge Management*, CIKM ’19, pp. 2241–2244, New York, NY, USA, 2019. Associa-
 611 tion for Computing Machinery. ISBN 9781450369763. doi: 10.1145/3357384.3358118. URL
 612 <https://doi.org/10.1145/3357384.3358118>.

613 Kalervo Järvelin and Jaana Kekäläinen. Cumulated gain-based evaluation of ir techniques. *ACM
 614 Trans. Inf. Syst.*, 20(4):422–446, October 2002. ISSN 1046-8188. doi: 10.1145/582415.582418.
 615 URL <https://doi.org/10.1145/582415.582418>.

617 Siwon Kim, Kukjin Choi, Hyun-Soo Choi, Byunghan Lee, and Sungroh Yoon. Towards a rigor-
 618 ous evaluation of time-series anomaly detection. *arXiv preprint arXiv:2109.05257*, 2021. URL
 619 <https://arxiv.org/abs/2109.05257>.

620 Elias Kyriakides and Marios Polycarpou. *Intelligent monitoring, control, and security of critical
 621 infrastructure systems*, volume 565. Springer, 2015.

623 Chih-Yu Lai, Fan-Keng Sun, Zhengqi Gao, Jeffrey H. Lang, and Duane S. Boning. Nominality
 624 score conditioned time series anomaly detection by point/sequential reconstruction. In *Advances
 625 in Neural Information Processing Systems (NeurIPS)*, 2023. URL <https://openreview.net/forum?id=ljgM3vNqfQ>.

627 Kwei-Herng Lai, Daochen Zha, Junjie Xu, Yue Zhao, Guanchu Wang, and Xia Hu.
 628 Revisiting time series outlier detection: Definitions and benchmarks. In J. Van-
 629 schoren and S. Yeung (eds.), *Proceedings of the Neural Information Processing
 630 Systems Track on Datasets and Benchmarks*, volume 1, 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/file/ec5decca5ed3d6b8079e2e7e7bacc9f2-Paper-round1.pdf.

633 Dan Li, Dacheng Chen, Baihong Jin, Lei Shi, Jonathan Goh, and See-Kiong Ng. Mad-gan:
 634 Multivariate anomaly detection for time series data with generative adversarial networks. In
 635 *Artificial Neural Networks and Machine Learning – ICANN 2019: Text and Time Series*,
 636 volume 11731 of *Lecture Notes in Computer Science*, pp. 703–716. Springer, 2019. doi:
 637 10.1007/978-3-030-30490-4_56. URL https://link.springer.com/chapter/10.1007/978-3-030-30490-4_56.

639 Zhihan Li, Youjian Zhao, Jiaqi Han, Ya Su, Rui Jiao, Xidao Wen, and Dan Pei. Multivariate
 640 Time Series Anomaly Detection and Interpretation using Hierarchical Inter-Metric and Tem-
 641 poral Embedding. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Dis-
 642 covery & Data Mining*, pp. 3220–3230. ACM, August 2021. ISBN 978-1-4503-8332-5. doi:
 643 10.1145/3447548.3467075.

644 Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *Proceedings of the 8th IEEE
 645 International Conference on Data Mining (ICDM)*, pp. 413–422, 2008. doi: 10.1109/ICDM.
 646 2008.17. URL <https://cs.nju.edu.cn/zhouzh/zhouzh.files/publication/icdm08b.pdf>.

648 Qinghua Liu and John Paparrizos. The elephant in the room: Towards a reliable time-series anomaly
 649 detection benchmark. In *The Thirty-eighth Conference on Neural Information Processing Systems*
 650 *Datasets and Benchmarks Track*, 2024. URL <https://openreview.net/forum?id=R6kJtWstGy>.

651

652 Aditya P. Mathur and Nils Ole Tippenhauer. SWaT: A water treatment testbed for research and
 653 training on ICS security. In *Proceedings of the 2016 International Workshop on Cyber-Physical*
 654 *Systems for Smart Water Networks (CySWater)*, pp. 31–36. IEEE, 2016. doi: 10.1109/CySWater.
 655 2016.7469060.

656

657 Sasho Nedelkoski, Jasmin Bogatinovski, Ajay Kumar Mandapati, Soeren Becker, Jorge Car-
 658 doso, and Odej Kao. Multi-source distributed system data fornbsp;ai-powered analytics. pp.
 659 161–176, Berlin, Heidelberg, 2020. Springer-Verlag. ISBN 978-3-030-44768-7. doi: 10.1007/
 660 978-3-030-44769-4_13. URL https://doi.org/10.1007/978-3-030-44769-4_13.

661

662

663 Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth
 664 64 words: Long-term forecasting with transformers. In *The Eleventh International Confer-
 665 ence on Learning Representations*, 2023. URL <https://openreview.net/forum?id=Jbdc0vTOcol>.

666

667 John Paparrizos, Paul Boniol, Themis Palpanas, Ruey S. Tsay, Aaron Elmore, and Michael J.
 668 Franklin. Volume under the surface: a new accuracy evaluation measure for time-series
 669 anomaly detection. *Proc. VLDB Endow.*, 15(11):2774–2787, July 2022. ISSN 2150-8097. doi:
 670 10.14778/3551793.3551830. URL <https://doi.org/10.14778/3551793.3551830>.

671

672 Daehyung Park, Yuuna Hoshi, and Charles C. Kemp. A multimodal anomaly detector for robot-
 673 assisted feeding using an lstm-based variational autoencoder. *IEEE Robotics and Automation*
 674 *Letters*, 3(3):1544–1551, 2018. doi: 10.1109/LRA.2018.2801475.

675

676 Sridhar Ramaswamy, Rajeev Rastogi, and Kyuseok Shim. Efficient algorithms for mining outliers
 677 from large data sets. In *Proceedings of the 2000 ACM SIGMOD International Conference on*
 678 *Management of Data*, SIGMOD '00, pp. 427–438, New York, NY, USA, 2000. Association for
 679 Computing Machinery. ISBN 1581132174.

680

681 Bernhard Schölkopf, John C. Platt, John C. Shawe-Taylor, Alex J. Smola, and Robert C. Williamson.
 682 Estimating the support of a high-dimensional distribution. *Neural Comput.*, 13(7):1443–1471,
 683 July 2001. ISSN 0899-7667. doi: 10.1162/089976601750264965. URL <https://doi.org/10.1162/089976601750264965>.

684

685 Lifeng Shen, Zhuocong Li, and James T. Kwok. Timeseries anomaly detection using temporal
 686 hierarchical one-class network. In *Advances in Neural Information Processing Systems*, vol-
 687 ume 33, pp. 13016–13026, 2020. URL <https://proceedings.neurips.cc/paper/2020/file/97e401a02082021fd24957f852e0e475-Paper.pdf>.

688

689 Hyeok-Ki Shin, Woomyo Lee, Jeong-Han Yun, and Byung-Gil Min. Two ics security datasets and
 690 anomaly detection contest on the hil-based augmented ics testbed. In *Proceedings of the 14th*
 691 *Cyber Security Experimentation and Test Workshop*, CSET '21, pp. 36–40, New York, NY, USA,
 692 2021. Association for Computing Machinery. ISBN 9781450390651. doi: 10.1145/3474718.
 693 3474719. URL <https://doi.org/10.1145/3474718.3474719>.

694

695 Mei-Ling Shyu, Shu-Ching Chen, Kanoksri Sarinnapakorn, and LiWu Chang. A novel anomaly
 696 detection scheme based on principal component classifier. Technical report, Department of Elec-
 697 trical and Computer Engineering, University of Miami, Coral Gables, FL, USA, 2003. Technical
 698 Report.

699

700 Junho Song, Keonwoo Kim, Jeonglyul Oh, and Sungzoon Cho. Memto: Memory-guided trans-
 701 former for multivariate time series anomaly detection. In *Advances in Neural Information*
 702 *Processing Systems (NeurIPS)*, 2023. URL <https://openreview.net/forum?id=UFW67uduJd>.

702 Ya Su, Youjian Zhao, Chenhao Niu, Rong Liu, Wei Sun, and Dan Pei. Robust Anomaly Detection
 703 for Multivariate Time Series through Stochastic Recurrent Neural Network. In *Proceedings of*
 704 *the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp.
 705 2828–2837. ACM, July 2019. ISBN 978-1-4503-6201-6. doi: 10.1145/3292500.3330672.

706 Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In
 707 *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML’17,
 708 pp. 3319–3328. JMLR.org, 2017.

709 Siyi Tang, Jared Dunnmon, Khaled Kamal Saab, Xuan Zhang, Qianying Huang, Florian Dubost,
 710 Daniel Rubin, and Christopher Lee-Messer. Self-supervised graph neural networks for improved
 711 electroencephalographic seizure analysis. In *International Conference on Learning Representa-*
 712 *tions*, 2022. URL https://openreview.net/forum?id=k9bx1EfHI_-.

713 Giacomo Toscano and Maria Cristina Recchioni. Bias-optimal vol-of-vol estimation: the role of
 714 window overlapping. *Decisions in Economics and Finance*, 45(1):137–185, 2022. doi: 10.
 715 1007/s10203-021-00349-4. URL <https://link.springer.com/article/10.1007/s10203-021-00349-4>.

716 Ruey S Tsay, Daniel Peña, and Alan E Pankratz. Outliers in multivariate time series. *Biometrika*,
 717 87(4):789–804, 2000.

718 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
 719 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg,
 720 S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural*
 721 *Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

722 Chengsen Wang, Zirui Zhuang, Qi Qi, Jingyu Wang, Xingyu Wang, Haifeng Sun, and Jianxin Liao.
 723 Drift doesn’t matter: Dynamic decomposition with diffusion reconstruction for unstable multi-
 724 variate time series anomaly detection. In *Advances in Neural Information Processing Systems*
 725 (*NeurIPS*), 2023. URL <https://openreview.net/forum?id=aW5bSuduF1>.

726 Jiehui Xu, Haixu Wu, Jianmin Wang, and Mingsheng Long. Anomaly transformer: Time series
 727 anomaly detection with association discrepancy. In *International Conference on Learning Repre-*
 728 *sentations (ICLR)*, 2022. URL <https://arxiv.org/abs/2110.02642>.

729 Yiyuan Yang, Chaoli Zhang, Tian Zhou, Qingsong Wen, and Liang Sun. Dcdetector: Dual attention
 730 contrastive representation learning for time series anomaly detection. In *Proceedings of the 29th*
 731 *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD ’23, pp. 3033–3045,
 732 New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701030. doi:
 733 10.1145/3580305.3599295. URL <https://doi.org/10.1145/3580305.3599295>.

734 George Zerveas, Srideepika Jayaraman, Dhaval Patel, Anuradha Bhamidipaty, and Carsten Eickhoff.
 735 A transformer-based framework for multivariate time series representation learning. In *Proceed-
 736 ings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, KDD ’21,
 737 pp. 2114–2124. Association for Computing Machinery, 2021. ISBN 9781450383325.

738 Chuxu Zhang, Dongjin Song, Yuncong Chen, Xinyang Feng, Cristian Lumezanu, Wei Cheng,
 739 Jingchao Ni, Bo Zong, Haifeng Chen, and Nitesh V. Chawla. A deep neural network for un-
 740 supervised anomaly detection and diagnosis in multivariate time series data. In *Proceedings of*
 741 *the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applica-*
 742 *tions of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances*
 743 *in Artificial Intelligence*, AAAI’19/IAAI’19/EAAI’19. AAAI Press, 2019. ISBN 978-1-57735-
 744 809-1.

745

746

747

748

749

750

751

752

753

754

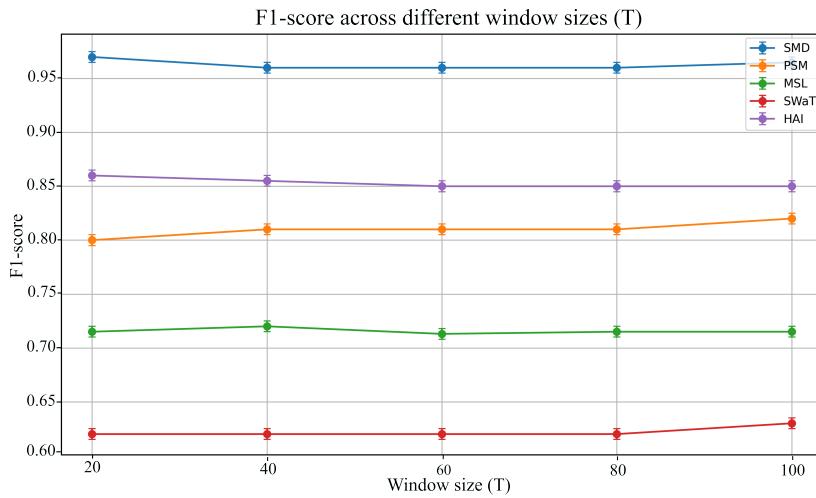
755

756 APPENDIX TO: LOW RANK TRANSFORMER FOR MULTIVARIATE TIME SERIES
 757 ANOMALY DETECTION AND LOCALIZATION
 758

759 **A TRAINING DETAILS**

760 **ALoRa-T architecture design and hyperparameter selection:** The ALoRa-T model consists of
 761 three self-attention layers with skip connections. Each layer uses $H = 8$ attention heads across all
 762 datasets. Training is performed using the ADAM optimizer with a learning rate of 10^{-4} , and early
 763 stopping is applied to prevent overfitting. A fixed regularization parameter of $\lambda_{\text{reg}} = 10$ is used for
 764 all datasets.

765 **Window Size T :** The window size T is an important hyperparameter that determines how much
 766 past information the model can utilize at each time step. To assess its impact, we conduct an ablation
 767 study evaluating ALoRa-Det on the validation set with different values of T , as shown in Figure 5.



788 Figure 5: Validation F1-scores across different window sizes for each dataset. The average over 5
 789 runs is reported along with the standard deviation.

790 The figure shows that ALoRa-Det’s performance is not sensitive to the choice of window size. It is
 791 worth noting that smaller T values lead to faster training and inference, while larger values result in
 792 longer processing times. Based on these considerations, we select the following window sizes for
 793 each dataset: SMD (20), PSM (100), MSL (20), SWaT (20), and HAI (20). For the MSDS dataset
 794 used for localization, the window size is set to 100.

795 **Selection of h_1 , used in eq. (10) :** We present the process followed for selecting the threshold
 796 parameter h_1 . As discussed in Section 5, ALoRa-T incorporates a low-rank regularization term,
 797 which encourages the self-attention matrices $S^{(L)}$ to have a reduced-rank structure. The selection
 798 procedure is automated and data-driven, as it must adapt to each dataset. With this procedure, h_1 is
 799 automatically determined for any new dataset based on the spectral properties of the self-attention
 800 matrices observed under normal operating conditions. Specifically, we analyze the distribution of the
 801 fourth and fifth largest eigenvalues of $S^{(L)}$ across the training sequence and track their trajectories.
 802 The threshold h_1 is then automatically chosen as the **maximum value** observed between these two
 803 eigenvalue trajectories. This ensures that only the dominant eigenvalues, typically corresponding
 804 to a rank of 3–4, are preserved. The choice to preserve the 3–4 largest eigenvalues is informed by
 805 the findings of Geshkovski et al. (2023). Figure 6 illustrates this selection process. Based on this
 806 analysis, we set $h_1 = 10^{-2}$ for SMD, 10^{-3} for PSM, 3×10^{-2} for SWaT and MSL, and 2×10^{-1}
 807 for HAI.

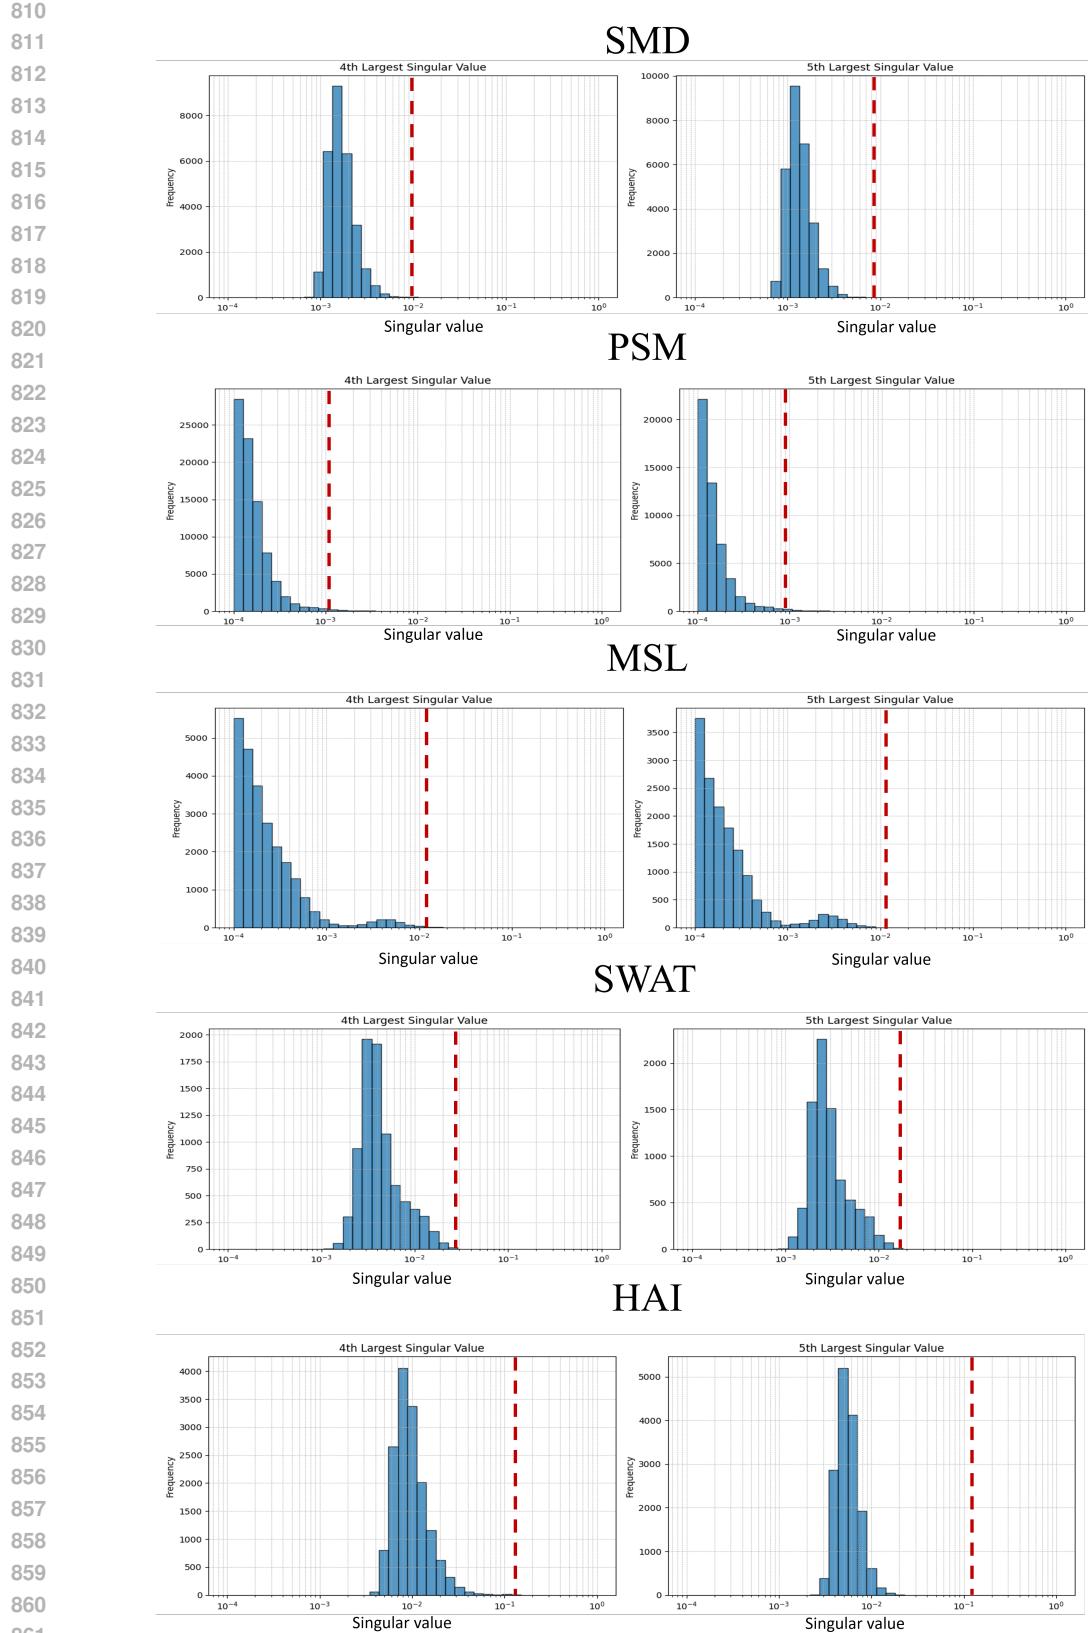


Figure 6: Distribution of the fourth and fifth largest eigenvalues of $S^{(L)}$ computed over normal training data. The red dashed line indicates the selected h_1 value, chosen to zero out smaller eigenvalues and ensure a typical rank no larger than 4.

864 **Selection of h_2 :** The parameter h_2 is the threshold used to convert raw anomaly scores into binary
865 decision labels. Such a threshold is required for **all** anomaly detection methods to provide anomaly
866 predictions. To avoid introducing bias from any particular threshold-selection strategy and to ensure
867 a **fair comparison across all methods**, we follow a widely adopted evaluation protocol: for each
868 method, including our own, we report the performance obtained using the **best achievable thresh-**
869 **old**, i.e., the threshold that yields the highest performance. This approach is standard in the anomaly
870 detection literature (Dai et al., 2024; Lai et al., 2023). In summary, selecting h_2 in this manner
871 guarantees **fairness** when comparing baselines: it evaluates the **quality of the anomaly scores pro-**
872 **duced by each method**, rather than the particular threshold-selection mechanism used to convert
873 those scores into labels.

874

875

876

877

878

879

880

881

882

883

884

885

886

887

888

889

890

891

892

893

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

918 **B EXPERIMENTAL SETTING**
919920 **B.1 DATASETS**
921922 The datasets used in our experiments are publicly available and can be downloaded from the links
923 provided below:
924925 • **Secure Water Treatment (SWaT)** Li et al. (2019): This dataset captures data from a
926 real-world industrial water treatment system monitored by 51 sensors over 11 days. The
927 first 7 days contain normal operations, while the last 4 include 41 manually injected
928 anomalies.
929 *Download:* [https://itrust.sutd.edu.sg/itrust-labs-home/](https://itrust.sutd.edu.sg/itrust-labs-home/itrust-labs_swat/)
930 *itrust-labs_swat/*931 • **Server Machine Dataset (SMD)** Su et al. (2019): A large-scale dataset with 38 features
932 collected over five weeks from server machines in a major Internet company, covering
933 various system performance metrics.
934 *Download:* [https://github.com/NetManAIOps/OmniAnomaly/tree/](https://github.com/NetManAIOps/OmniAnomaly/tree/master/ServerMachineDataset)
935 *master/ServerMachineDataset*936 • **Pooled Server Metrics (PSM)** Abdulaal et al. (2021): Collected internally at eBay, this
937 dataset contains 25 dimensions from application server nodes, including 13 weeks of train-
938 ing data and 8 weeks of test data.
939 *Download:* <https://github.com/eBay/RANSynCoders/tree/main/data>940 • **Mars Science Laboratory (MSL)** Hundman et al. (2018): Provided by NASA, this dataset
941 includes telemetry data with 55 features recorded from the MSL rover during space mis-
942 sions.
943 *Download:* <https://github.com/khundman/telem anom>944 • **HIL-based Augmented ICS Security Dataset (HAI)** Shin et al. (2021): Collected from a
945 realistic ICS testbed enhanced with a Hardware-In-the-Loop simulator, this dataset contains
946 78 dimensions from sensors and actuators under normal conditions and during 38 simulated
947 cyber-attacks.
948 *Download:* <https://github.com/icsdataset/hai>949 • **Multi-Source Distributed System (MSDS)** Nedelkoski et al. (2020): Collected from a
950 cloud-based OpenStack testbed, this dataset contains multi-source monitoring
951 data—including metrics, logs, and traces—under normal and fault-injected conditions. It
952 includes root cause labels, making it suitable for anomaly localization task.
953 *Download:* <https://zenodo.org/records/3484801>954 Table 5 summarizes key statistics of the datasets used in our experiments, including the number of
955 features, training and test samples, and the anomaly proportion in the test set.
956957 Table 5: Summary of the five benchmark datasets used in our experiments. ‘Dim’ denotes the
958 number of features. ‘Train’ and ‘Test’ indicate the number of samples in the training and test sets.
959 ‘Anomaly rate (%)’ shows the percentage of anomalies in the test set.
960

	Dim	Application	Train	Test	Anomaly rate(%)
SWaT	51	Water	495,000	449,919	0.121
SMD	38	Server	28479	28479	0.156
PSM	25	Server	132481	87841	0.278
MSL	55	Space	58317	73729	0.105
HAI	78	Power	921603	402005	0.223
MSDS	10	AIOps	29268	29286	0.72

968
969 **Dataset Preprocessing:** For each dataset, we first normalize each time series individually before
970 passing the data to the model. This normalization is performed using the mean and standard devi-
971 ation computed from the training set. The same normalization parameters are then applied to the
972

972 corresponding test set. For the SWaT and HAI datasets, which contain a very large number of time
 973 steps due to high-frequency (per-second) sampling, we downsample the data to one-minute intervals
 974 by averaging the values over every 60-second window.
 975

976 B.2 BASELINES 977

978 We reproduced the results of all baseline methods using their official or publicly available imple-
 979 ments on GitHub. For each baseline, we followed the configuration settings recommended in
 980 their respective papers to ensure optimal performance. The baselines and their implementations can
 981 be found at the following links:

- 982 • **kNN**: <https://github.com/yzhao062/pyod>
- 983 • **PCA**: <https://github.com/yzhao062/pyod>
- 984 • **LOF**: <https://github.com/yzhao062/pyod>
- 985 • **OCSVM**: <https://github.com/yzhao062/pyod>
- 986 • **IForest**: <https://github.com/yzhao062/pyod>
- 987 • **OmniAnomaly**: <https://github.com/NetManAIOps/OmniAnomaly>
- 988 • **InterFusion**: <https://github.com/ryu-ichiro/InterFusion>
- 989 • **Anomaly Transformer**: <https://github.com/thuml/Anomaly-Transformer>
- 990 • **DAEMON**: <https://github.com/Sherlock-C/DAEMON>
- 991 • **DCdetector**: <https://github.com/DAMO-DI-ML/KDD2023-DCdetector>
- 992 • **MEMTO**: <https://github.com/dreamgonfly/memto>
- 993 • **NPSR**: <https://github.com/1ai-chihyu/NPSR>
- 994 • **SARAD**: <https://github.com/ZhihaoDai/SARAD>
- 995 • **D3R**: <https://github.com/Wang-Xinyu666/D3R>

1001 B.3 EVALUATION METRICS 1002

1003 **Detection:** The evaluation of multivariate time series anomaly detection has attracted growing re-
 1004 search attention, largely due to the challenge of measuring model performance in a manner that
 1005 aligns with real-world temporal decision-making. Anomalies usually span continuous ranges rather
 1006 than isolated points. As a result, point-wise evaluation that treats each timestamp independently is
 1007 inappropriate because it ignores temporal continuity and unfairly penalizes slight delays or near-
 1008 misses in detection. Traditionally, the point-adjustment method was used to address this limitation,
 1009 where detecting any point within an anomalous segment suffices to consider the entire segment as
 1010 an observed anomaly (Xu et al., 2022; Song et al., 2023; Shen et al., 2020). However, this ap-
 1011 proach often inflates performance metrics, as even a random anomaly scoring method can achieve
 1012 performance comparable to that of a well-informed method, which effectively captures the temporal
 1013 characteristics of anomalies (Kim et al., 2021; Huet et al., 2022). This limitation of point-adjusted
 1014 metrics is further illustrated in Figure 7.

1015 Recognizing these limitations, many studies (Liu & Paparrizos, 2024; Hwang et al., 2019) have sug-
 1016 gested range-based evaluation metrics as the most appropriate choice for multivariate time series
 1017 anomaly detection. For instance, Hwang et al. (2019) proposed the range-based F_1 , precision, and
 1018 recall (RF_1 , R-P, R-R). An advancement over these are the affiliation-based F1, precision, and recall
 1019 (Huet et al., 2022), which demonstrate superior performance and robustness compared to the afore-
 1020 mentioned metrics. Unlike traditional metrics, this approach accounts for the temporal alignment
 1021 between predicted and actual anomalies. Precision is computed as the average directed distance
 1022 from predicted anomaly events to their nearest ground truth counterparts, while recall measures the
 1023 average directed distance from ground truth events to their closest predictions. For a more detailed
 1024 explanation of how these metrics are computed, we refer the reader to the original paper (Huet et al.,
 1025 2022). Paparrizos et al. (2022) introduced VUS-ROC and VUS-PR as range-aware alternatives
 1026 to traditional AUC. However, because MTS anomaly detection is a highly imbalanced classification
 1027 problem, F_1 -like metrics are the most appropriate and reliable for evaluation, while VUS metrics can

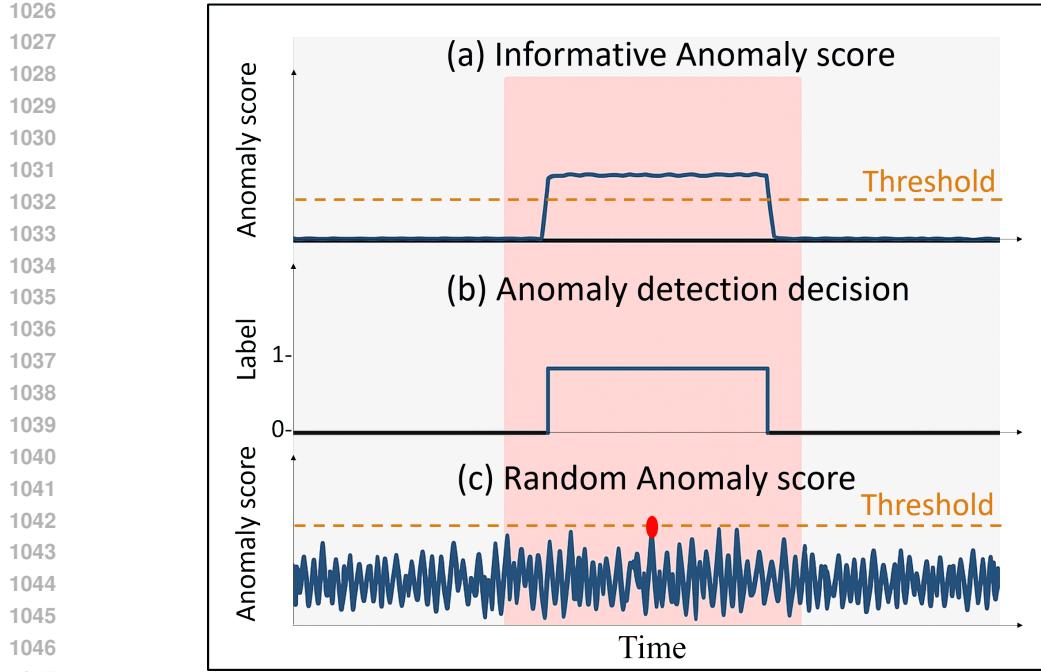


Figure 7: Line (a) shows the anomaly scores from an informative anomaly scoring approach, while Line (c) represents a random scoring method. The middle section displays the binary decisions (where anomaly = 1, otherwise 0) made by both methods using the point adjustment technique, which ultimately led to the same outcome. The red segments indicate the ground truth abnormal time steps, while the red dot in Line (c) marks the only time step that, by chance, exceeded the anomaly threshold.

provide additional insights but should be used together with F1-like metrics. Given the superior robustness and effectiveness of affiliation-based metrics, we adopt the affiliation-based F_1 -score as our primary evaluation metric. Results for the range-based F_1 -score (RF_1), along with the VUS-ROC (V_{ROC}) and VUS-PR (V_{PR}) metrics, are reported in Table 11. Nevertheless, due to the highly imbalanced nature of MTS anomaly localization, F1-like metrics remain the most appropriate choice, as they explicitly capture the trade-off between precision and recall.

Localization: At each time step t , anomaly localization is evaluated by identifying which features (time series) contribute to the detected anomaly. The ground-truth for localization is provided as a binary label vector $G_i \in \{0, 1\}^N$, where $G_i = 1$ indicates that the i -th feature is anomalous at time t , and $G_i = 0$ otherwise.

After computing the anomaly score $AS_t^{(i)}$ for each feature using eq.(13), features are ranked in descending order of their scores to identify those most likely responsible for the anomaly. The top- k features are selected as:

$$\Gamma_{t@P\%} = \text{Top-}k \text{ features ranked by } AS_t^{(i)}, \quad \text{where } k = \lceil |G_i| \times P\% \rceil$$

For example, if there are 3 anomalous features (i.e., $|G_i| = 3$) and we evaluate at $P = 150$, then $k = 5$. The localization performance is then quantified using the *Hit Rate* (HR) at $P\%$, which measures the fraction of truly anomalous features that appear in the top- k predictions:

$$HR_{t@P\%} = \frac{|G_i \cap \Gamma_{t@P\%}|}{|G_i|} \quad (14)$$

1080
 1081 **Normalized Discounted Cumulative Gain (NDCG):** Let $r_j \in \{0, 1\}$ indicate whether the j -th
 1082 ranked feature in $\Gamma_{t@P\%}$ is truly anomalous (i.e., belongs to G_i). The *Discounted Cumulative Gain*
 1083 (*DCG*) and its ideal counterpart (*IDCG*) are defined as:

$$\text{DCG}_{t@P\%} = \sum_{j=1}^k \frac{r_j}{\log_2(j+1)}$$

$$\text{IDCG}_t = \sum_{j=1}^{|G_i|} \frac{1}{\log_2(j+1)}$$

1091 The resulting *NDCG* is:

$$\text{NDCG}_{t@P\%} = \frac{\text{DCG}_{t@P\%}}{\text{IDCG}_t}, \quad (15)$$

1096 NDCG ranges from 0 to 1, with higher values indicating better-ranked localization.

1099 **Interpretation Score (IPS):** To evaluate anomaly localization at the segment level, we use the
 1100 *Interpretation Score (IPS)* Li et al. (2021). For each anomaly segment S_i , we first compute a single
 1101 score for each feature by taking the maximum anomaly score within the segment:

$$AS_{S_i}^{(j)} = \max_{t \in S_i} AS_t^{(j)} \quad (16)$$

1105 We then identify the top-ranked features as the predicted anomalous dimensions P_{S_i} , and compare
 1106 them against the ground-truth anomalous features G_{S_i} . The IPS measures the average proportion of
 1107 correctly predicted features across all segments:

$$\text{IPS} = \frac{1}{N} \sum_{i=1}^N \frac{|G_{S_i} \cap P_{S_i}|}{|G_{S_i}|}, \quad (17)$$

1112 where N is the total number of anomalous segments. Each segment is equally weighted in the final
 1113 score.

1115 C COMPUTATIONAL ANALYSIS

1117 To evaluate the practical efficiency of **ALoRa-Det**, we compare its model size (number of learnable
 1118 parameters), total training time (in seconds), and inference time per sample (in milliseconds) against
 1119 three state-of-the-art baselines: MEMTO, SARAD, and D^3R . The results across three benchmark
 1120 datasets are reported in Table 6.

1122 Table 6: Computational efficiency comparison. The table reports the number of learnable parameters
 1123 (in millions), total training time (s), and inference time per sample (ms).

1125 Method	1126 SMD (d=38) Params / Inf/ms / Train(s)	1126 SWaT (d=51) Params / Inf/ms / Train(s)	1126 HAI (d=78) Params / Inf/ms / Train(s)
1127 MEMTO	5.9M / 1.10 / 108	5.9M / 5.71 / 57	6M / 6.8 / 96
1128 SARAD	9.6M / 0.11 / 147	9.6 / 0.16 / 652	9.6M / 0.48 / 126
1129 D3R	52.3M / 4.83 / 994	52.3 / 6.27 / 283	52.5 M / 21.6 / 1077
1130 ALoRa-Det	3.2 M / 0.13 / 75	3.2M / 0.12 / 45	3.2M / 0.13 / 36

1132 The results in Table 6 indicate that **ALoRa-Det** almost always outperforms the competing methods
 1133 across all three metrics. Unlike other approaches, the total number of parameters remains constant

1134 as input dimensionality increases, while both inference time per sample and training time remain
 1135 consistently low. These advantages become more evident as the dimensionality of the MTS grows,
 1136 highlighting the scalability of the proposed method to high-dimensional settings. Taken together,
 1137 these findings confirm that **ALoRa-Det** is not only more efficient than existing baselines but also
 1138 highly suitable for real-time applications.

1139

1140

D ABLATION STUDIES - SUPPLEMENT

1141

1142

D.1 LIGHTWEIGHT MTS EMBEDDING PARAMETER EFFICIENCY

1143

1144

Table 7 shows the parameter efficiency of the lightweight embedding used in our method, comparing
 it to standard fully dense filters.

1145

1146

Table 7: Parameter comparison between standard fully dense filters with model dimension, $d_{\text{model}} =$
 512 and the proposed Lightweight MTS Embedding module.

1147

1148

1149

Input Size d	Standard Weights	Custom Weights
10	15,360	90
20	30,720	380
40	61,440	1,024
50	76,800	1,024
100	153,600	1,024
200	307,200	1,024

1150

1151

In addition, we compare the standard and Lightweight-MTS embedding methods in terms of detection performance to demonstrate that the proposed lightweight approach does not compromise effectiveness. On the contrary, it maintains, and in some cases even improves performance. Figure 8 presents the F_1 scores achieved by both embeddings, along with boxplots showing the distribution of reconstruction errors on normal test samples, providing further insight into the reconstruction behavior of each method.

1152

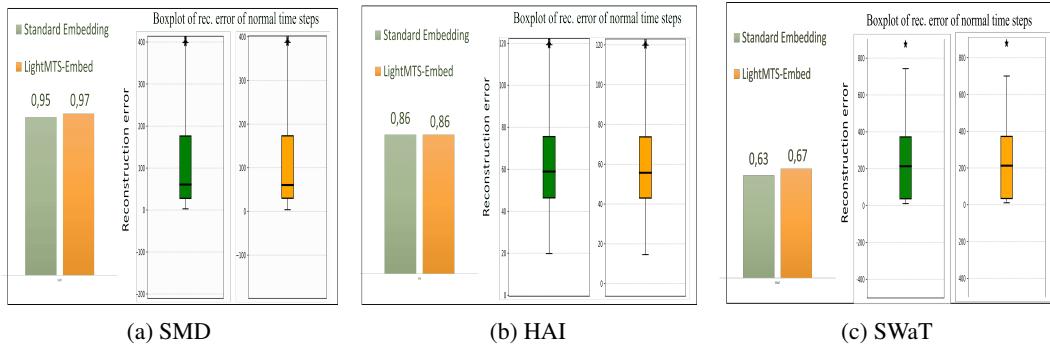
1153

1154

1155

1156

1157



1158

1159

1160

1161

1162

1163

1164

1165

Figure 8: Model performance under the standard and lightweight embeddings for the SMD, HAI, and SWaT datasets. In each subplot, the figure on the left shows the F_1 score comparison, and the figure on the right presents the boxplot of reconstruction errors on normal test samples.

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

Why LightMTS embedding module use pairwise interactions: The proposed LightMTS embedding module was designed to reduce complexity by restricting interactions to **pairwise** relationships between time series. Increasing the number of time series involved in each interaction, from pairs to triplets or higher-order groups, would lead to a **combinatorial explosion** in the number of possible interactions. In such cases, the complexity would become comparable to that of standard fully connected or convolutional embeddings.

1188 **Top- K Selection for the LightMTS-Embed Module:** The choice to select the top-512 pairs fol-
 1189 lowers standard design practice in Transformer architectures, where the embedding dimension is com-
 1190 monly set to 512, as introduced in the original Transformer model (Vaswani et al., 2017). This setting
 1191 provides a strong balance between representational capacity and computational efficiency, enabling
 1192 consistently low training and inference times across all datasets (Table 6). To assess the sensitivity
 1193 of the model to this parameter, an ablation study for different values of K , is presented in Table D.1.
 1194 The results indicate that the model is not highly sensitive to this hyperparameter, as performance
 1195 remains similar for $K = 512$ and $K = 256$, with the latter offering slightly lower accuracy but im-
 1196 proved computational efficiency. However, as the model’s computational efficiency is already very
 1197 strong, as demonstrated in Table 6, $K = 512$ remains a robust default choice. It generalizes well
 1198 across datasets from diverse domains and provides sufficient representation-learning capacity even
 1199 for multivariate time-series datasets with very high dimensionality.
 1200

1200 **Table 8: Ablation Study on the Top- K pairs selection:** The table reports the F_1 score average over
 1201 5 runs with the standard deviation for different values of the parameter K .
 1202

K	SMD	PSM	MSL	SWaT	HAI
512	0.97 ± 0.012	0.82 ± 0.014	0.72 ± 0.011	0.68 ± 0.010	0.86 ± 0.015
256	0.965 ± 0.014	0.82 ± 0.014	0.72 ± 0.011	0.67 ± 0.011	0.854 ± 0.015
128	0.95 ± 0.012	0.805 ± 0.015	0.713 ± 0.012	0.66 ± 0.012	0.85 ± 0.013
64	0.94 ± 0.014	0.78 ± 0.015	0.67 ± 0.012	0.65 ± 0.013	0.84 ± 0.012

D.2 ABLATION STUDY FOR REMOVING FEED-FORWARD LAYERS

This ablation study examines how omitting feed-forward layers affects both computational efficiency and anomaly detection and localization performance. As stated in Proposition 1, feed-forward layers do not alter the latent representation, indicating that their contribution to performance is minimal. The experimental results in Table 9 confirm this theoretical insight. Specifically, when feed-forward layers are included, performance either decreases or remains unchanged, while computational cost increases substantially.

Table 9: **Ablation Study for Removing Feed-Forward Layers:** The table reports the number of learnable parameters in millions (Cost), the anomaly detection performance using the F_1 -score (Det.), and the localization performance using the Hit-Rate@100 metric (Loc.), reported only for the datasets where localization labels are available.

FFN	SMD			PSM			MSL			SWaT			HAI		
	Cost	Det.	Loc.	Cost	Det.	Loc.	Cost	Det.	Loc.	Cost	Det.	Loc.	Cost	Det.	
✓ (Yes)	6.3	0.96	0.56	2.2	0.79	0.56	6.3	0.72	0.56	6.3	0.67	0.03	6.3	0.854	
✗ (No)	3.2	0.97	0.50	1.1	0.82	0.50	3.2	0.72	3.2	0.68	0.042	3.2	0.86		

D.3 ALORA-LOSS PLACEMENT

The proposed ALoRa loss (Eq. 7) is applied to the average self-attention matrix across heads rather than to each head individually. To justify this choice, we provide an ablation study comparing the two approaches in terms of both computational cost and performance.

As shown in Table 10, both variants achieve nearly identical detection performance across all datasets. However, applying the ALoRa loss to the *average* self-attention matrix is slightly more computationally efficient, which is why it was preferred. Importantly, both approaches are capable of achieving our goal of reducing the rank while maintaining the effectiveness of the model. As a final note, even if the head-wise application of the ALoRa loss had been preferred, the proposed method would still remain more computationally efficient than other state-of-the-art methods.

1242 Table 10: **Ablation Study on ALoRA-Loss Placement:** Head-Wise vs Averaged Self-Attention.
 1243 Reported are the training time (in seconds) required for each case and the corresponding detection
 1244 performance (F1-score).

ALoRA-Loss Placement	SMD		SWaT		HAI	
	Training time	F_1	Training time	F_1	Training time	F_1
Head-Wise ALoRA-Loss	155	0.97	98	0.67	78	0.86
OURS: Averaged Attention ALoRA-Loss	75	0.97	0.45	0.68	0.36	0.86

E ANOMALOUS EFFECT ATTRIBUTION VIA EFFECT WEIGHTS

To examine how anomalies in one time series affect others through shared latent representations, we design a controlled simulation using a simple self-attention model. This setup also demonstrates how the *effect weights* introduced in Section 4 quantify influence in both latent and reconstruction spaces.

We generate bivariate i.i.d. data $X \in \mathbb{R}^{T \times 2}$, where:

$$x_t^{(1)} \sim \begin{cases} \mathcal{N}(\mu_1 + \Delta, \sigma_1^2), & \text{for } t_1 \leq t < t_2 \\ \mathcal{N}(\mu_1, \sigma_1^2), & \text{otherwise} \end{cases}, \quad x_t^{(2)} \sim \mathcal{N}(\mu_2, \sigma_2^2) \quad \forall t$$

We set $T = 500$, $t_1 = 200$, and $t_2 = 300$. The model consists of a 2-layer, 1-head self-attention mechanism. We fix the value matrices $W_V^{(l)}$ for layers $l = 1, 2$, as well as the output projection matrix W_{out} , while the remaining parameters are kept learnable:

$$W_V^{(1)} = I_2, \quad W_V^{(2)} = \begin{bmatrix} 0.2 & 0.7 \\ 0.8 & 0.3 \end{bmatrix}, \quad W_{\text{out}} = \begin{bmatrix} 0.1 & 0.9 \\ 0.9 & 0.1 \end{bmatrix}$$

This controlled setup allows us to observe how a localized anomaly in one input series propagates through the latent space and affects the reconstructed time series. As shown in Fig. 9, the effect weights E_{1j} and C_{1j} reflect how the anomaly in the first input series propagates through the latent space and influences the reconstructed time series, respectively. Together, they provide meaningful and interpretable attributions of anomaly influence across the model’s internal and output representations.

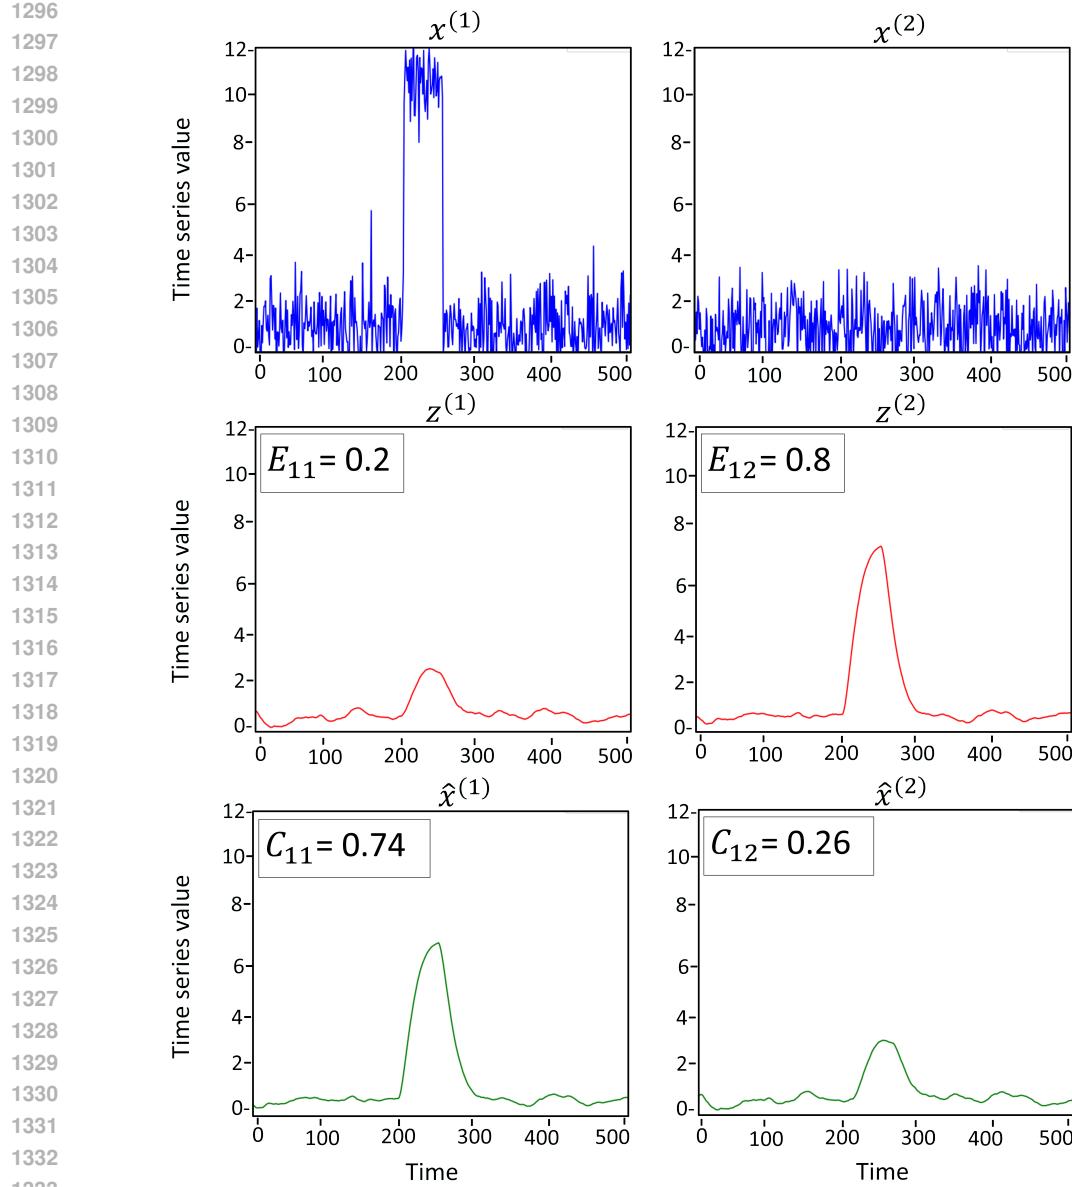


Figure 9: Averaged results over 100 simulation runs. **Top:** Input time series with an anomaly introduced in series 1 between time steps 200 and 300. **Middle:** Latent representations; effect weights E_{1j} quantify the contribution of the anomaly in series 1 to each latent dimension. **Bottom:** Reconstructed time series; effect weights C_{1j} reflect how the anomaly influences the reconstruction of each output series.

F EXPERIMENTS - SUPPLEMENT

To assess the performance of the proposed ALORA-DET method, we provide additional results using three evaluation metrics. The selection of these metrics is explained in Appendix B.3. Table 11 presents the performance of the proposed method and all compared methods using the range-based F1-score (R_{F1}), as well as the VUS-ROC (V_{ROC}) and VUS-PR (V_{PR}).

Effectiveness of the ALORA-T Score Across Anomaly Types: To assess the effectiveness of the ALORA-T score, and to further provide empirical evidence that the attention rank increases in the presence of anomalies, we examine how the rank of the metric—and thus the behavior of

1350 Table 11: **Range-aware evaluation metrics.** RF_1 denotes Range-based F1-score; V_{ROC} and V_{PR}
 1351 denote the Volume Under the Surface ROC and precision, respectively. Best scores are highlighted
 1352 in bold, and second-best are underlined.

1353

1354	Method	SMD			PSM			MSL			SWaT			HAI		
		RF_1	V_{ROC}	V_{PR}												
1356	KNN	0.15	0.42	0.09	0.31	0.40	0.28	0.11	0.37	0.08	0.10	0.44	0.35	0.12	0.09	0.12
1357	PCA	0.18	0.65	0.11	0.48	0.58	0.42	0.15	0.62	0.20	0.13	0.68	0.51	0.21	0.63	0.13
1358	IsolationForest	0.19	0.68	0.1	0.23	0.54	0.33	0.14	0.57	0.17	0.13	0.42	0.12	0.17	0.69	0.12
1359	OC-SVM	0.11	0.58	0.08	0.19	0.53	0.37	0.13	0.59	0.18	0.03	0.62	0.52	0.23	0.64	0.14
1360	OmniAnomaly	0.17	0.33	0.07	0.14	0.29	0.05	0.12	0.50	0.11	0.021	0.37	0.12	0.20	0.73	0.22
1361	InterFusion	0.18	0.36	0.08	0.19	0.34	0.12	0.14	0.53	0.12	0.09	0.40	0.14	0.28	0.75	0.23
1362	A.T.	0.11	0.50	0.10	0.19	0.26	0.37	0.14	0.51	0.12	0.13	0.61	0.22	0.09	0.59	0.21
1363	MEMTO	0.26	0.52	0.11	0.29	0.50	0.29	0.13	0.50	0.12	0.14	0.68	0.39	0.16	0.60	0.22
1364	NPSR	<u>0.43</u>	0.89	0.48	0.49	<u>0.63</u>	<u>0.47</u>	<u>0.20</u>	0.53	0.13	0.33	0.83	0.64	<u>0.20</u>	0.75	0.44
1365	D^3R	0.1	<u>0.91</u>	<u>0.57</u>	0.2	0.65	0.46	0.18	<u>0.61</u>	0.15	0.09	0.84	0.38	0.1	<u>0.89</u>	<u>0.61</u>
1366	SARAD	0.34	0.80	0.16	0.25	0.62	0.42	0.12	0.53	0.13	0.171	0.59	0.33	0.07	0.70	0.21
1367	ALoRa-Det	0.50	0.93	0.58	<u>0.45</u>	0.69	0.50	0.28	0.58	0.15	<u>0.23</u>	<u>0.78</u>	<u>0.57</u>	0.62	0.92	0.62

1364

1365

1366

1367 the ALoRa-T score—changes under different anomaly types. These anomaly types are the ones
 1368 formalized and defined in the studies (Lai et al., 2021) and (Tsay et al., 2000).

1369

1370

1371

1372

1373

1374

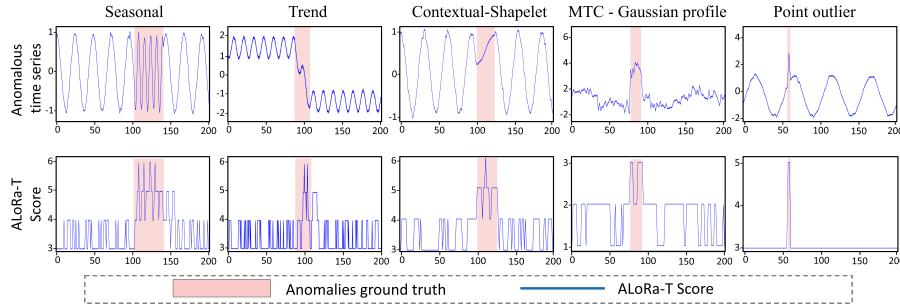
1375

1376

1377

1378

1379



1380

1381

1382

1383 Figure 10: ALoRa-T scores for different anomaly types, following the definitions and categorizations presented in (Lai et al., 2021) and (Tsay et al., 2000). The plotted ALoRa-T scores are averaged over 5 runs.

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

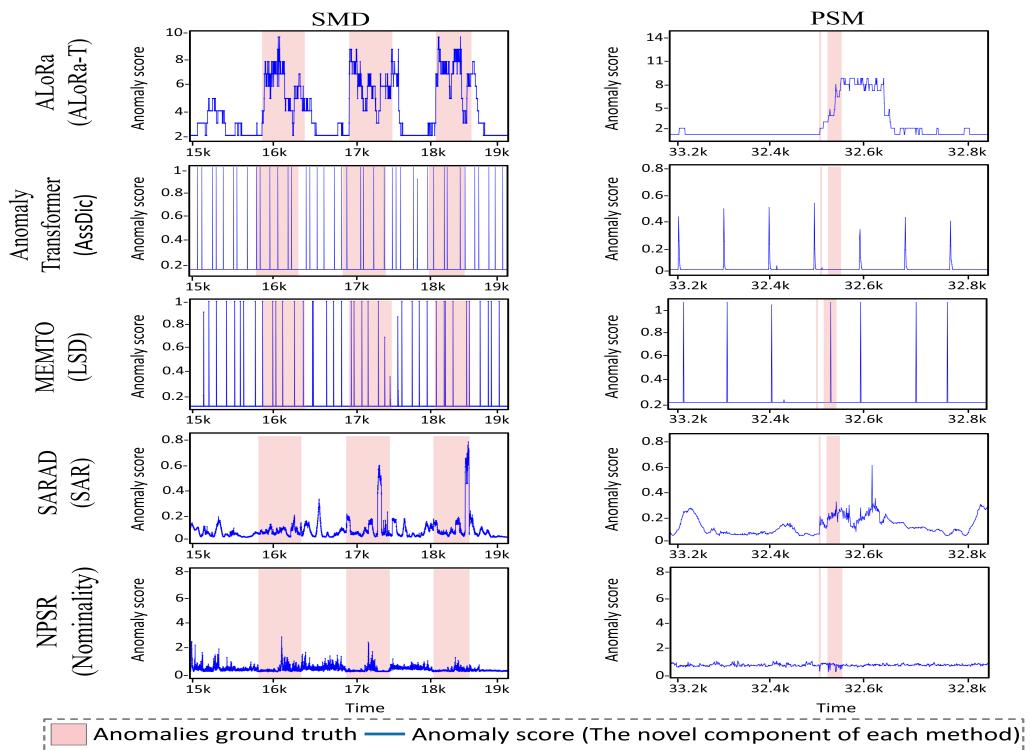
1401

1402

1403

Based on Fig. 10, there is clear evidence that, across a variety of anomaly types, the rank of the self-attention matrices, and consequently the corresponding ALoRa-T score, increases in the presence of anomalies. This consistent behavior across different anomaly categories demonstrates the robustness and generalizability of the ALoRa-T score in detecting anomalous patterns.

1404
 1405 **Comparison of ALoRa-T Scores with State-of-the-Art Methods:** To support and extend the
 1406 findings in the main paper, this appendix presents additional visualizations to demonstrate the
 1407 generalizability of the ALoRa-T anomaly scoring method. First, we provide additional abnormal
 1408 segments from the SMD and PSM datasets (Fig. 11) to reinforce our earlier results. Second, we include
 1409 anomaly scores from the SWaT and HAI datasets (Fig. 12) to illustrate ALoRa-T’s effectiveness
 1410 across multiple domains. Based on Fig. 12 and 11, the ALoRa-T score continues to stand out for
 1411 its effectiveness to capture the temporal characteristics of anomalies, providing timely and highly
 1412 informative responses to anomalous patterns. The score rises quickly and remains aligned with
 1413 the duration and structure of the anomalies. In some cases, the ALoRa-T score remains high for
 1414 a short period of time, even after the labeled anomaly segment ends. However, this behavior can
 1415 be meaningfully interpreted: the self-attention matrices at post-anomaly timesteps still incorporate
 1416 abnormal information from earlier points. The score gradually decreases after T timesteps—where
 1417 T is the window size, once the attention mechanism no longer includes the earlier abnormal data.
 1418 This demonstrates the effectiveness of the ALoRa-T score to accurately detect anomalies and to
 1419 provide interpretable signals for both true and false alarms. In contrast, MEMTO and Anomaly
 1420 Transformer still behave similarly to their performance in the main results (see Fig. 4), where their
 1421 outputs resemble random guessing. This makes their detection signals unreliable. SARAD and
 1422 NPSR produce more informative patterns. However, they frequently fail to capture the temporal
 1423 dynamics of anomalies accurately, often resulting in delayed or unsuccessful detection.
 1424
 1425
 1426
 1427
 1428
 1429
 1430
 1431
 1432
 1433
 1434
 1435
 1436
 1437
 1438
 1439
 1440
 1441
 1442
 1443
 1444
 1445



1446 Figure 11: Anomaly scores for the SMD and PSM, datasets. The red segment indicates the ground
 1447 truth of anomalies
 1448
 1449
 1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457

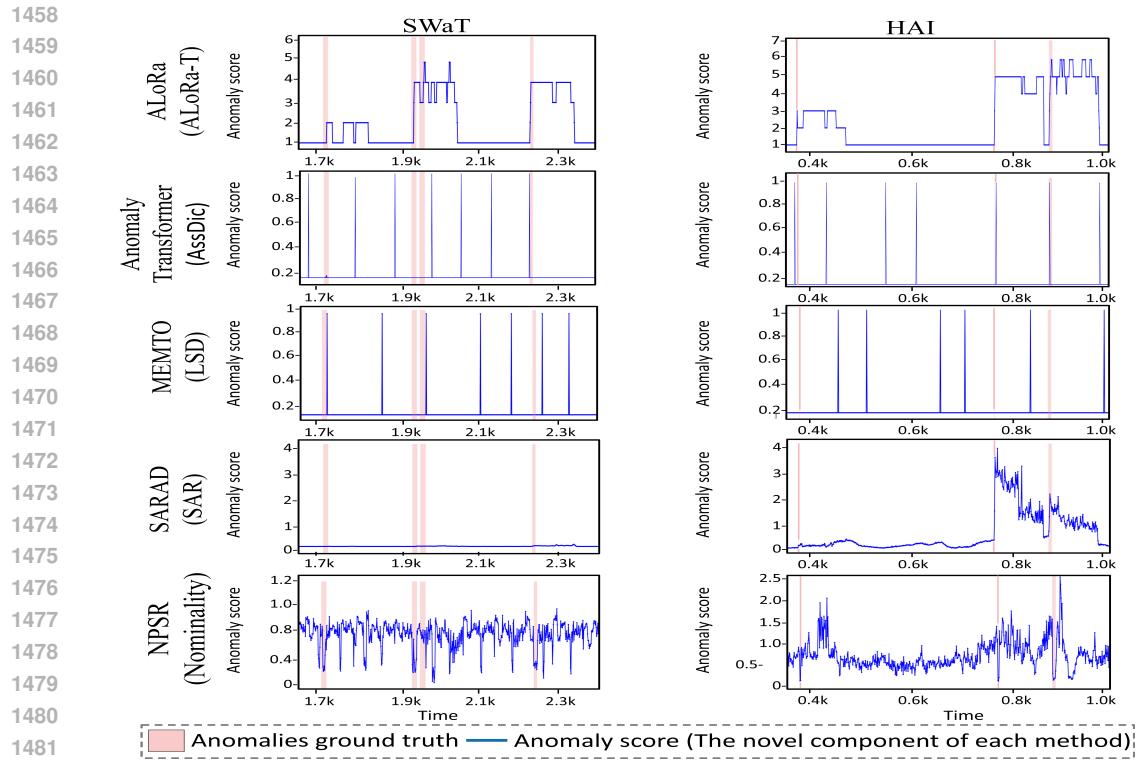


Figure 12: Anomaly scores for the SWat and HAI, datasets. The red segment indicates the ground truth of anomalies

1485
 1486
 1487
 1488
 1489
 1490
 1491
 1492
 1493
 1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511

1512 **Standard Deviation of Anomaly Diagnosis Results:** We report the standard deviations of the
 1513 anomaly diagnosis results presented in Section 6, as well as for the additional evaluation metrics
 1514 presented in Table 11. Specifically, Table 12 shows the standard deviations of the detection results
 1515 reported in Table 1, while Table 14 reports the standard deviations for the results in Table 11. Finally,
 1516 Table 13 presents the standard deviations of the localization results shown in Table 2.

1517
 1518 **Table 12: Standard deviation of detection metrics**, for the results in Table 1. P, R and F_1 denote
 1519 the standard deviation of Precision, Recall, and F_1 -score respectively.
 1520

Method	SMD			PSM			MSL			SWaT			HAI		
	P	R	F_1												
KNN	0.006	0.004	0.005	0.003	0.002	0.002	0.009	0.006	0.008	0.003	0.002	0.002	0.006	0.004	0.005
PCA	0.009	0.007	0.008	0.006	0.004	0.005	0.003	0.002	0.002	0.009	0.007	0.008	0.003	0.002	0.002
LOF	0.003	0.002	0.002	0.009	0.007	0.008	0.006	0.004	0.005	0.006	0.005	0.005	0.003	0.002	0.002
OC-SVM	0.006	0.004	0.005	0.006	0.004	0.005	0.003	0.002	0.002	0.009	0.007	0.008	0.009	0.007	0.008
IsolationForest	0.009	0.007	0.008	0.003	0.002	0.002	0.006	0.004	0.005	0.006	0.004	0.005	0.009	0.007	0.008
OmniAnomaly	0.022	0.018	0.020	0.012	0.009	0.010	0.022	0.017	0.020	0.011	0.009	0.010	0.012	0.008	0.010
InterFusion	0.012	0.009	0.010	0.021	0.018	0.020	0.012	0.009	0.010	0.021	0.018	0.020	0.012	0.009	0.010
A.T.	0.021	0.023	0.020	0.020	0.03	0.030	0.020	0.016	0.018	0.017	0.014	0.016	0.02	0.016	0.017
DAEMON	0.014	0.009	0.012	0.010	0.009	0.011	0.011	0.010	0.011	0.018	0.017	0.018	0.013	0.008	0.011
DCdetector	0.022	0.021	0.021	0.023	0.024	0.023	0.017	0.020	0.019	0.016	0.015	0.015	0.019	0.018	0.018
MEMTO	0.065	0.055	0.060	0.009	0.006	0.007	0.009	0.007	0.008	0.025	0.019	0.022	0.053	0.045	0.050
NPSR	0.003	0.002	0.002	0.002	0.001	0.001	0.002	0.001	0.001	0.003	0.002	0.002	0.004	0.003	0.003
D^3R	0.005	0.003	0.004	0.004	0.003	0.003	0.002	0.001	0.001	0.008	0.006	0.007	0.006	0.004	0.005
SARAD	0.10	0.028	0.022	0.13	0.08	0.021	0.021	0.044	0.018	0.025	0.022	0.013	0.017	0.10	0.110
ALoRa-Det	0.011	0.011	0.012	0.013	0.014	0.014	0.012	0.011	0.011	0.012	0.011	0.010	0.010	0.009	0.015

1534
 1535
 1536
 1537
 1538 **Table 13: Standard deviation of localization scores**, for the results in Table 2. The Interfusion
 1539 method applies only to segment-based localization, so the standard deviation is reported only for the
 1540 IPS metric.

Method	SMD				MSDS				SWaT									
	HR@P 100	HR@P 150	NDCG@P 100	NDCG@P 150	IPS@P 100	IPS@P 150	HR@P 100	HR@P 150	NDCG@P 100	NDCG@P 150	IPS@P 100	IPS@P 150						
MEMTO	0.020	0.019	0.024	0.027	0.022	0.041	0.031	0.048	0.021	0.039	0.005	0.018	0.0012	0.0035	0.0003	0.0012	0.0011	0.028
OMNI	0.041	0.041	0.045	0.045	0.028	0.029	0.04	0.03	0.03	0.04	0.002	0.002	0.02	0.025	0.033	0.022	0.01	0.01
Interfusion	-	-	-	-	0.029	0.031	-	-	-	-	0.001	0.001	-	-	-	-	0.01	0.01
SARAD	0.018	0.017	0.023	0.025	0.021	0.043	0.03	0.05	0.02	0.04	0.006	0.02	0.001	0.003	0.0002	0.001	0.001	0.003
DAEMON	0.05	0.08	0.07	0.06	0.006	0.01	0.06	0.06	0.05	0.05	0.01	0.01	0.007	0.006	0.005	0.005	0.01	0.009
AERCA	0.02	0.01	0.01	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.02	0.02
ALoRa-Loc	0.05	0.05	0.06	0.06	0.09	0.07	0.01	0.05	0.01	0.02	0.001	0.003	0.006	0.005	0.006	0.001	0.01	0.007

1548
 1549
 1550
 1551 **Table 14: Standard deviation** for the results in Table 11. RF_1 denotes Range-based F_1 -score; V_{ROC}
 1552 and V_{PR} denote the Volume Under the Surface ROC and PR, respectively.

Method	SMD			PSM			MSL			SWaT			HAI		
	RF_1	V_{ROC}	V_{PR}												
KNN	0.006	0.008	0.004	0.005	0.006	0.005	0.007	0.009	0.006	0.004	0.006	0.005	0.006	0.007	0.005
PCA	0.009	0.012	0.006	0.008	0.010	0.007	0.006	0.009	0.005	0.010	0.013	0.007	0.007	0.009	0.006
IsolationForest	0.012	0.014	0.007	0.010	0.012	0.009	0.009	0.011	0.006	0.011	0.013	0.008	0.010	0.012	0.007
OC-SVM	0.010	0.011	0.006	0.011	0.012	0.008	0.007	0.010	0.008	0.012	0.014	0.009	0.011	0.013	0.007
OmniAnomaly	0.020	0.022	0.012	0.017	0.020	0.011	0.019	0.021	0.013	0.018	0.020	0.011	0.019	0.022	0.013
InterFusion	0.018	0.020	0.011	0.019	0.022	0.012	0.017	0.019	0.011	0.018	0.021	0.012	0.020	0.022	0.013
A.T.	0.015	0.018	0.010	0.017	0.020	0.012	0.016	0.018	0.011	0.019	0.021	0.012	0.018	0.020	0.012
MEMTO	0.025	0.028	0.015	0.021	0.025	0.013	0.023	0.027	0.015	0.026	0.030	0.016	0.022	0.026	0.014
NPSR	0.007	0.009	0.005	0.008	0.010	0.006	0.009	0.011	0.006	0.007	0.009	0.005	0.008	0.010	0.006
D^3R	0.011	0.013	0.007	0.010	0.012	0.007	0.011	0.013	0.007	0.012	0.014	0.008	0.011	0.013	0.007
SARAD	0.013	0.016	0.009	0.015	0.018	0.010	0.014	0.017	0.010	0.015	0.018	0.011	0.014	0.017	0.010
ALoRa-Det	0.009	0.011	0.006	0.010	0.012	0.007	0.009	0.011	0.006	0.010	0.012	0.007	0.009	0.011	0.006

1566 PSEUDOCODES FOR ALORA-DET AND ALORA-LOC:

1567

Algorithm 1 ALoRa-Det: Attention Low-Rank Transformer for MTS Anomaly Detection

1568

Input: $X \in \mathbb{R}^{T \times d}$: Multivariate Time Series

1569

 H : Number of attention heads, L : Number of Transformer layers.

1570

 d_{model} : Embedding dimension, λ_{reg} : ALoRa regularization coefficient.

1571

Output: Anomaly scores $\text{AS}(x_t) \in \mathbb{R}$ for all t

1572

Stage 1: LightMTS-Embed \triangleright see Section 5

1573

1: Define kernel size m (default: 3) and time series pair set $\mathcal{P} = \{(i, j) \mid 1 \leq i < j \leq d\}$

1574

2: Define sparse kernel tensor $W \in \mathbb{R}^{|\mathcal{P}| \times d \times m}$

1575

3: $\hat{X}_{:,k} = \text{Conv1D}(X; W_k) \quad \forall k \in [1, |\mathcal{P}|]$

1576

Stage 2: ALoRa-T — Attention Low-Rank Transformer

1577

4: Initialize total ALoRa loss: $\mathcal{L}_{\text{ALoRa}} \leftarrow 0$

1578

5: **for** $l = 1, \dots, L$ **do**

1579

6: $Z^{(l)} \leftarrow \text{MHA}(X^{(l-1)}) + X^{(l-1)}, X^{(0)} = \hat{X} \quad \triangleright \text{MHA + Residual connection}$

1580

7: $X^{(l)} \leftarrow Z^{(l)} \quad \triangleright$ (Optionally apply an activation function, e.g., GELU)

1581

8: Compute $\mathcal{L}_{\text{ALoRa}}(S^{(l)}) \quad \triangleright$ Eq. (7), $S^{(l)} = \frac{1}{H} \sum_{h=1}^H S_h^{(l)}$

1582

9: $\mathcal{L}_{\text{ALoRa}} \leftarrow \mathcal{L}_{\text{ALoRa}} + \mathcal{L}_{\text{ALoRa}}(S^{(l)})$

1583

10: **end for**

1584

Stage 3: Reconstruction and Training:

1585

11: Reconstruction: $\hat{X} = Z^{(L)} W^{\text{out}} \in \mathbb{R}^{T \times d}$

1586

12: Training Objective: $\mathcal{L}_{\text{total}} = \|X - \hat{X}\|_2^2 + \lambda_{\text{reg}} \cdot \mathcal{L}_{\text{ALoRa}} \quad \triangleright$ Eq. (8)

1587

Stage 4: Inference - Anomaly Detection

1588

13: Anomaly score for time t : $\text{AS}(x_t) = \|x_t - \hat{x}_t\|_2^2 \cdot \text{ALoRa-T}(x_t) \quad \triangleright$ Eq. 9

1589

14: **return** $\text{AS}(x_t)$

1590

1591

1592

1593

1594

1595

Algorithm 2 ALoRa-Loc: Attention Low-Rank Transformer for MTS Anomaly Localization

1596

Stage 1: ALoRa-T Parameter Tracking During Final Epoch of Training

1597

1: During the last epoch of ALoRa-Det training:

1598

2: Track LightMTS-Embed convolution kernels $W_k \in \mathbb{R}^{d \times m}$ for all $k = 1, \dots, d_{\text{model}}$

1599

3: **for** each Transformer layer $l = 1, \dots, L$ **do**

1600

4: Track value projection matrix $W^{(\mathcal{V},l)} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{model}}}$

1601

 \triangleright If MHA is used: then $W^{(\mathcal{V},l)} = [W^{(\mathcal{V},1,l)}, \dots, W^{(\mathcal{V},H,l)}] \cdot W^{\text{V-proj},l}$, where each $W^{(\mathcal{V},h,l)} \in \mathbb{R}^{d_{\text{model}} \times \frac{d_{\text{model}}}{H}}$ and $W^{\text{V-proj},l} \in \mathbb{R}^{\frac{d_{\text{model}}}{H} \times d_{\text{model}}}$, H is the number of heads

1602

5: **end for**

1603

6: Keep track of output projection matrix $W^{\text{out}} \in \mathbb{R}^{d_{\text{model}} \times d} \quad \triangleright$ (Eq. (6))

1604

Stage 2: Compute Input-to-Output Contribution Matrices : \triangleright (see Section 4)

1605

7: Compute contribution matrices $B \in \mathbb{R}^{d_{\text{model}} \times d_{\text{model}}}$ and $E \in \mathbb{R}^{d \times d_{\text{model}}}$ \triangleright see Section 4

1606

8: Contribution of input time series to the reconstructed ones: $C \in \mathbb{R}^{d \times d}$ \triangleright Eq. (12)

1607

Stage 3: Inference — Anomaly Localization

1608

9: ALoRa-Loc: $\text{LAS}_t^{(i)} = \sum_{j=1}^d C_{ij} \cdot \|x_t^{(j)} - \hat{x}_t^{(j)}\|_2^2$

1609

10: ALoRa-Loc (top- k): $\text{LAS}_t^{(i,topk)} = \sum_{j \in \mathcal{I}_k^{(i)}} C_{ij} \cdot \|x_t^{(j)} - \hat{x}_t^{(j)}\|_2^2$

1610

11: $\triangleright \mathcal{I}_k^{(i)}$: indices of top- k values in row i of C

1611

12: ALoRa-Loc*: $\text{LAS}_t^{*(i)} = \|x_t^{(i)} - \hat{x}_t^{(i)}\|_2^2$

1612

13: **return** $\text{LAS}_t^{(i)}$, $\text{LAS}_t^{(i,topk)}$, and $\text{LAS}_t^{*(i)}$ for all i, t

1613

1614

1615

1616

1617

1618

1619

1620 **G ANALYTICAL PROOFS**
 1621

1622 [Proof of Proposition 1] To begin the proof, we first define the Space-Time Autoregressive (STAR)
 1623 model. Let $\mathbf{y}_t \in \mathbb{R}^{1 \times d}$, for $t = 1, \dots, N$, denote the multivariate observation at time step t , where
 1624 $y_t^{(i)} \in \mathbb{R}$ represents the value of the i -th series at time t . The STAR model assumes that each time
 1625 series is influenced by its own past values as well as by the past values of spatially neighboring
 1626 series.

1627 **STAR Model Definition.** For a STAR model with $p + 1$ temporal lags, and $q + 1$ spatial lags, the
 1628 model is defined as

1629
$$\mathbf{y}_t = \sum_{k=0}^p \sum_{l=0}^q a_{tk}^{(l)} \mathbf{y}_{t-k} \mathbf{B}^{(l)} = \sum_{l=0}^q \mathbf{A}_t^{(l)} \mathbf{X} \mathbf{B}^{(l)} \quad (18)$$

 1630
 1631

1632 where $\mathbf{B}^{(l)} \in \mathbb{R}^{d \times d}$ is the spatial weight matrix of order l , and $B_{ij}^{(l)}$ represents the spatial
 1633 dependency of time series i on time series j in the l -th spatial lag. The coefficient $a_{tk}^{(l)}$ denotes the
 1634 temporal weight for lag k at time t , corresponding to the l -th spatial lag. The final equality follows
 1635 by defining the matrix $\mathbf{X} = [\mathbf{y}_{t-p}, \dots, \mathbf{y}_t]^\top \in \mathbb{R}^{p \times d}$. In the case where the spatial dependence can
 1636 be represented as a single lag (i.e., the spatial lags are collapsed into one), the model simplifies to:
 1637

1638 **STAR model with $p + 1$ temporal lags and one spatial lag:**
 1639

1640
$$\mathbf{y}_t = \sum_{k=0}^p a_{tk} \mathbf{y}_{t-k} \mathbf{B} = \mathbf{A}_t \mathbf{X} \mathbf{B}, \quad (19)$$

 1641
 1642

1643 where $\mathbf{A}_t \in \mathbb{R}^{1 \times p}$ is the vector of temporal weights used for modeling time step t . Finally, the j -th
 1644 component of \mathbf{y}_t , denoted by $y_t^{(j)}$, can be written as:

1645
$$y_t^{(j)} = \sum_{k=0}^p \sum_{i=1}^d A_{tk} B_{ij} y_{t-k}^{(i)} = \sum_{i=1}^d B_{ij} \left(\sum_{k=0}^p A_{tk} y_{t-k}^{(i)} \right) \quad (20)$$

 1646
 1647

1648 The STAR model can be fitted by minimizing the least square error between the observed data \mathbf{y}_t
 1649 and the fitted values. We note that all the learnable weights are not input-dependent.

1650 **Statement 1:** When no skip connections are used in the self-attention mechanism, the update rule
 1651 at layer l is defined as:

1652
$$\mathbf{Z}^{(l)} = \mathbf{S}^{(l)} \mathbf{Z}^{(l-1)} \mathbf{W}^{(V,l)}, \quad \text{where,} \quad \mathbf{Z}^{(0)} = \tilde{\mathbf{X}}$$

 1653
 1654

1655 By unrolling the above equation, it follows that after L layers of self-attention, the latent represen-
 1656 tation at time step t can be expressed as:

1657
$$\mathbf{Z}_t^{(L)} = \mathbf{S}_t^{(L)} \mathbf{S}^{(L-1)} \dots \mathbf{S}^{(1)} \tilde{\mathbf{X}} \mathbf{W}^{(V,1)} \dots \mathbf{W}^{(V,L)} \quad (21)$$

 1658
 1659

1660 where each $\mathbf{S}^{(i)} \in \mathbb{R}^{T \times T}$ is the SA-matrix at layer i , and $\mathbf{S}_t^{(L)} \in \mathbb{R}^{1 \times T}$ is the t 'th row (last one) of
 1661 the final SA-matrix. Each $\mathbf{W}^{(V,i)} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{model}}}$ is the value projection matrix at layer i . By defining
 1662 the product of attention matrices up to layer L at time step t as $\mathbf{A}_t = \mathbf{S}_t^{(L)} \mathbf{S}^{(L-1)} \dots \mathbf{S}^{(1)} \in \mathbb{R}^{1 \times T}$,
 1663 and the product of all value projection matrices as $\mathbf{B} = \mathbf{W}^{(V,1)} \dots \mathbf{W}^{(V,L)} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{model}}}$, the final
 1664 latent representation is compactly expressed by:

1665
$$\mathbf{Z}_t = \mathbf{A}_t \tilde{\mathbf{X}} \mathbf{B} \in \mathbb{R}^{1 \times d_{\text{model}}} \quad (22)$$

 1666

1667 Then, by expanding the above equation for each component, we can derive its analytical expression.
 1668 The expression for each **latent space** time series at time step t is given by:

1669
 1670
$$z_t^{(j)} = \sum_{k=1}^{d_{\text{model}}} b_{kj} \left(\sum_{q=1}^t a_{tq} \tilde{x}_q^{(k)} \right),$$

 1671
 1672

1673 By comparing the derived equation with Eq. (20), it is evident that the two models share the same
 1674 structure. The key difference lies in how the weights a_{tq} are obtained: traditional STAR models use

1674 fixed lag weights estimated by minimizing a loss function (e.g., mean squared error), whereas the
 1675 Transformer computes these weights dynamically through the attention mechanism, with the queries
 1676 Q and keys K determining them in real time.
 1677

Statement 2:

1678 When skip connections are applied in the self-attention mechanism, the update rule at layer l is
 1679 defined as:
 1680

$$Z^{(l)} = \tilde{Z}^{(l)} + Z^{(l-1)}, \quad \text{where } \tilde{Z} = S^{(l)} Z^{(l-1)} W^{(V,l)}. \quad (23)$$

1682 To clarify the structure of the derived Eq. (5), assume without loss of generality that the total number
 1683 of layers is $L = 3$. Then, the final representation can be written as:
 1684

$$\begin{aligned} z_t^{(3)} = & y_t + \underbrace{S_t^{(1)} \tilde{X} W^{(V,1)}}_{\mathbf{A}_t^{(1)} \quad \mathbf{B}^{(1)}} + \underbrace{S_t^{(2)} \tilde{X} W^{(V,2)}}_{\mathbf{A}_t^{(2)} \quad \mathbf{B}^{(2)}} + \underbrace{S_t^{(3)} \tilde{X} W^{(V,3)}}_{\mathbf{A}_t^{(3)} \quad \mathbf{B}^{(3)}} \\ & + \underbrace{S_t^{(2)} S_t^{(1)} \tilde{X} W^{(V,1)} W^{(V,2)}}_{\mathbf{A}_t^{(4)} \quad \mathbf{B}^{(4)}} \\ & + \underbrace{S_t^{(3)} S_t^{(1)} \tilde{X} W^{(V,1)} W^{(V,3)}}_{\mathbf{A}_t^{(5)} \quad \mathbf{B}^{(5)}} \\ & + \underbrace{S_t^{(3)} S_t^{(2)} \tilde{X} W^{(V,2)} W^{(V,3)}}_{\mathbf{A}_t^{(6)} \quad \mathbf{B}^{(6)}} \\ & + \underbrace{S_t^{(3)} S_t^{(2)} S_t^{(1)} \tilde{X} W^{(V,1)} W^{(V,2)} W^{(V,3)}}_{\mathbf{A}_t^{(7)} \quad \mathbf{B}^{(7)}}. \end{aligned} \quad (24)$$

1702 Each component of Eq. (24) has the form $\mathbf{A}_t^{(j)} \tilde{X} \mathbf{B}^{(j)}$ for $j = 0, \dots, 2^L - 1$, where $\mathbf{A}_t^{(j)} \in \mathbb{R}^{1 \times T}$ de-
 1703 notes the row vector appearing to the left of \tilde{X} , and $\mathbf{B}^{(j)} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{model}}}$ denotes the matrix product
 1704 appearing to the right of \tilde{X} in the j -th term. For example, for $j = 4$: $\mathbf{A}_t^{(4)} = S_t^{(2)} S_t^{(1)}$ and $\mathbf{B}^{(4)} =$
 1705 $W^{(V,1)} W^{(V,2)}$. As shown in the proof of Statement 1, each such component induces a structure
 1706 analogous to that of a STAR model. Thus, the overall latent representation Z_t can be interpreted as a
 1707 **linear combination of multiple STAR-like processes**, where each component has its own temporal
 1708 (a_{tj}) and spatial weights (b_{kj}). In other words, each term $\mathbf{A}_t^{(j)} \tilde{X} \mathbf{B}^{(j)}$ defines a distinct STAR-like
 1709 model with potentially different spatial and temporal lag structures.
 1710

1711 We can obtain a more specific interpretation of this representation by comparing Eq. (24) with the
 1712 general definition of a STAR model in Eq. (18). In particular, Eq. (24) corresponds to a STAR
 1713 model with $p = T$ (the window size) temporal lags and $q = 8 = \binom{L}{2}$ spatial lags, where $\mathbf{B}^{(0)} = \mathbf{I}_d$
 1714 and $\mathbf{A}_t^{(0)} = \mathbf{e}^{(T)} = [0, 0, \dots, 1]$. The explicit structure of each $\mathbf{A}_t^{(l)}$ and $\mathbf{B}^{(l)}$ is indicated by the
 1715 underbraces in Eq. (24).

Statement 3:

1718 The feed-forward layer, linearly combines values across all series at the same time step to produce
 1719 a new time series $y_t^{(j)}$. We now show that the transformed representation for each time series, after
 1720 applying the feed-forward layer, can be written in the same structure as in Eq. (22): For ease of
 1721 presentation let $T = 2$ and the latent dimension $d_{\text{model}} = 3$. The new j -th time series at time step t
 1722 is defined as: $y_t^{(j)} = w_{1j} z_t^{(1)} + w_{2j} z_t^{(2)} + w_{3j} z_t^{(3)}$.
 1723

1724 Expanding each $z_t^{(i)}$:

$$y_t^{(j)} = w_{1j} \left[b_{11} \left(a_{t,t-1} \tilde{x}_{t-1}^{(1)} + a_{tt} \tilde{x}_t^{(1)} \right) + b_{21} \left(a_{t,t-1} \tilde{x}_{t-1}^{(2)} + a_{tt} \tilde{x}_t^{(2)} \right) + b_{31} \left(a_{t,t-1} \tilde{x}_{t-1}^{(3)} + a_{tt} \tilde{x}_t^{(3)} \right) \right]$$

$$\begin{aligned}
& + w_{2j} \left[b_{12} \left(a_{t,t-1} \tilde{x}_{t-1}^{(1)} + a_{tt} \tilde{x}_t^{(1)} \right) + b_{22} \left(a_{t,t-1} \tilde{x}_{t-1}^{(2)} + a_{tt} \tilde{x}_t^{(2)} \right) + b_{32} \left(a_{t,t-1} \tilde{x}_{t-1}^{(3)} + a_{tt} \tilde{x}_t^{(3)} \right) \right] \\
& + w_{3j} \left[b_{13} \left(a_{t,t-1} \tilde{x}_{t-1}^{(1)} + a_{tt} \tilde{x}_t^{(1)} \right) + b_{23} \left(a_{t,t-1} \tilde{x}_{t-1}^{(2)} + a_{tt} \tilde{x}_t^{(2)} \right) + b_{33} \left(a_{t,t-1} \tilde{x}_{t-1}^{(3)} + a_{tt} \tilde{x}_t^{(3)} \right) \right]
\end{aligned}$$

By regrouping the terms,

$$\begin{aligned}
& \left(a_{t,t-1} \tilde{x}_{t-1}^{(1)} + a_{tt} \tilde{x}_t^{(1)} \right) \underbrace{\left(w_{1j} b_{11} + w_{2j} b_{12} + w_{3j} b_{13} \right)}_{\tilde{b}_{1j}} \\
& + \left(a_{t,t-1} \tilde{x}_{t-1}^{(2)} + a_{tt} \tilde{x}_t^{(2)} \right) \underbrace{\left(w_{1j} b_{21} + w_{2j} b_{22} + w_{3j} b_{23} \right)}_{\tilde{b}_{2j}} \\
& + \left(a_{t,t-1} \tilde{x}_{t-1}^{(3)} + a_{tt} \tilde{x}_t^{(3)} \right) \underbrace{\left(w_{1j} b_{31} + w_{2j} b_{32} + w_{3j} b_{33} \right)}_{\tilde{b}_{3j}}
\end{aligned}$$

so for general latent space dimension d_{model} and for general window size T :

$$y_t^{(j)} = \sum_{k=1}^{d_{\text{model}}} \left[\tilde{b}_{kj} \left(\sum_{q=1}^t a_{tq} \tilde{x}_q^{(k)} \right) \right], \quad \text{where } \tilde{b}_{kj} = \sum_{r=1}^{d_{\text{model}}} w_{rj} b_{kr}. \quad (25)$$

Therefore, from the first statement of the proposition, it follows that the structure of the latent space time series remains unchanged, that is, it retains a STAR-like form. The only difference is that the new weights \tilde{b}_{kj} now determine the contribution of each input time series, replacing the original weights b_{kj} .

G.1 DERIVATION OF CONTRIBUTION WEIGHTS FROM INPUT SPACE TO LATENT AND RECONSTRUCTION SPACES (RELATED TO SECTION 5.2)

The latent representation at the final layer is given by Eq. (5). In Eq. (24), we provide an analytical expression for the case of a three-layer model. Each component of Eq. (24) has the form $A_t^{(k)} \tilde{\mathbf{X}} \mathbf{B}^{(k)}$ for $k = 0, \dots, 2^L - 1$, where $A_t^{(k)} \in \mathbb{R}^{1 \times T}$ denotes the row vector appearing to the left of $\tilde{\mathbf{X}}$, and $\mathbf{B}^{(k)} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{model}}}$ denotes the matrix product appearing to the right of $\tilde{\mathbf{X}}$ in the k -th term. For example, for $k = 4$: $A_t^{(4)} = \mathbf{S}_t^{(2)} \mathbf{S}_t^{(1)}$ and $\mathbf{B}^{(4)} = \mathbf{W}^{(V,1)} \mathbf{W}^{(V,2)}$. Each of these components contributes to the overall contribution weights we aim to derive. As a first step, we consider how each embedding-space time series contributes to the representation of the j -th time series (column) in the k -th component of Eq. (24) ($A_t^{(k)} \tilde{\mathbf{X}} \mathbf{B}^{(k)}$). The j -th time series (column) in the k -th component, can be expressed as:

$$z_t^{(j,k)} = \sum_{r=1}^{d_{\text{model}}} b_{rj}^{(k)} \left(\sum_{q=1}^t a_{tq}^{(k)} \tilde{x}_q^{(r)} \right),$$

where $b_{rj}^{(k)}$ denotes the (r, j) -th entry of $\mathbf{B}^{(k)}$ and measures how strongly the r -th embedding-space time series contributes to the j -th latent series in the k -th component. The coefficient $a_{tq}^{(k)}$ is the (t, q) -th entry of $A_t^{(k)}$ and serves as a temporal weighting factor. Importantly, the dependence on $a_{tq}^{(k)}$ does not alter the interpretation of $b_{rj}^{(k)}$. To illustrate this, we consider the case $T = 2$ and $d_{\text{model}} = 3$. In this setting, the three latent space time series are given by:

$$\begin{aligned}
z_t^{(1,k)} &= b_{11}^{(k)} (a_{t,t-1}^{(k)} \tilde{x}_{t-1}^{(1)} + a_{t,t}^{(k)} \tilde{x}_t^{(1)}) + b_{21}^{(k)} (a_{t,t-1}^{(k)} \tilde{x}_{t-1}^{(2)} + a_{t,t}^{(k)} \tilde{x}_t^{(2)}) + b_{31}^{(k)} (a_{t,t-1}^{(k)} \tilde{x}_{t-1}^{(3)} + a_{t,t}^{(k)} \tilde{x}_t^{(3)}) \\
z_t^{(2,k)} &= b_{12}^{(k)} (a_{t,t-1}^{(k)} \tilde{x}_{t-1}^{(1)} + a_{t,t}^{(k)} \tilde{x}_t^{(1)}) + b_{22}^{(k)} (a_{t,t-1}^{(k)} \tilde{x}_{t-1}^{(2)} + a_{t,t}^{(k)} \tilde{x}_t^{(2)}) + b_{32}^{(k)} (a_{t,t-1}^{(k)} \tilde{x}_{t-1}^{(3)} + a_{t,t}^{(k)} \tilde{x}_t^{(3)})
\end{aligned}$$

1782 $z_t^{(3,k)} = b_{13}^{(k)}(a_{t,t-1}^{(k)}\tilde{x}_{t-1}^{(1)} + a_{t,t}^{(k)}\tilde{x}_t^{(1)}) + b_{23}^{(k)}(a_{t,t-1}^{(k)}\tilde{x}_{t-1}^{(2)} + a_{t,t}^{(k)}\tilde{x}_t^{(2)}) + b_{33}^{(k)}(a_{t,t-1}^{(k)}\tilde{x}_{t-1}^{(3)} + a_{t,t}^{(k)}\tilde{x}_t^{(3)})$
 1783

1784 In all three equations above: The temporal attention weights $a_{t,t-1}^{(k)}$ and $a_{t,t}^{(k)}$ are shared across all
 1785 parentheses in each equation and are also shared across all equations, as they remain the same
 1786 for $z_t^{(1,k)}, z_t^{(2,k)}, z_t^{(3,k)}$. What differentiates how much each embedding space time series ($\tilde{x}^{(i)}$)
 1787 contributes to each latent space series ($z^{(j,k)}$) is only the contribution weight $b_{ij}^{(k)}$.
 1788

1789 Then, by summing the contribution weights over all components of Eq. (24), we obtain the matrix
 1790 that characterizes how embedding-space time series contribute to the latent space, defined as
 1791

$$\mathbf{B} = \prod_{i=1}^L (\mathbf{W}^{(V,i)} + I).$$

1794 Finally, by accounting for the embedding module, the contribution weights from the raw input time
 1795 series to the latent space time series are given by:
 1796

$$E_{ij} = \sum_{k=1}^{d_{\text{model}}} \left(\sum_{l=-\frac{m-1}{2}}^{\frac{m-1}{2}} w_{i,l}^{(k)} \right) b_{kj} \quad (26)$$

1802 Then, the reconstruction of each time series is produced as follows: $\widehat{x_t^{(k)}} = \sum_{j=1}^{d_{\text{model}}} w_{jk}^{\text{out}} z_t^{(L,j)}$
 1803

1804 This explains why the final derived contribution weights of the input time series to the reconstructed
 1805 time series are given by:
 1806

$$C_{ij} = \sum_{k=1}^{d_{\text{model}}} w_{kj}^{\text{out}} \cdot b_{ik}$$

1810
 1811
 1812
 1813
 1814
 1815
 1816
 1817
 1818
 1819
 1820
 1821
 1822
 1823
 1824
 1825
 1826
 1827
 1828
 1829
 1830
 1831
 1832
 1833
 1834
 1835