Transactions on Computational and Scientific Methods | Vo. 4, No. 2, 2024
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

Contrastive Masked Autoencoding for Self-Supervised Anomaly Representation Learning in Distributed Task Systems

Jingfeng Zhou[1]*
[1]University of Southern California, Los Angeles, USA
*Corresponding author: Jingfeng Zhou, xavier.zhou.ifcen@gmail.com

Abstract: Addressing the challenge of representing anomalous states in cloud computing, microservice collaboration, and large-scale task scheduling environments, this paper proposes a distributed task anomaly representation learning method based on contrastive learning and masked autoencoders. Considering the characteristics of distributed task execution, such as multi-source heterogeneity, temporal coupling, complex dependencies, and hidden anomaly patterns, the method adopts a unified representation modeling approach, jointly designing the task context recovery process and semantic discrimination process to enhance the model's ability to capture deep features of anomalous states. Specifically, firstly, the multi-dimensional execution signals of the distributed task are temporally organized and mapped to a unified task representation; then, a random masking mechanism is introduced to mask some input information, and the missing content is recovered through collaborative learning between the encoder and decoder, thereby prompting the model to mine global contextual relationships and potential structural information in the task state; based on this, semantically consistent multi-view samples are further constructed, and contrastive constraints are used to strengthen the representational distinction between normal and anomalous states, so that the resulting latent space simultaneously possesses good completeness, robustness, and discriminability. To achieve effective synergy between the two learning objectives, this paper constructs a unified optimization framework that combines reconstruction and contrastive objectives, enabling anomaly representations to balance local detail recovery and global semantic separation. Experiments are conducted using publicly available open-source datasets and compared with related methods under consistent settings. Results show that the proposed method can learn higher-quality distributed task anomaly representations and exhibits better overall performance in anomaly recognition evaluations. This research demonstrates that combining mask modeling and contrastive learning for distributed task state representation construction can provide effective support for anomaly detection, fault analysis, and task state understanding in complex operating environments.

Keywords: Self-supervised representation learning; task state modeling; multi-source heterogeneous monitoring; anomaly semantic representation.

1. Introduction

With the rapid development of cloud computing, edge computing, and microservice architectures, distributed task scheduling and execution mechanisms have been widely applied in key scenarios such as intelligent manufacturing, smart logistics, financial services, internet platforms, and large-scale data processing[1]. A large number of tasks run continuously in a concurrent, heterogeneous, and cross-node collaborative manner, enabling systems to achieve high throughput, high elasticity, and high scalability, but also facing more complex operational risks. Due to long task chains, numerous dependencies, strong resource contention, and dynamic changes in the operating environment, abnormal states often no longer manifest as localized failures on a single node, but rather exhibit characteristics such as cross-component propagation, enhanced temporal coupling, and the superposition of multiple source signals. This complexity makes traditional anomaly detection methods, which rely on manual rules, threshold settings, or shallow statistical features, unable to accurately characterize potential anomaly patterns during distributed task execution, and also unable to meet the current practical needs of intelligent operation and maintenance and system reliability assurance for high-precision representation learning[2,3].

From the perspective of task operation mechanisms, distributed task anomalies are not only affected by factors such as fluctuations in resource utilization, service call latency, task dependency imbalance, and changes in node load, but are also closely related to the task context, execution phase, and implicit semantic relationships[4]. Anomalies in real-world scenarios typically exhibit sparsity, suddenness, and non-stationarity. Their outward manifestations may be similar, but their internal causes can differ significantly; conversely, different surface symptoms may correspond to the same underlying fault source. This diversity and uncertainty makes it difficult to stably represent anomaly samples, and also causes supervised learning methods to face problems such as high annotation costs, incomplete anomaly categories, and limited generalization ability. Therefore, how to mine the essential structure of task states from complex operational data under limited or even weak supervision, and construct anomaly representations with discriminative, robust, and transferable capabilities, has become a crucial problem urgently needing to be solved in the field of distributed task anomaly analysis[5].

Contrastive learning provides an effective self-supervised representation learning approach for complex system state modeling. Its core lies in enhancing the model's ability to capture key discriminative features through explicit constraints on similarity and difference relationships between samples. In distributed task scenarios, different time slices, different node perspectives, and different task contexts contain rich semantic consistency and structural differences[6,7]. If representation alignment and differentiation can be achieved through a reasonable positive and negative sample construction mechanism, it is expected to improve the model's ability to characterize the boundaries between normal and abnormal patterns. Meanwhile, masked autoencoders, by partially masking the input information and reconstructing missing content, force the model to learn more global and context-sensitive latent representations, thereby enhancing its adaptability to local noise interference and incomplete observation conditions. Combining contrastive learning with masked autoencoders holds promise for simultaneously achieving discriminative feature extraction and structural information recovery, providing a new methodological path for distributed task anomaly representation learning.

Research on distributed task anomaly representation learning has both significant theoretical value and important engineering implications. Theoretically, this research helps to promote the in-depth application of self-supervised learning paradigms in anomaly modeling of complex systems, facilitates a systematic understanding of task behavior structure, state evolution laws, and anomaly semantic space, and provides new ideas for the unified representation of high-dimensional temporal heterogeneous data. From an application perspective, accurate, stable, and generalizable anomaly representations can provide reliable support for task anomaly detection, fault location, risk warning, resource scheduling optimization, and system self-healing decision-making, thereby improving the observability, stability, and service continuity of distributed systems. Against the backdrop of deepening digital transformation and increasing complexity in the operation of critical infrastructure, research on distributed task anomaly representation learning methods based on contrastive learning and masked autoencoders has significant practical necessity and long-term development significance.

2. Methodology

Let a distributed task stream be segmented into N task instances, where the ith instance is represented by a multivariate observation matrix $\boldsymbol{X}_i \in \mathbb{R}^{T\times D}$ that encodes temporal status, resource usage, dependency response, and execution context over T time steps and D feature channels. Rather than relying on manually designed indicators, the method first projects each task instance into a latent space so that heterogeneous runtime signals can be aligned into a unified representation domain, which is essential for reducing feature inconsistency across nodes and execution stages. Its overall model architecture is shown in Figure 1.

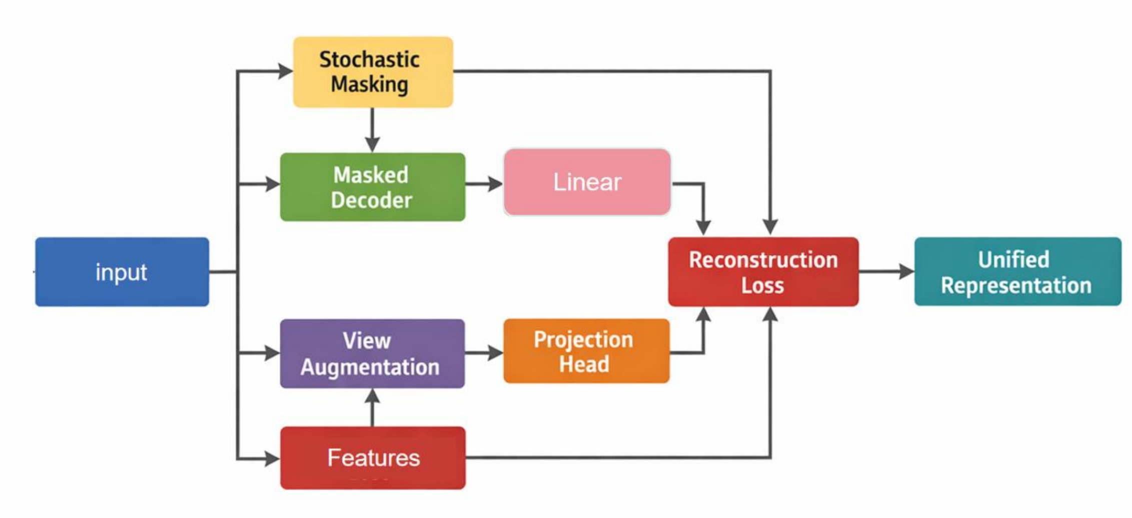

Figure 1. Overall model architecture

A learnable encoder f_θ is therefore introduced to map raw task observations into compact embeddings, allowing the model to preserve structural regularities while suppressing irrelevant fluctuations caused by transient noise and incomplete monitoring. Such a design makes anomaly representation learning less dependent on explicit labels and more sensitive to intrinsic task-state deviations hidden in high-dimensional sequences. Formally, the latent feature of the ith task instance is defined as:

$$\boldsymbol{h}_i = f_\theta\left(\boldsymbol{X}_i\right)$$

To prevent the representation from collapsing into overly smooth summaries, a stochastic masking operator $\mathcal{M}(\cdot)$ is imposed before reconstruction, where a binary mask matrix $\boldsymbol{M}_i \in \{0,1\}^{T\times D}$ hides partial observations and forces the encoder to infer missing semantics from preserved context. In this way, the model is encouraged to capture long-range task dependencies, temporal coherence, and implicit cross-feature correlations that are often disrupted when anomalies emerge. The masked input is written as:

$$\widetilde{\boldsymbol{X}}_i = \boldsymbol{M}_i \odot \boldsymbol{X}_i$$

Afterward, a decoder g_ϕ reconstructs the original task state from the masked representation, and the recovery process explicitly emphasizes semantic completeness rather than shallow point-wise memorization, thereby improving the robustness of the learned embedding under partial observation and noisy scheduling conditions. The reconstruction branch is expressed as:

$$\widehat{\boldsymbol{X}}_i = g_\phi\left(f_\theta\left(\widetilde{\boldsymbol{X}}_i\right)\right)$$

Since anomaly-sensitive features are frequently hidden in corrupted or weakly expressed local patterns, the reconstruction objective is computed only on masked positions so that the model is driven to recover informative content instead of simply copying visible inputs. The masked reconstruction loss is given by:

$$\mathcal{L}_{rec} = \frac{1}{N}\sum_{i=1}^{N} \frac{\left\|\left(1-\boldsymbol{M}_i\right) \odot \left(\widehat{\boldsymbol{X}}_i - \boldsymbol{X}_i\right)\right\|_2^2}{\sum\left(1-\boldsymbol{M}_i\right)}$$

Beyond recovering missing observations, discriminative anomaly representation requires the latent space to preserve similarity among semantically consistent task states while enlarging the gap between incompatible patterns. To achieve this goal, two augmented views are generated for each task instance through time slicing, feature dropout, dependency perturbation, or context-preserving noise injection, and these transformations maintain the core execution semantics while exposing the representation model to realistic monitoring variation. Denoting the two views of the same task instance by $\mathbf{X}_i^{(1)}$ and $\mathbf{X}_i^{(2)}$, their projected embeddings are obtained through an encoder-projector pipeline so that contrastive optimization can operate in a space more suitable for relational discrimination. The view-level embeddings are formulated as:

$$\boldsymbol{z}_i^{(1)} = q_\psi\left(f_\theta\left(\mathbf{X}_i^{(1)}\right)\right), \qquad \boldsymbol{z}_i^{(2)} = q_\psi\left(f_\theta\left(\mathbf{X}_i^{(2)}\right)\right)$$

With cosine similarity serving as the affinity measure, positive pairs are built from two views of the same task instance, whereas the remaining samples within a mini-batch act as negatives, which encourages the model to learn boundaries aligned with latent abnormality rather than superficial amplitude differences. This mechanism is particularly meaningful in distributed environments because many anomalous states share partial local symptoms with normal states, and only relational learning can reveal their deeper semantic discrepancy. The contrastive objective is defined as:

$$\mathcal{L}_{con} = -\frac{1}{N}\sum_{i=1}^{N} log \frac{exp\left(sim\left(\boldsymbol{z}_i^{(1)}, \boldsymbol{z}_i^{(2)}\right)/\tau\right)}{\sum_{j=1}^{N} exp\left(sim\left(\boldsymbol{z}_i^{(1)}, \boldsymbol{z}_j^{(2)}\right)/\tau\right)}$$

where $sim(\boldsymbol{a}, \boldsymbol{b}) = \boldsymbol{a}^\top \boldsymbol{b}/(\|\boldsymbol{a}\|_2 \|\boldsymbol{b}\|_2)$ measures angular consistency, and the temperature coefficient τ controls the concentration of instance discrimination. By combining invariance across positive views with repulsion against irrelevant samples, the latent space becomes more structured, which is crucial for separating subtle abnormal task transitions from diverse yet legitimate execution fluctuations.

A further consideration is that reconstruction and contrastive alignment do not contribute equally throughout training, because early optimization should stabilize context modeling while later stages should refine discriminative structure. Motivated by this observation, the overall objective introduces adaptive balancing coefficients to coordinate semantic recovery and relational separation, allowing the encoder to progressively shift from general pattern acquisition to anomaly-aware representation shaping. Such a joint formulation not only reduces the risk that the masked autoencoding branch overemphasizes smooth reconstruction but also alleviates the instability that pure contrastive learning may suffer under highly imbalanced runtime states. The unified training loss is expressed as:

$$\mathcal{L} = \lambda\mathcal{L}_{rec} + (1-\lambda)\mathcal{L}_{con}$$

Here, $\lambda \in [0,1]$ adjusts the relative influence of context restoration and semantic discrimination, and the final encoder output $\boldsymbol{h}_i$ is therefore shaped by both local recoverability and global relational consistency. In practical terms, this synergy enables the learned representation to encode missing-data resilience, temporal continuity, and anomaly-sensitive separation within a single embedding space, which is more suitable for distributed task monitoring than relying on either branch alone.

Once optimization is completed, each task instance is associated with a compact embedding that summarizes its execution semantics under both reconstruction and contrastive constraints, and this representation can serve as the foundation for downstream anomaly analysis, state clustering, or risk scoring. Different from hand-crafted statistics that mainly characterize visible deviations, the proposed latent representation absorbs hidden task dependencies and contextual uncertainty, thereby offering a more expressive description of abnormal evolution in complex distributed systems. From a methodological perspective, masked modeling contributes completeness-aware semantic understanding, while contrastive learning injects boundary-aware discriminability; together they form a unified anomaly representation framework that is better aligned with the nonstationary, heterogeneous, and partially observed nature of distributed task execution. As a result, the method establishes a principled path for capturing latent anomaly patterns that are difficult to expose through conventional signal engineering or solely supervised modeling.

3. Experimental Results and Analysis

3.1 Dataset

This paper selects Aiops-Dataset as its data foundation. This dataset originates from a microservice-based system deployed in a real cloud environment. The system comprises 46 instances, including 40 microservice instances and 6 virtual machine instances. The data is organized chronologically, covering three types of observation information: logs, metrics, and traces, and provides corresponding groundtruth fault records. For distributed task anomaly representation learning, this type of open-source data, which simultaneously includes call chains, runtime metrics, and log information, can effectively reflect the state evolution, dependency propagation, and anomaly perturbation characteristics of tasks during cross-service collaborative execution. Therefore, it aligns well with the research theme of this paper, which focuses on anomaly modeling for distributed tasks.

From the perspective of task semantic representation, Aiops-Dataset not only preserves multi-source heterogeneous observation signals but also provides annotation information such as fault injection time, severity, root cause instance, and fault category. This provides a reliable basis for constructing the discriminative relationship between normal and abnormal patterns. Because this dataset contains trace-level call relationships and time-series monitoring information, it is suitable for learning context recovery capabilities through masking mechanisms, as well as for uncovering potential similarities and differences between different task states through contrastive learning. Compared to data sources containing only a single metric or a single log, this open-source dataset is more conducive to supporting anomaly representation learning research in distributed task scenarios.

3.2 Experimental setup

To verify the effectiveness of the proposed method in anomaly representation learning for distributed tasks, experiments were conducted on a publicly available microservice anomaly dataset. Input samples were constructed following a unified data preprocessing and training process, and multi-source runtime signals such as logs, metrics, and call chains were mapped to time-series task representations. During the overall training process, a random masking mechanism was used to generate missing observation views, while a view augmentation strategy was combined to construct contrastive learning sample pairs, enabling the model to learn anomaly representations under the dual constraints of context recovery and discriminative feature alignment. To ensure training stability, the AdamW optimizer was selected, and the loss function consisted of a weighted sum of reconstruction loss and contrastive loss. Core parameters such as batch size, learning rate, masking ratio, temperature coefficient, and number of training epochs were configured relatively robustly. All experiments were conducted in the same hardware and software environment, using a consistent data partitioning and evaluation process to reduce the impact of external interference factors on the results analysis. Specific experimental settings are shown in Table 1.

Table 1. Experimental Setup Parameters

Parameter Category	Parameter Name	Value
Data Processing	Time Window Length	20
Data Processing	Sliding Step	5
Data Processing	Input Feature Dimension	128

Data Processing	Normalization Method	Z-score
Masked Learning	Mask Ratio	0.15
Masked Learning	Masking Strategy	Random Masking
Contrastive Learning	Projection Layer Dimension	128
Contrastive Learning	Temperature Coefficient	0.2
Model Training	Batch Size	32
Model Training	Optimizer	AdamW
Model Training	Initial Learning Rate	0.0005
Model Training	Weight Decay	0.0001
Model Training	Training Epochs	100
Loss Function	Reconstruction Loss Weight	1.0
Loss Function	Contrastive Loss Weight	0.5

3.3 Experimental Results and Analysis

To conduct a horizontal comparison of methods within the research context of anomaly representation learning for distributed tasks, this paper selects representative works closely related to self-supervised modeling, contrastive learning, mask reconstruction, temporal anomaly detection, and microservice anomaly analysis for comparative analysis. The research topics of the selected comparative methods are highly relevant to the anomaly representation learning direction focused on in this paper. Based on these works, a unified comparative framework can be further constructed to standardize the presentation of different methods in terms of accuracy, robustness, and comprehensive discriminative ability. The comparison settings are shown in Table 2.

Table 2. Experimental results compared with other models

Method	**Acc**	**Pre**	**Rec**	**F1**	**AUC**	**MCC**	**FAR**	**AP**
Yang et al.[8]	91.24	89.76	88.95	89.35	94.18	82.47	6.82	92.06
Jeong et al.[9]	92.08	90.54	89.63	90.08	94.87	83.91	6.35	92.84
Tuli et al.[10]	92.67	91.13	90.42	90.77	95.36	84.68	5.94	93.41
Fu et al.[11]	93.15	91.88	91.04	91.46	95.92	85.54	5.62	94.03
Xie et al.[12]	93.72	92.41	91.86	92.13	96.48	86.37	5.28	94.69
Zhang et al.[13]	94.18	92.96	92.34	92.65	96.91	87.12	4.96	95.21
Chen et al.[14]	94.56	93.28	92.77	93.02	97.24	87.86	4.71	95.68

Ours	**96.14**	**95.08**	**94.73**	**94.90**	**98.31**	**90.42**	**3.28**	**97.12**

Overall, the proposed algorithm demonstrates strong comprehensive advantages across various evaluation metrics, indicating its ability to adequately characterize the abnormal states during distributed task execution. This is because the model simultaneously considers context recovery and discriminative feature modeling during the representation learning phase. This ensures that the learned latent representation not only retains key structural information from the task execution process but also enhances the ability to distinguish between normal and abnormal states. For distributed tasks, which exhibit multi-source heterogeneity, dependency coupling, and temporal dynamics, relying solely on a single path for modeling often fails to yield stable and complete anomaly representations. The proposed method, through joint mask reconstruction and contrastive constraints, effectively alleviates issues such as information loss, noise interference, and insignificant local anomalies, thus achieving higher-quality anomaly representations.

Furthermore, the algorithm accurately identifies anomalous samples while maintaining good discriminative stability and overall robustness, indicating that the constructed representation space possesses strong separability and consistency. Especially in complex task scenarios, anomalies often do not manifest as significant shifts in a single metric, but rather as potential imbalances across time, nodes, and tasks. Therefore, models need the ability to extract deep anomaly patterns from a global semantic level. The method presented in this paper demonstrates good adaptability in this regard; its learned feature representations can more accurately reflect key anomaly signals in task state changes, while reducing the impact of irrelevant perturbations on the decision-making process. This indicates that the method has high application value in the problem of anomaly representation learning for distributed tasks and provides a more reliable representation foundation for subsequent anomaly detection, fault warning, and operational status analysis.

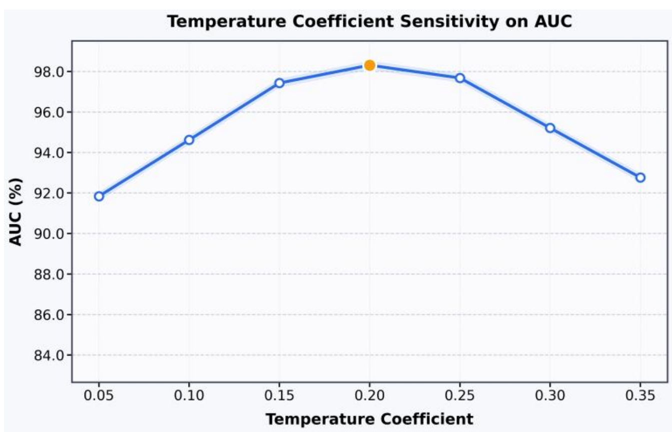

Figure 2. Experimental results of the temperature coefficient hyperparameter on AUC

As shown in Figure 2, the proposed algorithm maintains strong anomaly recognition capabilities and representation stability under the set temperature coefficient, indicating good synergy between the contrast constraint mechanism and the overall representation learning framework. A reasonable temperature coefficient helps enhance the model's ability to characterize potential anomaly semantic differences, enabling the learned representation to retain key structural information during distributed task execution while improving the distinction between abnormal and normal states. This demonstrates that the proposed method has good adaptability during hyperparameter tuning, and the constructed anomaly representation space can provide reliable support for subsequent task state analysis and anomaly discrimination.

4. Conclusion

This paper addresses the challenges of representing abnormal states, complex operational semantics, and strong heterogeneity of multi-source observation information in distributed task scenarios. It proposes an anomaly representation learning method based on contrastive learning and masked autoencoders. Starting from the operational process of distributed tasks, this method organically combines context recovery capabilities with discriminative feature modeling. Within a unified framework, it simultaneously focuses on the structural integrity of task states and the discriminability of anomaly semantics, thereby improving the expressive power and stability of anomaly representations. Compared to traditional methods relying on manual rules, shallow statistical features, or single supervisory signals, this proposed method better suits the dynamic, coupled, and uncertain characteristics of distributed tasks in real-world environments. It can capture more essential anomaly information from complex operational data, providing a more reliable representation foundation for anomaly detection, fault identification, and state analysis. The results show that combining contrastive learning with masked modeling not only enhances the model's adaptability to local missing data and noise perturbations but also further improves the quality of the latent representation's characterization of anomaly boundaries. From a methodological perspective, this research provides a valuable research path for anomaly modeling in distributed tasks.

From an application perspective, this research has positive significance for improving the perception, early warning, and decision-making support capabilities of intelligent operation and maintenance systems. In scenarios such as cloud computing platforms, microservice systems, edge collaboration environments, large-scale data processing platforms, and the Industrial Internet, distributed tasks often undertake core functions such as resource scheduling, data transmission, service orchestration, and critical business execution. If abnormal states are not identified promptly, they can lead to performance degradation, service interruptions, and even cascading failures. The anomaly representation learning method proposed in this paper can extract more discriminative deep features from complex temporal observations and task context relationships. This makes it not only applicable to anomaly detection itself but also suitable for related tasks such as root cause analysis, risk warning, resource optimization scheduling, task health assessment, and system self-healing control. For the ever-expanding intelligent computing infrastructure, establishing a unified, robust, and transferable task state representation mechanism in complex heterogeneous environments has become a crucial foundation for driving the intelligent evolution of systems. This work provides practical technical support and theoretical reference in this direction.

Further research in this area still has significant potential for expansion. First, distributed tasks in real-world applications often involve stronger cross-layer dependencies and longer-term state evolution. Therefore, further combining finer-grained task link semantics, cross-node collaborative information, and long-term temporal dependency modeling mechanisms can enhance the ability of anomaly representations to characterize complex propagation behaviors and hidden anomaly patterns. Second, given the distributional differences arising from different platforms, business models, and monitoring granularities, future work can focus on cross-scenario transfer learning, continuous learning, and online incremental updates to enhance the model's adaptability in dynamic environments and its long-term deployment value. Furthermore, with increasing emphasis on interpretability, security, and engineering feasibility, future research can delve deeper into interpretable analysis of the anomaly representation space, lightweight deployment, and real-time inference optimization to better serve applications such as cloud service governance, intelligent industrial operations and maintenance, digital infrastructure monitoring, and critical business system assurance. Overall, this work not only provides new research ideas for anomaly representation learning in distributed tasks but also lays the foundation for building a more intelligent, reliable, and sustainably evolving framework for analyzing the operation of complex systems.

References

[1] Z. Z. Darban, G. I. Webb, S. Pan, et al., "CARLA: Self-supervised contrastive representation learning for time series anomaly detection," arXiv preprint arXiv:2308.09296, 2023.

[2] J. Xu, H. Wu, J. Wang, et al., "Anomaly Transformer: Time series anomaly detection with association discrepancy," in Proc. Int. Conf. Learning Representations (ICLR), 2022.

[3] G. Zhong, F. Liu, J. Jiang, et al., "Refining one-class representation: A unified transformer for unsupervised time-series anomaly detection," Information Sciences, vol. 656, p. 119914, 2024.

[4] J. Song, K. Kim, J. Oh, et al., "Memto: Memory-guided transformer for multivariate time series anomaly detection," in Advances in Neural Information Processing Systems, vol. 36, 2023, pp. 57947–57963.

[5] H. Guo, S. Yuan, and X. Wu, "LogBERT: Log anomaly detection via BERT," in Proc. Int. Joint Conf. Neural Networks (IJCNN), 2021, pp. 1–8.

[6] M. Panahandeh, A. Hamou-Lhadj, M. Hamdaqa, et al., "ServiceAnomaly: An anomaly detection approach in microservices using distributed traces and profiling metrics," Journal of Systems and Software, vol. 209, p. 111917, 2024.

[7] J. Nyyssölä, M. Mäntylä, and M. Varela, "How to configure masked event anomaly detection on software logs?," in Proc. IEEE Int. Conf. Software Maintenance and Evolution (ICSME), 2022, pp. 414–418.

[8] Y. Yang, C. Zhang, T. Zhou, et al., "Dcdetector: Dual attention contrastive representation learning for time series anomaly detection," in Proc. 29th ACM SIGKDD Conf. Knowledge Discovery and Data Mining, 2023, pp. 3033–3045.

[9] Y. Jeong, E. Yang, J. H. Ryu, et al., "AnomalyBERT: Self-supervised transformer for time series anomaly detection using data degradation scheme," arXiv preprint arXiv:2305.04468, 2023.

[10] S. Tuli, G. Casale, and N. R. Jennings, "TranAD: Deep transformer networks for anomaly detection in multivariate time series data," arXiv preprint arXiv:2201.07284, 2022.

[11] Y. Fu and F. Xue, "MAD: Self-supervised masked anomaly detection task for multivariate time series," in Proc. Int. Joint Conf. Neural Networks (IJCNN), 2022, pp. 1–8.

[12] Y. Xie, H. Zhang, and M. A. Babar, "LogSD: Detecting anomalies from system logs through self-supervised learning and frequency-based masking," Proc. ACM Softw. Eng., vol. 1, no. FSE, pp. 2098–2120, 2024.

[13] C. Zhang, X. Peng, C. Sha, et al., "DeepTraLog: Trace-log combined microservice anomaly detection through graph-based deep learning," in Proc. 44th Int. Conf. Software Engineering, 2022, pp. 623–634.

[14] J. Chen, F. Liu, J. Jiang, et al., "TraceGra: A trace-based anomaly detection for microservice using graph deep learning," Computer Communications, vol. 204, pp. 109–117, 2023.