



**UNIVERSITY
OF TURKU**

This is a self-archived – parallel published version of an original article. This version may differ from the original in pagination and typographic details. When using please cite the original.

This is the peer reviewed version of the following article:

CITATION: J. Chen, X. Tan, S. Rahardja, J. Yang and S. Rahardja, "Joint Selective State Space Model and Detrending for Robust Time Series Anomaly Detection," in IEEE Signal Processing Letters, vol. 31, pp. 2050-2054, 2024, doi: 10.1109/LSP.2024.3438078.

which has been published in final form at

DOI: <http://dx.doi.org/10.1109%2FLSP.2024.3438078>

© 2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works

Joint Selective State Space Model and Detrending for Robust Time Series Anomaly Detection

Junqi Chen, Xu Tan, Sylwan Rahardja, *Graduate Student Member, IEEE*,
 Jiawei Yang, *Member, IEEE*, and Susanto Rahardja, *Fellow, IEEE*

Abstract—Deep learning-based sequence models are extensively employed in Time Series Anomaly Detection (TSAD) tasks due to their effective sequential modeling capabilities. However, the ability of TSAD is limited by two key challenges: (i) the ability to model long-range dependency and (ii) the generalization issue in the presence of non-stationary data. To tackle these challenges, an anomaly detector that leverages the selective state space model known for its proficiency in capturing long-term dependencies across various domains is proposed. Additionally, a multi-stage detrending mechanism is introduced to mitigate the prominent trend component in non-stationary data to address the generalization issue. Extensive experiments conducted on real-world public datasets demonstrate that the proposed methods surpass all 12 compared baseline methods.

Index Terms—Time series anomaly detection, Selective state space model, Time series detrending

I. INTRODUCTION

As the volume of data generated continues to grow exponentially, Time Series Anomaly Detection (TSAD) has garnered significant attention due to its increasing demand in various real-world applications such as intrusion detection, disaster warning, and medical diagnosis [1]–[5]. TSAD aims to identify irregular points or subsequences collectively referred to as *anomalies*. Typically, samples with high reconstruction errors of deep neural networks (DNNs) trained exclusively on normal data are detected as anomalies. Given the temporal correlations inherent in time series [6], DNN-based sequence models are considered the most suitable approach for TSAD tasks [7]–[14]. However, two key challenges remain in these sequence model-based TSAD: (i) the ability to model long-range dependencies and (ii) the generalization issue for non-stationary data.

Normal behavior in time-series data typically involves long-term dependencies [1]. To capture these dependencies and enhance the modeling of normal behavior, various sequence models including Temporal Convolutional Networks (TCNs) [15], Recurrent Neural Networks (RNNs) [8], [9] and Transformers [11], [12] have been explored for TSAD. However,

existing methods still encounter difficulties due to their intrinsic characteristics, such as limited context window, high memory costs, or unstable gradient flow. Recent research on the selective state space model (S6) [16] has addressed the shortcomings of the aforementioned sequence models and demonstrated its excellent long-term modeling capabilities in other domains. Moreover, the selective nature of the S6 model allows it to discard abnormal information and generate reliable reconstructed output for TSAD. Despite offering a potential solution to model long-range dependencies, there has been no research exploring its application in TSAD.

The generalization issue for non-stationary data stems from trends in non-stationary data. The erratic distribution of trends can result in significant fluctuations in data magnitude, leading to erroneously high reconstruction errors in regions with previously unseen trend patterns in the training set, ultimately causing false alarms [14]. Traditional time series decomposition methods such as Seasonal-Trend decomposition using LOESS (STL) [17] and Hodrick–Prescott (HP) trend filter [18] are frequently used to mitigate the impact of these trends. However, achieving optimal decomposed results may not yield the best detection performance. To address this, several approaches have tried to integrate decomposition methods into detectors for joint optimization. Nevertheless, most of them rely on Moving Average (MA) with a fixed kernel size [14], [19], [20], lacking the flexibility for broader scenarios. Other methods rely on DNN models for detrending [21], [22]. Though they allow end-to-end optimization, their dependence on training data reintroduces generalization issues.

To tackle the challenges outlined above, an innovative detector constructed using the S6 model and integrated with a multi-stage detrending mechanism is proposed. The contributions of this paper is as follows: (i) Introduce a novel detector that utilizes the S6 model to effectively model long-range dependencies in TSAD. (ii) Propose a multi-stage detrending mechanism that can generate reliable decomposed results. (iii) Evaluate the proposed method on three benchmark datasets and demonstrate the superiority of the proposed method over recent State-of-the-art (SOTA) methods.

II. RELATED WORK

Sequence models such as TCN, RNN, and Transformer [23]–[25] had greatly improved TSAD tasks. However, these methods struggled with long-range dependency modeling, limiting detection performance. Fig. 1 illustrates how different sequence models gather context information. As shown

Junqi Chen, Xu Tan and Susanto Rahardja are with the School of Marine Science and Technology, Northwestern Polytechnical University, Xi'an, 710072, China. (e-mail: jqchen@ieee.org; xutan@ieee.org; susantorahardja@ieee.org).

Sylwan Rahardja is with the School of Computing, University of Eastern Finland, FI-80101 Joensuu, Finland. (e-mail: sylwanrahardja@ieee.org).

Jiawei Yang is with the Department of Computing University of Turku, 20014, Turku, Finland (e-mail: jiaweiyang@ieee.org).

Susanto Rahardja is also with the Engineering Cluster, Singapore Institute of Technology, 10 Dover Drive, Singapore 138683.

(Corresponding author: Susanto Rahardja.)

in Fig. 1(a), TCN's context was constrained by the kernel size k , which made it unsuitable for capturing long-term dependencies. Conversely, the Transformer could capture long context using the Self-Attention (SA) mechanism as shown in Fig. 1(b). However, SA was memory and time-intensive during training and inference because it did not compress context information, which made it impractical for modeling long dependencies. Although RNN could efficiently compress context information using a finite state shown in Fig. 1(c), it suffered from unstable gradient issues with long sequences [26]. The State Space Model (SSM), similar to RNN, resolves gradient issues by employing solely linear transformations, as shown in Fig. 1(d).

The Structured State Space Sequence model (S4) [27] was a well-known SSM that combines the benefits of RNN with the stability of linear transformations. This allows S4 to compress context states into a finite size while ensuring stable gradients, making it suitable for modeling long-range dependencies. However, its parameters remain constant through time, forming a Linear Time Invariance (LTI) model, limiting its effectiveness in TSAD tasks. To remove the LTI constraint, Gu et al. [16] proposed a selective SSM called S6 by introducing input-dependent parameters. S6 exhibited superior performance in modeling long-range dependencies in other domains, including natural language processing [16], computer vision [28], and speech [29]. Additionally, its selection mechanism has the potential to enhance TSAD performance by adaptively selecting the appropriate context information for producing reliable anomaly scores. Despite its effectiveness across various domains, its application in TSAD remains to be unexplored.




Fig. 1. Illustration of different sequence models gathering context information, where k denotes the kernel size, t denotes the time-step, $\mathbf{x}_{1:t-1}$ represents the context information, \mathbf{x}_t and \mathbf{y}_t represent the input and output in current time-step, respectively. (a) TCN's context was constrained by the kernel size k . (b) SA in the Transformer could capture long-term context but lacked information compression. (c-d) RNN and SSM could compress context information by a finite state.

III. PRELIMINARIES

A. Problem Statement

Given a multivariate time series $\mathbf{x} \in \mathbb{R}^{T \times D}$ with length T for training, where each observation $\mathbf{x}_t \in \mathbb{R}^D$ contains D features. TSAD aims to predict anomalous labels $\mathbf{y} = [y_1, \dots, y_T]^T$ for unseen test time series $\hat{\mathbf{x}}$ of length \hat{T} , where $y_t \in \{0, 1\}$. This prediction is based on anomaly scores $\mathbf{a} = [a_1, \dots, a_{\hat{T}}]^T$ generated by the detector and a predefined threshold a_{th} , where $y_t = 1$ if $a_t > a_{th}$, otherwise 0.

B. S6 Model

A typical SSM model maps a univariate sequence $x_t \in \mathbb{R}$ to the output $y_t \in \mathbb{R}$ through a hidden state $\mathbf{h}_t \in \mathbb{R}^N$ with

four parameters $\Delta \in \mathbb{R}$, $\mathbf{A}, \mathbf{B}, \mathbf{C} \in \mathbb{R}^N$, where t denotes the t -th time-step and N is the number of states.

To remove the LTI constraint, three of the above parameters in S6 model were adjusted to vary with time, namely $\Delta \in \mathbb{R}^T$, $\mathbf{B}, \mathbf{C} \in \mathbb{R}^{T \times N}$, where T denotes the number of time-steps. The operation of the S6 model could be represented by a general formula applicable to multivariate time series data $\mathbf{x}, \mathbf{y} \in \mathbb{R}^{T \times D}$:

$$\begin{aligned} \mathbf{h}_{t,d} &= \bar{\mathbf{A}}_{t,d} \odot \mathbf{h}_{t-1,d} + \bar{\mathbf{B}}_{t,d} \cdot \mathbf{x}_{t,d}, \\ \mathbf{y}_{t,d} &= \mathbf{C}_t^T \cdot \mathbf{h}_{t,d}, \end{aligned} \quad (1)$$

where d denotes the d -th feature, \odot represents the Hadamard product, $\bar{\mathbf{A}}_{t,d} = f_A(\Delta_{t,d}, \mathbf{A}_d)$ and $\bar{\mathbf{B}}_{t,d} = f_B(\Delta_{t,d}, \mathbf{B}_t)$. Here, $f_A(\cdot)$ and $f_B(\cdot)$ represent discretization rules for \mathbf{A} and \mathbf{B} . It is suggested that $f_A(\cdot)$ followed the zero-order hold rule, while $f_B(\cdot)$ follows the Euler rule [16], expressed as:

$$\begin{aligned} f_A(\Delta_{t,d}, \mathbf{A}_d) &= \exp(\Delta_{t,d} \cdot \mathbf{A}_d), \\ f_B(\Delta_{t,d}, \mathbf{B}_t) &= \Delta_{t,d} \cdot \mathbf{B}_t. \end{aligned} \quad (2)$$

C. Time Series Detrending

A common time series model with trend and seasonality could be formulated as $\mathbf{x}_t = \boldsymbol{\tau}_t + \mathbf{s}_t + \mathbf{r}_t$, $t = 1, \dots, T$, where \mathbf{x}_t denotes the observation at the t -th time-step, $\boldsymbol{\tau}_t$ is the trend component, \mathbf{s}_t is the seasonal component with a period of k , and \mathbf{r}_t represents the residual component. In this paper, emphasis was placed on trend removal and HP trend filter was adopted as the detrending method, expressed as:

$$\hat{\boldsymbol{\tau}}_t = \arg \min_{\boldsymbol{\tau}_t} \sum_{t=1}^T (\mathbf{x}_t - \boldsymbol{\tau}_t)^2 + \lambda \sum_{t=2}^{T-1} (\boldsymbol{\tau}_{t+1} - 2\boldsymbol{\tau}_t + \boldsymbol{\tau}_{t-1})^2, \quad (3)$$

where λ is a hyper-parameter controlling the smoothness.

IV. METHOD

Fig. 2 illustrates the overall framework of the proposed method, comprising a HP trend filter, multiple Decomposition-based Mamba (DMamba) blocks, and a final output module. Initially, input data was processed by the trend filter to extract initial trends and seasonality components. The initial seasonality was refined through DMamba blocks, while the initial trend was merged with trends from DMamba blocks. Finally, the fused trend and refined seasonality were summed to form the reconstructed output.




Fig. 2. The structure of the proposed method includes a HP trend filter, multiple DMamba blocks, and an output module.

A. HP Trend Filter

The input data was initially segmented into subsequences using sliding windows of length W . Each input window $\mathbf{x}^{\text{in}} \in \mathbb{R}^{W \times D}$ underwent detrending using HP trend filter. To ensure the extraction of global trends, the current input window was first combined with historical windows, and then the HP-based trend component τ_{HP} was computed via Eq. (3). The HP-based seasonality was further derived as $\mathbf{s}_{\text{HP}} = \mathbf{x}^{\text{in}} - \tau_{\text{HP}}$.

B. DMamba Blocks




Fig. 3. The structure of Mamba block.

The HP-based seasonality \mathbf{s}_{HP} was first projected to a D_m -dimensional space by an embedding network $\text{Emb}(\cdot)$, and then it underwent L DMamba blocks. Each block comprised a Mamba block [16] and an Adaptive MA (AMA) module. The input of each block was the output of the previous block, denoted as $\mathbf{x}^l = \mathbf{y}^{l-1}$ with $\mathbf{y}^0 = \text{Emb}(\mathbf{s}_{\text{HP}})$. For clarity, the superscript l was omitted in the following content. The input \mathbf{x} first entered the Mamba block shown in Fig. 3 and could be formulated by Eq. (4):

$$\begin{aligned} \mathbf{u}, \mathbf{g} &= \sigma(\text{Conv}(\mathbf{x} \cdot \mathbf{W}_2)), \sigma(\mathbf{x} \cdot \mathbf{W}_1), \\ \mathbf{z} &= \text{S6}(\mathbf{u}), \\ \mathbf{x}_{\text{SSM}} &= (\mathbf{g} \odot \mathbf{z}) \cdot \mathbf{W}_3, \end{aligned} \quad (4)$$

where $\mathbf{W}_1, \mathbf{W}_2 \in \mathbb{R}^{D_m \times 2D_m}$, $\mathbf{W}_3 \in \mathbb{R}^{2D_m \times D_m}$ are trainable projection matrices, $\text{Conv}(\cdot)$ represents a convolution layer with a kernel size of 4, σ denotes the SiLU [30] activation function, and $\text{S6}(\cdot)$ denotes the S6 operation detailed in algorithm 1. The three parameters $\Delta, \mathbf{B}, \mathbf{C}$ in Algorithm 1 could be computed as:

$$\Delta = \text{Softplus}(\mathbf{u} \cdot \mathbf{W}_\Delta), \quad (5)$$

$$\mathbf{B} = \mathbf{u} \cdot \mathbf{W}_B, \quad (6)$$

$$\mathbf{C} = \mathbf{u} \cdot \mathbf{W}_C, \quad (7)$$

where $\text{Softplus}(\cdot)$ denotes the Softplus activation function, $\mathbf{W}_B, \mathbf{W}_C \in \mathbb{R}^{2D_m \times N}$ and $\mathbf{W}_\Delta \in \mathbb{R}^{2D_m \times 2D_m}$ are trainable projection matrices.

Algorithm 1 S6 operation

Input: $\mathbf{u} : (W, D)$ with W time-steps and D features ;

Output: $\mathbf{x}_{\text{SSM}} : (W, D)$;

- 1: $\mathbf{A} : (D, N) \leftarrow$ Parameter with N states
 - 2: $\Delta : (W, D) \leftarrow f_\Delta(\mathbf{u})$ using Eq. (5)
 - 3: $\mathbf{B} : (W, N) \leftarrow f_B(\mathbf{u})$ using Eq. (6)
 - 4: $\mathbf{C} : (W, N) \leftarrow f_C(\mathbf{u})$ using Eq. (7)
 - 5: $\bar{\mathbf{A}}, \bar{\mathbf{B}} : (W, D, N) \leftarrow \text{discretize}(\Delta, \mathbf{A}, \mathbf{B})$ using Eq. (2)
 - 6: $\mathbf{y} \leftarrow \text{SSM}(\bar{\mathbf{A}}, \bar{\mathbf{B}}, \mathbf{C})(\mathbf{u})$ using Eq. (1)
 - 7: **return** \mathbf{x}_{SSM} .
-

The AMA module then took \mathbf{x}_{SSM} as input and performed trend removal through average pooling:

$$\tau_{\text{MA}} = \text{AvgPool}(\mathbf{x}_{\text{SSM}}, k) \quad (8)$$

with a kernel size of k . The kernel size is adaptively estimated by $k = \lfloor W/n \rfloor$, where $\lfloor \cdot \rfloor$ denotes the floor function, $n = \arg \max(\mathbf{x}_f^H \odot \mathbf{x}_f)$ indicates the peak position of the power density function and \mathbf{x}_f is the Fourier transformation of \mathbf{x}_{SSM} . The output of the current block is then obtained by $\mathbf{y} = \mathbf{x}_{\text{SSM}} - \tau_{\text{MA}}$.

C. Final Reconstructed Output

The output of the final DMamba block served as the reconstructed seasonal component i.e., $\mathbf{s} = \mathbf{y}^L \cdot \mathbf{W}_s$, where $\mathbf{W}_s \in \mathbb{R}^{D_m \times D}$ is a trainable projection matrix. The final trend is a combination of all extracted trends:

$$\tau = \tau_{\text{HP}} + \text{Conv}\left(\sum_{l=1}^L \tau_{\text{MA}}^l\right), \quad (9)$$

where $\text{Conv}(\cdot)$ is a CNN layer with an input channel of D_m , an output channel of D , and a kernel size of 3. In the training phase, the objective was to minimize the Mean Square Error (MSE) between the original input and the reconstructed output:

$$\mathcal{L} = \|\mathbf{x}^{\text{in}} - (\mathbf{s} + \tau)\|_2^2. \quad (10)$$

During inference, the MSE serves as the anomaly score.

V. EXPERIMENTS

A. Setup

TABLE I
DATASETS DESCRIPTION.

Dataset	Train	Test	Dimensions	Contamination (%)
NASA	6,329	18,743	25/55	9.52
SMD	128,267	128,270	38	7.11
SWaT	6,840	7,500	25	12.63

1) *Datasets*: Experiments were conducted on three publicly available datasets. The **NASA** dataset [8] contained sensor metrics recorded by Soil Moisture Active Passive satellite (SMAP) and Mars Science Laboratory rover (MSL). Three non-trivial subsets were utilized, namely *A-4*, *T-1*, and *C-2*. as suggested in [12]. The **Server Machine Dataset (SMD)** [9] includes multiple server machine metrics from a large Internet company. Emphasis was placed on non-trivial subsets *1-1*, *1-6*, *2-1*, *3-2*, and *3-7*, as suggested [12]. Lastly, the **Secure Water Treatment (SWaT)** [38] was a dataset that spans 11 days of continuous operation. In the experiments, SWaT was downsampled to the minute level and binary features were discarded. Detailed information is listed in Tab. I.

2) *Baselines*: The proposed method was compared to traditional and DNN-based methods. Traditional methods include LOF [31], CBLOF [32], OCSVM [33], IForest [34], HBOS [35], LODA [36] and ECOD [37]. DNN-based methods comprised two RNN-based methods: Omnianomaly [9] and FGANomaly [10], two Transformer-based methods: AnomalyTrans [11] and DCdetector [13], and a time series decomposition-based method: D3R [14].

TABLE II
COMPARISON RESULTS ON THREE DATASETS. THE BEST RESULTS WERE **BOLDED** WHILE THE SECOND BEST WERE UNDERLINED.

Method	NASA			SWaT			SMD			Avg.		
	P-AF	R-AF	F1-AF	P-AF	R-AF	F1-AF	P-AF	R-AF	F1-AF	P-AF	R-AF	F1-AF
LOF [31]	0.4020	0.6313	0.4886	0.8450	0.2855	0.4268	0.7002	0.7376	0.6959	0.6491	0.5515	0.5371
CBLOF [32]	0.4078	0.6036	0.4818	0.6861	0.0510	0.0949	0.8227	0.7131	0.6698	0.6389	0.4559	0.4155
OCSVM [33]	0.4735	0.4739	0.4341	0.4558	0.4316	0.4434	0.7923	0.5003	0.4900	0.5739	0.4686	0.4558
IForest [34]	0.3811	0.6353	0.4744	0.9977	0.1424	0.2492	0.7905	0.6589	0.6359	0.7231	0.4789	0.4532
HBOS [35]	0.5014	0.9926	0.6661	0.9835	0.1476	0.2567	0.7848	0.5329	0.5425	0.7566	0.5577	0.4884
LODA [36]	0.4736	0.9551	0.6326	0.9777	0.0263	0.0513	0.7379	0.6064	0.5674	0.7297	0.5293	0.4171
ECOD [37]	0.3806	0.6238	0.4720	0.9940	0.1122	0.2016	0.6358	0.5551	0.5349	0.6701	0.4304	0.4028
Omnianomaly [9]	0.5074	0.9999	0.6732	0.5085	0.9738	0.6681	0.4957	0.9221	0.6438	0.5039	0.9653	0.6617
FGANomaly [10]	0.5101	1.000	<u>0.6756</u>	0.5338	0.9639	0.6871	0.6283	0.9937	0.7659	0.5574	0.9859	<u>0.7095</u>
Anomalytrans [11]	0.5024	0.9792	0.6639	0.5965	0.5929	0.5947	0.5398	0.9415	0.6847	0.5462	0.8379	0.6478
DCdetector [13]	0.4848	0.9777	0.6471	0.4753	0.5175	0.4955	0.5019	0.5367	0.6412	0.4873	0.6773	0.5946
D3R [14]	0.4334	0.9776	0.5868	0.6241	0.7970	<u>0.7000</u>	0.7506	0.9328	<u>0.8295</u>	0.6027	0.9025	0.7054
Proposed method	0.5282	0.9766	0.6829	0.6441	0.8641	0.7380	0.8061	0.9040	0.8403	0.6595	0.9149	0.7537

3) *Implementation details:* All baselines were implemented based on PyOD [39] or their publicly available codes. We utilized Peak-Over-Threshold (POT) [40] for threshold selection. The smoothness parameter λ was set to 10^4 . The block size L was 3, the model size D_m was 32, the state size was 16 and the window size W was 100. The training was conducted for 7 epochs with a batch size of 128, using the AdamW [41] optimizer with a learning rate of 10^{-4} . Given that point-wise metrics can be influenced by lengthy anomalies and point-adjustment strategy [9] could lead to misguided rankings [42], the recently proposed affiliation-based precision (P-AF), recall (R-AF), and F1-score (F1-AF) [43] were adopted as evaluation metrics. These metrics offer evaluations at the level of abnormal events rather than individual points, making them suitable for TSAD. Since F1-AF represents the harmonic mean of P-AF and R-AF, it was employed to evaluate the overall performance, a common practice in the TSAD community.

B. Results

The proposed method was compared to 12 baselines, as listed in Tab. II. Generally, sequence model-based methods outperformed traditional methods due to their ability to model temporal dependencies. The proposed method surpassed all compared methods on the three datasets, achieving a 6.2% relative improvement in average F1-AF compared to the best baseline method. This demonstrated that the proposed S6 model-based detector was suitable for TSAD tasks and superior to previous SOTA detectors with RNN or Transformer backbones. Meanwhile, for the SWaT and SMD datasets containing significant trend variation, D3R and the proposed method achieved notable performance improvement by incorporating time series decomposition approaches compared to other methods. Overall, the experimental results showcased the robustness of the proposed method in non-stationary scenarios and its ability to detect complex time series anomalies.

To further verify the necessity of the HP trend filter and the AMA mechanism, an ablation analysis was conducted and the results were shown in Tab. III. The results demonstrated that applying only one might not bring improvement. For instance, solely applying the HP trend filter in the NASA dataset or only utilizing AMA in the SWaT dataset led to a

TABLE III
ABLATION STUDY OF THE HP TREND FILTER AND AMA. THE BEST RESULTS WERE **BOLDED** WHILE THE SECOND BEST WERE UNDERLINED.

Modules		F1-AF				
HP Trend Filter	AMA	NASA	SWaT	SMD	Avg.	
✗	✗	0.6717	0.6968	0.7780	0.7155	
✗	✓	<u>0.6777</u>	0.6662	0.7782	0.7074	
✓	✗	0.6700	0.7302	<u>0.8110</u>	0.7371	
✓	✓	0.6829	0.7380	0.8403	0.7537	

performance drop. However, when both were applied, the HP trend filter provided stable input for AMA’s period estimation, where AMA and the detector could symbiotically optimize performance and capture more complex patterns, resulting in the best performance over all datasets. Additionally, a visualization in Fig. 4 clearly showcased the effectiveness of the proposed detrending mechanisms.




Fig. 4. A visualization from the SWaT dataset, which contains non-stationary data. The first row shows the input signal and extracted trend. The second row displays anomaly scores from various detectors. Baseline methods were influenced by trends, leading to false alarms, while our proposed method generated reliable scores by incorporating detrending mechanisms.

VI. CONCLUSION

This paper proposed solutions to effectively tackle two critical challenges in TSAD: modeling long-range dependency and generalizing to non-stationary data, which are bottlenecks for sequence models in TSAD. Our solutions involve leveraging a recently published S6 model to capture long-term context and introducing a novel multi-stage detrending mechanism to provide stable input for the sequence model. Experimental results underscored the suitability of the S6 model for TSAD and highlighted the importance of the proposed detrending mechanism. Furthermore, this study laid the foundation for developing more advanced S6 model-based detectors for TSAD.

REFERENCES

- [1] Z. Z. Darban, G. I. Webb, S. Pan, C. C. Aggarwal, and M. Salehi, "Deep learning for time series anomaly detection: A survey," *arXiv preprint arXiv:2211.05244*, 2022.
- [2] J. Yang, G. I. Choudhary, S. Rahardja, and P. Fränti, "Classification of interbeat interval time-series using attention entropy," *IEEE Transactions on Affective Computing*, vol. 14, no. 1, pp. 321–330, 2020.
- [3] J. Yang, X. Tan, and S. Rahardja, "Mipo: How to detect trajectory outliers with tabular outlier detectors," *Remote sensing*, vol. 14, no. 21, p. 5394, 2022.
- [4] N. LaRosa, J. Farber, P. Venkatasubramaniam, R. Blum, and A. Al Rashdan, "Separating sensor anomalies from process anomalies in data-driven anomaly detection," *IEEE Signal Processing Letters*, vol. 29, pp. 1704–1708, 2022.
- [5] Z. Wang, Y. Zhang, G. Wang, and P. Xie, "Main-auxiliary aggregation strategy for video anomaly detection," *IEEE Signal Processing Letters*, vol. 28, pp. 1794–1798, 2021.
- [6] K.-H. Lai, D. Zha, J. Xu, Y. Zhao, G. Wang, and X. Hu, "Revisiting time series outlier detection: Definitions and benchmarks," in *Thirty-fifth conference on neural information processing systems datasets and benchmarks track (round 1)*, 2021.
- [7] K. Cheng, Y. Liu, and X. Zeng, "Learning graph enhanced spatial-temporal coherence for video anomaly detection," *IEEE Signal Processing Letters*, vol. 30, pp. 314–318, 2023.
- [8] K. Hundman, V. Constantinou, C. Laporte, I. Colwell, and T. Soderstrom, "Detecting spacecraft anomalies using lstms and nonparametric dynamic thresholding," in *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 2018, pp. 387–395.
- [9] Y. Su, Y. Zhao, C. Niu, R. Liu, W. Sun, and D. Pei, "Robust anomaly detection for multivariate time series through stochastic recurrent neural network," in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2019, pp. 2828–2837.
- [10] B. Du, X. Sun, J. Ye, K. Cheng, J. Wang, and L. Sun, "Gan-based anomaly detection for multivariate time series using polluted training set," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 12, pp. 12208–12219, 2021.
- [11] J. Xu, H. Wu, J. Wang, and M. Long, "Anomaly transformer: Time series anomaly detection with association discrepancy," in *International Conference on Learning Representations*, 2021.
- [12] S. Tuli, G. Casale, and N. R. Jennings, "Tranad: deep transformer networks for anomaly detection in multivariate time series data," *Proceedings of the VLDB Endowment*, vol. 15, no. 6, pp. 1201–1214, 2022.
- [13] Y. Yang, C. Zhang, T. Zhou, Q. Wen, and L. Sun, "Dcdetector: Dual attention contrastive representation learning for time series anomaly detection," in *Proceedings of the 29th ACM SIGKDD international conference on knowledge discovery & data mining*, 2023, pp. 3033–3045.
- [14] C. Wang, Z. Zhuang, Q. Qi, J. Wang, X. Wang, H. Sun, and J. Liao, "Drift doesn't matter: Dynamic decomposition with diffusion reconstruction for unstable multivariate time series anomaly detection," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [15] Y. He and J. Zhao, "Temporal convolutional networks for anomaly detection in time series," in *Journal of Physics: Conference Series*, vol. 1213, no. 4. IOP Publishing, 2019, p. 042050.
- [16] A. Gu and T. Dao, "Mamba: Linear-time sequence modeling with selective state spaces," *arXiv preprint arXiv:2312.00752*, 2023.
- [17] R. B. Cleveland, W. S. Cleveland, J. E. McRae, I. Terpenning *et al.*, "Stl: A seasonal-trend decomposition," *J. Off. Stat.*, vol. 6, no. 1, pp. 3–73, 1990.
- [18] R. J. Hodrick and E. C. Prescott, "Postwar us business cycles: an empirical investigation," *Journal of Money, credit, and Banking*, pp. 1–16, 1997.
- [19] H. Wu, J. Xu, J. Wang, and M. Long, "Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting," *Advances in neural information processing systems*, vol. 34, pp. 22419–22430, 2021.
- [20] T. Zhou, Z. Ma, Q. Wen, X. Wang, L. Sun, and R. Jin, "Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting," in *International conference on machine learning*. PMLR, 2022, pp. 27268–27286.
- [21] Z. Zhang, R. Wang, R. Ding, and Y. Gu, "Unravel anomalies: an end-to-end seasonal-trend decomposition approach for time series anomaly detection," in *2024 International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2024, pp. 5415–5419.
- [22] S. Qin, J. Zhu, D. Wang, L. Ou, H. Gui, and G. Tao, "Decomposed transformer with frequency attention for multivariate time series anomaly detection," in *2022 International Conference on Big Data*. IEEE, 2022, pp. 1090–1098.
- [23] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv preprint arXiv:1803.01271*, 2018.
- [24] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [26] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *International conference on machine learning*. PMLR, 2013, pp. 1310–1318.
- [27] A. Gu, K. Goel, and C. Re, "Efficiently modeling long sequences with structured state spaces," in *International Conference on Learning Representations*, 2021.
- [28] Y. Liu, Y. Tian, Y. Zhao, H. Yu, L. Xie, Y. Wang, Q. Ye, and Y. Liu, "Vmamba: Visual state space model," *arXiv preprint arXiv:2401.10166*, 2024.
- [29] X. Jiang, C. Han, and N. Mesgarani, "Dual-path mamba: Short and long-term bidirectional selective structured state space models for speech separation," *arXiv preprint arXiv:2403.18257*, 2024.
- [30] S. Elfving, E. Uchibe, and K. Doya, "Sigmoid-weighted linear units for neural network function approximation in reinforcement learning," *Neural networks*, vol. 107, pp. 3–11, 2018.
- [31] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "Lof: identifying density-based local outliers," in *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, 2000, pp. 93–104.
- [32] Z. He, X. Xu, and S. Deng, "Discovering cluster-based local outliers," *Pattern recognition letters*, vol. 24, no. 9-10, pp. 1641–1650, 2003.
- [33] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson, "Estimating the support of a high-dimensional distribution," *Neural computation*, vol. 13, no. 7, pp. 1443–1471, 2001.
- [34] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in *2008 eighth IEEE international conference on data mining*. IEEE, 2008, pp. 413–422.
- [35] M. Goldstein and A. Dengel, "Histogram-based outlier score (hbos): A fast unsupervised anomaly detection algorithm," *KI-2012: poster and demo track*, vol. 1, pp. 59–63, 2012.
- [36] T. Pevný, "Loda: Lightweight on-line detector of anomalies," *Machine Learning*, vol. 102, pp. 275–304, 2016.
- [37] Z. Li, Y. Zhao, X. Hu, N. Botta, C. Ionescu, and G. H. Chen, "Ecod: Unsupervised outlier detection using empirical cumulative distribution functions," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 12, pp. 12181–12193, 2022.
- [38] A. P. Mathur and N. O. Tippenhauer, "Swat: A water treatment testbed for research and training on ics security," in *2016 international workshop on cyber-physical systems for smart water networks (CySWater)*. IEEE, 2016, pp. 31–36.
- [39] Y. Zhao, Z. Nasrullah, and Z. Li, "Pyod: A python toolbox for scalable outlier detection," *Journal of machine learning research*, vol. 20, no. 96, pp. 1–7, 2019.
- [40] A. Siffer, P.-A. Fouque, A. Termier, and C. Largouet, "Anomaly detection in streams with extreme value theory," in *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery & data mining*, 2017, pp. 1067–1075.
- [41] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [42] S. Kim, K. Choi, H.-S. Choi, B. Lee, and S. Yoon, "Towards a rigorous evaluation of time-series anomaly detection," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 7, 2022, pp. 7194–7201.
- [43] A. Huet, J. M. Navarro, and D. Rossi, "Local evaluation of time series anomaly detection algorithms," in *Proceedings of the 28th ACM SIGKDD international conference on knowledge discovery & data mining*, 2022, pp. 635–645.