

Finding Time Series Anomalies Using Granular-Ball Vector Data Description

Lifeng Shen¹, Liang Peng¹, Ruiwen Liu¹, Shuyin Xia^{1*}, Yi Liu²

¹Chongqing Key Laboratory of Computational Intelligence, Key Laboratory of Cyberspace Big Data Intelligent Security, Ministry of Education, Sichuan-Chongqing Co-construction Key Laboratory of Digital Economy Intelligence and Key Laboratory of Big Data Intelligent Computing, Chongqing University of Posts and Telecommunications.

²Chongqing Ant Consumer Finance Co., Ltd., Ant Group.

shenlf@cqupt.edu.cn, llangpeng@foxmail.com, liuruiwen622@foxmail.com, xiasy@cqupt.edu.cn, larry.liuy@myxiaojin.cn

Abstract

Modeling normal behavior in dynamic, nonlinear time series data is challenging for effective anomaly detection. Traditional methods, such as nearest neighbor and clustering approaches, often depend on rigid assumptions, such as a predefined number of reliable neighbors or clusters, which frequently break down in complex temporal scenarios. To address these limitations, we introduce the Granular-ball One-Class Network (GBOC), a novel approach based on a data-adaptive representation called Granular-ball Vector Data Description (GVDD). GVDD partitions the latent space into compact, high-density regions represented by granular-balls, which are generated through a density-guided hierarchical splitting process and refined by removing noisy structures. Each granular-ball serves as a prototype for local normal behavior, naturally positioning itself between individual instances and clusters while preserving the local topological structure of the sample set. During training, GBOC improves the compactness of representations by aligning samples with their nearest granular-ball centers. During inference, anomaly scores are computed based on the distance to the nearest granular-ball. By focusing on dense, high-quality regions and significantly reducing the number of prototypes, GBOC delivers both robustness and efficiency in anomaly detection. Extensive experiments validate the effectiveness and superiority of the proposed method, highlighting its ability to handle the challenges of time series anomaly detection.

Code — <https://github.com/notshine/GBOC>

1 Introduction

Complex cyber-physical systems, such as power plants, data centers, and smart factories, rely on a multitude of sensors operating concurrently to continuously generate high-volume multivariate time series data. Timely and accurate detection of anomalies within such data streams is critical, as it enables effective monitoring of system health, prevention of potential failures, and reduction of operational risks. Time series anomaly detection, involves identifying deviations from expected behavior at the level of individual time steps, with diverse real-world manifestations (Schmidl,

Wenig, and Papenbrock 2022). Examples of such anomalous events include physical attacks on industrial systems (Mathur and Tippenhauer 2016), unpredictable robot behavior (Park, Hoshi, and Kemp 2018), faulty sensors from wide-sensor networks (Wang, Kuang, and Duan 2015; Rassam, Maarof, and Zainal 2018), cybersecurity attacks on server monitoring systems (Su et al. 2019; Abdulaal, Liu, and Lancewicki 2021), and spacecraft malfunctions observed via telemetry sensor data (Liu, Liu, and Peng 2016).

Describing normal data in time series is crucial for effective anomaly detection. Commonly used approaches, such as nearest neighbor and clustering methods, aim to characterize normal patterns. Nearest neighbor approaches, like KNN (Ramaswamy, Rastogi, and Shim 2000), detect anomalies based on proximity. However, they struggle with the dynamic and complex nature of time series data. For instance, in group anomalies, a set of abnormal points may still have enough “neighbors” within the group, making them appear normal. This issue arises because these methods rely on local density or proximity, failing to capture the global or temporal structure of the data.

Clustering-based methods, such as KShapeAD (Paparrizos and Gravano 2015, 2017; Boniol et al. 2021), group data points into clusters representing normal behavior. These approaches generally assume that normal data forms dense, well-defined clusters, while anomalies are located far from clusters or in sparse, low-density regions. One-class classification methods, such as OCSVM (Schölkopf et al. 2001) and SVDD (Tax and Duin 2004), further simplify this concept by modeling normal data as a single cluster, often enclosed within a hypersphere. To enhance the flexibility of normality modeling, DeepSVDD (Ruff et al. 2018) leverages deep learning to learn a more adaptive hypersphere, while THOC (Shen, Li, and Kwok 2020) introduces a multiscale vector data description through a differentiable hierarchical clustering process. Similarly, memory-based approaches, such as MEMTO (Song et al. 2023), extend clustering by incorporating a memory module to store representative prototypes (or centroids) of normal clusters.

While clustering approaches are conceptually intuitive, they face notable limitations. Most clustering methods require the number of clusters to be predefined, which is particularly problematic when the true data structure is unknown or involves complex, multimodal distributions. One-

*Corresponding Author.

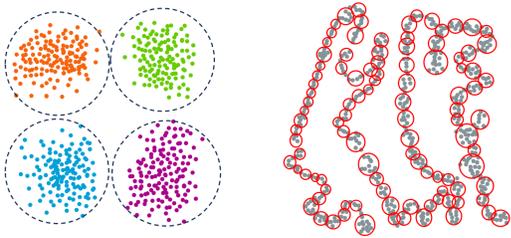


Figure 1: Latent space comparison between clustering-based modeling (left) and the proposed method preserving structural continuity (right). Clustering-based methods assume well-defined boundaries, which are unsuitable for the continuous and boundary-free distributions typical of time series data, requiring a more adaptive descriptive approach.

class classification methods address this by representing normal behavior with a single hypersphere, but this oversimplified assumption struggles to capture the diversity and complexity of multimodal patterns. Similarly, memory-based methods depend heavily on the quality and representativeness of stored prototypes. When these prototypes fail to encompass the full variability of normal behavior, the methods become biased and significantly less effective.

These limitations are especially pronounced in dynamic and nonlinear time series data. As illustrated in Figure 1 (right), time series sliding windows’ representations often exhibit structural continuity, characterized by smooth transitions between patterns. In contrast, the rigid cluster assignments depicted in Figure 1 (left) rely on discrete boundaries and separability, making them ill-suited to handle the continuous and dynamic nature of such representations.

To address the challenges of modeling normal behavior in complex and dynamic time series data, we propose the Granular-ball One-Class Network (GBOC). Unlike traditional methods, which rely on rigid assumptions such as predefined clusters or reliable neighbors, GBOC introduces the Granular-ball Vector Data Description (GVDD), a flexible representation that adaptively partitions the latent space into compact, high-density regions called granular-balls, illustrated in Figure 1 (right). These regions are constructed through a density-guided hierarchical splitting process and refined by pruning noisy or diffuse areas, ensuring robust modeling of normal behavior. During training, GBOC aligns samples with their nearest granular-ball centers to enhance representation compactness. At inference, anomaly scores are efficiently computed based on the distance to the nearest granular-ball center, leveraging the dense, high-quality regions. By reducing the number of prototypes compared to training samples, GBOC achieves computational efficiency and robustness. Experiments on diverse time series benchmarks demonstrate that GBOC outperforms traditional and deep learning-based methods in accuracy, scalability, and efficiency, offering a practical solution for real-world anomaly detection tasks.

Contribution. We propose the Granular-ball One-Class Network (GBOC), a novel time series anomaly detection method that introduces Granular-ball Vector Data Descrip-

tion (GVDD), the first extension of one-class methods by incorporating granular-ball computing. This design enables a time series adaptive vector data description, effectively alleviating the limitations of nearest-neighbor and clustering-based approaches. Extensive experiments demonstrate the superior performance of GBOC against strong baselines.

2 Preliminaries

Problem Formulation: A time series is a sequence of data points recorded at successive time intervals. It can be represented as: $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]^\top \in \mathbb{R}^{T \times d}$, where T is the number of time steps and d is the number of channels. A time series anomaly is a data point \mathbf{x}_t or subsequence $\mathbf{X}_{t_1:t_2} = [\mathbf{x}_{t_1}, \dots, \mathbf{x}_{t_2}]^\top$ (for $1 \leq t_1 \leq t_2 \leq T$) that deviates significantly from expected patterns. Anomalies are rare and misaligned with the statistical properties of normal data. Time series anomaly detection aims to identify time series anomaly points that deviate significantly from the expected value \hat{x}_t predicted by a detection model \mathcal{M} . For a predefined threshold $\delta_t > 0$, the detection criterion is: $\mathcal{M}(\mathbf{x}_t) > \delta_t$, where \mathcal{M} is trained to learn normal patterns from historical data and assigns an anomaly score to each observation. More related works on are provided in Appendix A.

Granular-ball Computing (GBC) (Xia et al. 2019) is a geometry-aware and flexible framework that models complex data by abstracting it into adaptive local regions, termed *granular-balls*. Inspired by the “global precedence” principle in cognitive science (Chen 1982), GBC enables multi-granularity cognitive modeling (Wang 2017) through compact and robust region-based representations. Initially developed for clustering, it has shown effectiveness in (Xia et al. 2020; Cheng et al. 2023; Xia et al. 2025a). Recently, GBC has been extended to point cloud registration (Hu et al. 2025), multiple kernel clustering (Xia et al. 2025b), open intent classification (MOGB) (Li et al. 2025), and multi-view contrastive learning (Su et al. 2025a), demonstrating its versatility. Notably, while MOGB uses granular-balls as multi-granularity decision boundaries to capture *class-wise structures* for open intent recognition, GBOC instead constructs adaptive density-based granular-balls to model *normal patterns without explicit class labels*, accommodating the continuous and boundary-free distributions typical of time series data.

Formally, a dataset $\{\mathbf{x}_t\}_{t=1}^n$ is represented by a set of granular-balls $GB = \{GB_j\}_{j=1}^m$, whose definition is given in Definition 1.

Definition 1 (Granular-ball (GB)) A granular-ball GB is a region in \mathbb{R}^d characterized by its center $\mathbf{c} \in \mathbb{R}^d$ and radius $\mathbf{r} \in \mathbb{R}_{\geq 0}$, derived from a finite set of points $\{P_i\}_{i=1}^{|GB|} \subset \mathbb{R}^d$ contained within it. The center \mathbf{c} is defined as the mean of all points in GB : $\mathbf{c} = \frac{1}{|GB|} \sum_{i=1}^{|GB|} \mathbf{x}_i$ and the radius \mathbf{r} is defined as the maximum Euclidean distance from any point $\mathbf{x}_i \in GB$ to the center \mathbf{c} :

$$\mathbf{r} = \max_{\mathbf{x}_i \in GB} \|\mathbf{x}_i - \mathbf{c}\|. \quad (1)$$

Based on the definition of a granular-ball, the construction of GBC follows a two-stage procedure. In the first

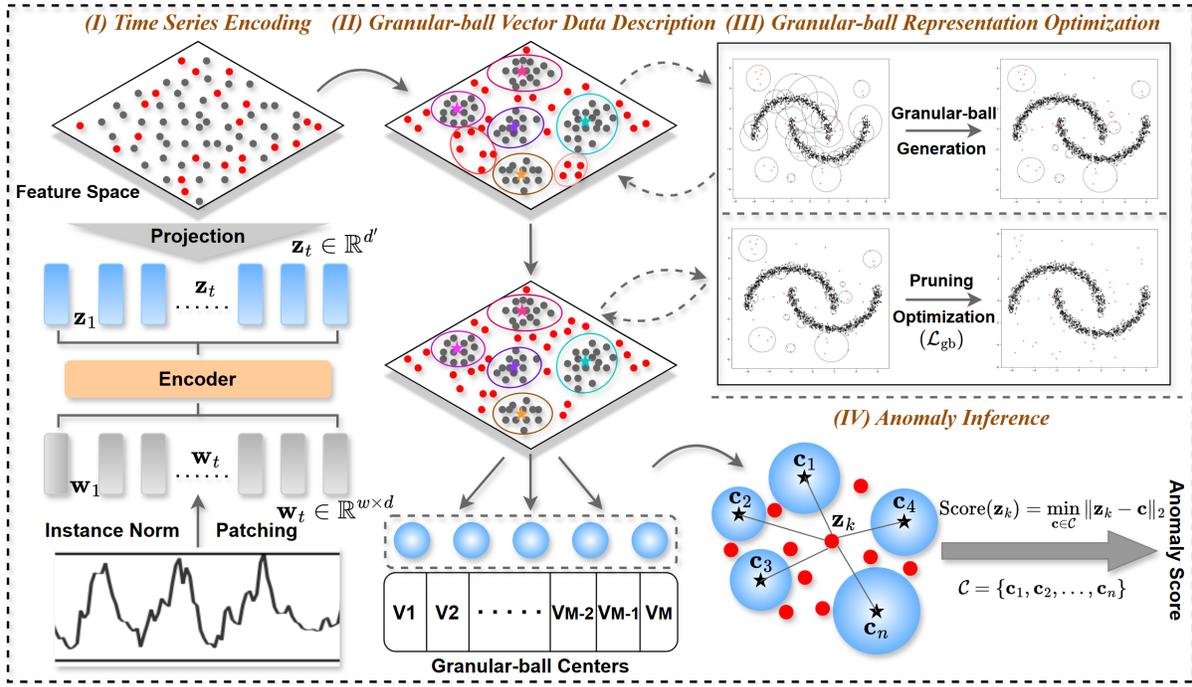


Figure 2: Illustration of the proposed Granular-ball One-Class Network (GBOC).

stage, the dataset $\{\mathbf{x}_t\}_{t=1}^n$ is coarsely partitioned using K-Means into approximately \sqrt{n} clusters, each forming an initial granular-ball. In the second stage, each granular-ball is recursively evaluated and refined into two child balls according to a density criterion: the *granular-ball distribution measure* (DM) (Xie et al. 2023), defined in Definition 2.

Definition 2 (Granular-ball Distribution Measure)

Given a granular-ball GB with center \mathbf{c} and radius r , the DM is defined by computing the ratio of the number $|GB|$ of data points and the sum radius s in GB , as follows:

$$DM(GB) = \frac{s}{|GB|}, \quad (2)$$

where $s = \sum_{\mathbf{x} \in GB} \|\mathbf{x} - \mathbf{c}\|$. A smaller DM value indicates better ball quality.

The granular-ball generation procedure is detailed in Algorithm 1, provided in Appendix A.2. As observed, any granular-ball GB containing at least s_{\min} points is partitioned into two sub-balls by applying the 2-Means clustering algorithm to the points within GB . Here, s_{\min} represents the minimum support requirement for a granular-ball to qualify for splitting, and it is empirically set to $s_{\min} = 8$.

Definition 3 (Refinement Criterion) A granular-ball GB is eligible for splitting if the weighted value of the decision metric (denoted as DM) after the split is strictly less than that of the original ball, i.e., $DM_w < DM(GB)$, and both resulting sub-balls satisfy the minimum support requirement.

The weighted decision metric after splitting is calculated as:

$$DM_w = \frac{|GB^{(1)}|}{|GB|} DM(GB^{(1)}) + \frac{|GB^{(2)}|}{|GB|} DM(GB^{(2)}), \quad (3)$$

where $GB^{(1)}$ and $GB^{(2)}$ are the two sub-balls obtained by applying 2-Means clustering to the data points within GB . This recursive refinement process continues until no further quality-improving splits are possible (Xie et al. 2024a). Upon termination, a set of granular-balls $\{GB_j\}_{j=1}^m$ is produced to cover the dataset.

3 Methodology

In this section, we formally elaborate on the granular-ball one-class network (GBOC) for time series anomaly detection. As shown in Figure 2, GBOC operates in four steps: (i) encoding time series into a latent space, (ii) grouping normal patterns into high-density granular-balls, (iii) refining these patterns with pruning and alignment, and (iv) detecting anomalies by measuring the distance to the closest granular-ball center. Next, we will explain each part step by step.

3.1 Time Series Encoding

As a common practice, the input time series \mathbf{X} is first segmented into overlapping windows $\{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_{N_w}\}$ using a sliding window mechanism, where each window has length w . Each window $\mathbf{w}_i \in \mathbb{R}^{w \times d}$ is encoded into a d' -dimensional latent representation

$$\mathbf{z}_i = f_\theta(\mathbf{w}_i), \quad (4)$$

producing the latent feature set $Z = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_{N_w}\}$ for the granular-ball computing module. In this work, a three-layer Long Short-Term Memory (LSTM) network is used as

the encoder. The final hidden states across layers are concatenated to form \mathbf{z}_i , enabling multilevel temporal dependency modeling. The encoder is modular and can be replaced with other models, such as Transformers (Vaswani et al. 2017), for different time series characteristics.

3.2 Granular-ball Vector Data Description

After obtaining initial representations, granular-ball computing is introduced to adaptively model the distribution of latent time series data by constructing *granular-ball vector descriptions* based on local data density and homogeneity, enabling flexible and effective normality modeling for one-class anomaly detection.

The granular-ball construction in GBOC follows the standard process described in Section 2 but operates in the latent representation space, which is optimized during training. Unlike traditional granular-balls in the original data space (Xia et al. 2019; Xie et al. 2024a,b), traditional granular-balls are not directly suitable for normality modeling due to the following reasons: (i) They are not integrated into the representation learning process, limiting their ability to effectively guide feature learning. (ii) Some granular-balls may be of low quality or confidence, and their inclusion could degrade the overall model performance.

To address these issues, we introduce two key components in our granular-ball representation optimization stage: (i) a pruning mechanism to remove low-quality granular-balls, and (ii) a joint learning strategy to align latent representations with high-quality granular-ball centers while maintaining temporal fidelity.

3.3 Granular-ball Representation Optimization

Elimination of Low-quality Granular-balls. To improve the compactness and reliability of the granular-ball representation, we propose a post-hoc pruning strategy that eliminates low-quality granular-balls characterized by abnormally large radii. Such granular-balls often correspond to overly coarse or poorly localized regions in the latent space, failing to capture fine-grained structural details and potentially introducing noise into the representation learning process.

The pruning process is governed by a dynamic threshold derived from the global distribution of granular-ball radii:

$$r_{\text{th}} = \mu \cdot \max \{ \text{median}(r), \text{mean}(r) \}, \quad (5)$$

where $r = \{r_j\}_{j=1}^{|GB|}$ represents the set of radii for all granular-balls. $\mu = 2$ empirically. Any granular-ball GB_j with radius $r_j > r_{\text{th}}$ is deemed diffuse and subsequently removed. This radius-based pruning strategy effectively discards noisy or low-confidence regions while retaining structurally meaningful and compact granular-balls. These preserved granular-balls more accurately reflect the true distribution of normal data within the latent space.

Nearest Granular-ball Center-Based Optimization. To ensure that learned representations are both semantically compact and temporally informative, we propose a joint optimization framework that combines a loss for preserving the

consistency of granular-ball vector data descriptions with a loss for time series reconstruction.

The granular-ball loss functions as a geometric alignment mechanism, enforcing compact semantic clustering in the latent space by pulling each feature vector toward its corresponding granular-ball centroid. This aggregation mechanism enhances the discriminability of the latent space by clustering samples from the normal mode together, forming a compact feature space representation:

$$\mathcal{L}_{\text{gb}} = \frac{1}{N} \sum_{i=1}^N \|\mathbf{z}_i - \mathbf{c}_{s(i)}\|_2^2, \quad (6)$$

where $s(i)$ denotes the index of the granular-ball assigned to the i -th sample, and $\mathbf{c}_{s(i)}$ is the corresponding center.

To complement this alignment objective and preserve temporal fidelity, we incorporate a *reconstruction loss* that ensures the encoder-decoder pipeline retains essential information from the input. This constraint prevents the model from collapsing representations into overly compressed spaces that lack discriminative features:

$$\mathcal{L}_{\text{rec}} = \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_i - g_\phi(\mathbf{z}_i)\|_2^2, \quad (7)$$

where g_ϕ denotes a light-weight MLP parameterized by ϕ .

Finally, the overall training objective is defined as:

$$\mathcal{L} = \lambda \cdot \mathcal{L}_{\text{rec}} + (1 - \lambda) \cdot \mathcal{L}_{\text{gb}}, \quad (8)$$

where λ is a trade-off coefficient (commonly set to 0.5) that balances the reconstruction fidelity and the strength of the structural alignment.

3.4 Anomaly Inference

During inference, each test instance \mathbf{w} is encoded into a representation $\mathbf{z} = f_\theta(\mathbf{w})$. The degree of deviation from the normal region is quantified by computing the Euclidean distance between \mathbf{z} and the nearest granular-ball center:

$$\text{Score}(\mathbf{z}) = \min_{\mathbf{c} \in \mathcal{C}} \|\mathbf{z} - \mathbf{c}\|_2, \quad (9)$$

where \mathcal{C} denotes the set of retained granular-ball centers. A higher anomaly score indicates a greater deviation from the known normal patterns. Samples with scores exceeding a predefined threshold are classified as anomalies. To determine anomaly labels in an unsupervised manner, we adopt a statistical thresholding strategy. Specifically, we apply the empirical 3σ rule (Montgomery 2020), whereby a test sample is flagged as anomalous if its score exceeds three standard deviations above the mean anomaly score observed on the evaluation set. This simple yet effective strategy provides a principle-based decision boundaries without relying on labeled anomalies (Hundman et al. 2018; Blázquez-García et al. 2021).

Discussions. Related clustering-based methods rely on rigid boundaries, predefined cluster structures, or simplistic single-hypersphere assumptions, making them less effective for dynamic, continuous, or multimodal time series. Even multiscale approaches like THOC (Shen, Li, and Kwok

Models	SMD Facility			YAHOO Synthetic			UCR Medical			IOPS WebService			WSD WebService		
	VP	VR	AF	VP	VR	AF	VP	VR	AF	VP	VR	AF	VP	VR	AF
PCA	0.482	0.821	0.666	0.024	0.370	NaN	0.724	0.998	0.996	0.202	0.660	0.291	0.109	0.878	0.222
KNN	0.023	0.837	NaN	0.281	0.948	NaN	0.856	1.000	0.994	0.222	0.957	0.531	0.011	0.035	NaN
IForest	0.407	0.804	0.650	0.027	0.832	0.263	0.025	0.876	NaN	0.335	0.917	0.478	0.024	0.614	0.173
MP	0.024	0.845	0.214	0.347	0.955	NaN	0.830	0.999	0.996	0.245	0.961	0.479	0.012	0.045	NaN
KShapeAD	0.019	0.799	NaN	0.048	0.383	NaN	0.998	1.000	0.995	0.121	0.816	0.298	0.010	0.006	NaN
GBMOD	0.510	0.845	0.931	0.221	0.625	0.586	0.810	0.921	0.921	0.221	0.625	0.583	0.301	0.801	0.782
GBDO	0.598	0.878	0.941	0.285	0.701	0.608	0.882	0.955	0.930	0.265	0.656	0.592	0.355	0.833	0.798
CNN	0.755	0.979	0.999	0.044	0.898	0.626	0.299	0.949	0.772	0.344	0.949	0.669	0.609	0.921	0.981
LSTMAD	0.782	0.991	0.999	0.087	0.959	0.413	0.151	0.981	0.691	0.386	0.983	0.916	0.249	0.817	0.834
TranAD	0.364	0.896	0.661	0.091	0.967	0.415	0.015	0.788	0.590	0.307	0.958	0.745	0.122	0.721	0.711
USAD	0.369	0.821	0.669	0.123	0.928	0.157	0.020	0.854	0.634	0.314	0.949	0.663	0.093	0.807	0.497
TimesNet	0.680	0.961	0.997	0.577	0.980	0.881	0.023	0.875	0.653	0.184	0.906	0.811	0.354	0.983	0.805
DeepSVDD	0.812	0.998	0.866	0.967	0.896	0.071	0.996	1.000	0.934	0.236	0.895	0.322	0.404	0.982	0.698
A.T.	0.017	0.494	0.401	0.098	0.864	0.651	0.135	0.755	0.655	0.124	0.891	0.685	0.354	0.393	0.511
THOC	0.272	0.891	0.663	0.048	0.698	0.275	0.513	0.979	0.836	0.407	0.906	0.331	0.025	0.555	0.092
MEMTO	0.314	0.962	0.663	0.074	0.868	0.393	0.630	0.878	0.453	0.180	0.912	0.743	0.021	0.501	NaN
GBOC	0.831	0.999	0.999	0.991	1.000	0.950	0.996	0.999	0.996	0.604	0.992	0.948	0.963	0.998	0.995

Table 1: Results on (univariate) time series anomaly detection datasets (VP: VUS-PR, VR: VUS-ROC, AF: Affiliation-F1).

2020), which employ hierarchical clustering for temporal features, are limited by fixed structures, reducing adaptability to density and temporal variations.

In contrast, the proposed GBOC dynamically forms granular-ball structures that adaptively capture both local density and global continuity without requiring predefined parameters or rigid boundaries. This ensures robust and effective anomaly detection in complex and noisy real-world time series distributions.

4 Experiment

In this section, we analyze the experimental results of the proposed method on seven widely used time series anomaly detection datasets and compare it with several state-of-the-art methods to demonstrate its effectiveness.

4.1 Experimental Settings

Datasets. We include both univariate and multivariate time series anomaly detection benchmark datasets. These datasets span a broad range of real-world domains, including industrial systems (SMD (Su et al. 2019)), web services (IOPS (Liu and Paparrizos 2024), WSD (Zhang et al. 2022)), healthcare (UCR (Wu and Keogh 2021), LTDB (Goldberger et al. 2000), SVDB (Greenwald, Patil, and Mark 1990)), environmental sensing (SMAP (Hundman et al. 2018), MSL (Hundman et al. 2018)), and synthetic monitoring (YAHOO (Laptev, Amizadeh, and Billawala 2015)). This diversity ensures that our evaluation captures both point and range-based anomalies, as well as varied temporal dynamics (e.g., high nonlinearity with noise, complex temporal variations, and distribution shifts) and different data characteristics. For all datasets, we follow the official train/tune/test splits provided by (Liu and Paparrizos 2024) to ensure fairness. See Appendix C.1 for more detailed descriptions of the datasets.

Compared Methods. We compare our method with 16 baselines, which are divided into two main groups: non-deep learning methods and deep learning methods. (i) The former baselines include PCA (Aggarwal 2017), KNN (Ramswamy, Rastogi, and Shim 2000), IForest (Liu, Ting, and Zhou 2008), MatrixProfile (Yeh et al. 2016), KShapeAD (Paparrizos and Gravano 2017), GBMOD (Cheng et al. 2025) and GBDO (Su et al. 2025b). (ii) The latter includes CNN (Munir et al. 2018), LSTMAD (Malhotra et al. 2015), TranAD (Tuli, Casale, and Jennings 2022), USAD (Audibert et al. 2020), TimesNet (Wu et al. 2023) and AnomalyTransformer (A.T.) (Xu et al. 2022). Moreover, we also include three competitive models such as DeepSVDD (Ruff et al. 2018), THOC (Shen, Li, and Kwok 2020), and MEMTO (Song et al. 2023). See Appendix C.2 for more detailed descriptions of the baselines.

Implementation Details. Following the protocols in (Liu and Paparrizos 2024), random seed is set for 2024, we tune the window size and the number of LSTM encoder layers for the proposed GBOC model based on validation performance. The window size is selected from 2, 5, 10, 50 based on the temporal resolution and dataset length, while the number of LSTM layers are set between 1 and 3, depending on input complexity. The hidden layer dimension is fixed at 32, and the loss weight λ is set to 0.5 by default. These settings remain consistent across datasets unless stated otherwise. All models are trained using the Adam optimizer with a learning rate of 10^{-4} and a batch size of 32. The hyperparameters of the baselines are shown in Appendix C.3. All experiments were conducted on a workstation equipped with an NVIDIA RTX 4090 GPU and 128 GB of RAM, using Python 3.10 and PyTorch 1.13.

Evaluation Metrics. We evaluate anomaly detection models using three metrics: VUS-PR (Paparrizos et al.

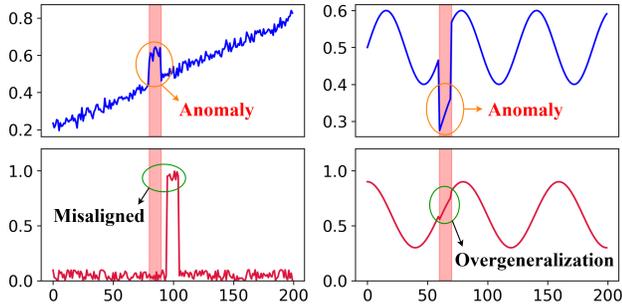


Figure 3: (a) shows high VUS but low Affiliation-F1 due to misaligned anomaly scores. (b) shows high VUS but NaN F1 as no predicted scores overlap with anomalies.

2022), VUS-ROC (Paparrizos et al. 2022), and Affiliation-F1 (Huet, Navarro, and Rossi 2022), reflecting ranking quality, temporal robustness, and range-based alignment. VUS-PR, based on the precision-recall manifold, is more suitable for rare-event scenarios than AUC-PR (Davis and Goadrich 2006). VUS-ROC introduces a tolerance window Δ to account for detection delays, treating nearby predictions as correct. Affiliation-F1 evaluates the alignment between predicted anomalies and ground truth intervals using a Gaussian distance kernel. However, some baselines achieve high VUS-PR or VUS-ROC scores but yield NaN in Affiliation-F1 when predictions fail to overlap with ground truth, as shown in Figure 3, we also provide a detailed analysis in Appendix F.1. Thus, these metrics complement each other to provide full insight into model performance.

4.2 Main Results

Univariate Detection. As shown in Table 1, GBOC consistently achieves strong performance on univariate time series anomaly detection tasks, securing the highest overall average score. It excels particularly on datasets such as SMD, YAHOO, IOPS, and WSD, while maintaining competitive accuracy across other benchmarks. These results highlight GBOC’s ability to handle diverse univariate anomaly detection challenges effectively.

Multivariate Detection. In Table 2 of Section 4.1 of the extended version, the proposed GBOC demonstrates exceptional performance in multivariate anomaly detection, outperforming several state-of-the-art baselines, including TimesNet, A.T., DeepSVDD, THOC, and MEMTO. Its robust results across various datasets reinforce its versatility and effectiveness in identifying anomalies in complex multivariate settings.

4.3 Robust Detection under Drift and Noise

This section examines four scenarios with different temporal dynamics: Type I (Clean, no drift or noise), Type II (Drift-only), Type III (Noise-only), and Type IV (Drift and noise). Table 2 summarizes the results:

Type I (Clean): Nearly all models perform well, with GBOC and KShapeAD achieving near-perfect accuracy, reflecting their ability to handle well-structured data.

Methods	I	II	III	IV
PCA	0.955	0.005	0.331	0.083
KNN	0.991	0.586	0.591	0.442
IForest	0.217	0.029	0.082	0.065
MP	0.995	0.594	0.640	0.450
KShapeAD	<u>1.000</u>	0.982	0.802	0.624
CNN	0.502	0.013	0.230	0.069
LSTMAD	0.503	0.005	0.280	0.143
TranAD	0.455	0.005	0.101	0.173
USAD	0.824	0.004	0.061	0.103
TimesNet	0.747	0.819	0.037	0.042
A.T.	0.479	0.011	0.532	0.474
DeepSVDD	0.824	0.153	<u>0.833</u>	<u>0.893</u>
THOC	0.569	0.006	0.088	0.075
MEMTO	0.782	0.028	0.121	0.031
GBOC	1.000	<u>0.977</u>	0.952	0.921

Table 2: VUS-PR results on four real-world anomaly types datasets from MSL (I), SMAP (II), and YAHOO (III, IV).

Type II (Drift-only): Models that rely on stationary assumptions, such as PCA, LSTM, and CNN, suffer significant performance declines due to their inability to adapt to distributional shifts. In contrast, GBOC and KShapeAD demonstrate strong adaptability, maintaining high performance despite the drift.

Type III (Noise-only): GBOC emerges as the most robust, leveraging its granular-ball-based modeling to isolate dense, high-quality regions and minimize false positives. This gives GBOC a significant edge over others. KShapeAD shows reasonable robustness but still experiences a marked performance drop due to its sensitivity to noise.

Type IV (Drift and noise): This is the most challenging scenario, where all models face severe degradation. However, GBOC remains the top performer, showcasing its exceptional ability to generalize under compounded corruption. DeepSVDD demonstrates commendable resilience, but most other baselines, including traditional clustering and KNN methods, struggle heavily.

Overall, GBOC demonstrates consistent top performance across scenarios, particularly in noisy or nonstationary environments, where traditional methods like KShapeAD degrade significantly.

4.4 Ablation Studies

In this section, we perform ablation studies to analyze: (i) the impact of granular-ball data description and (ii) the roles of losses \mathcal{L}_{rec} and \mathcal{L}_{gb} in Equation (8). More ablation study results are provided in Appendix E.2.

Granular-ball Data Description. Two GBOC variants were tested: (1) *w/o granular-ball computing*: Replaces granular-ball construction with K-Means. (2) *w/o granular-ball pruning*: Retains all granular-balls without pruning. As shown in Table 3, removing pruning degrades the performance by retaining noisy regions. Using K-Means results in greater drops due to its fixed clusters, making it less adaptable to varying data densities. In contrast, GBOC dynam-

GBC	Pruning	SMD	IOPS	UCR	YAHOO	WSD
×	×	0.755	0.554	0.921	0.823	0.911
✓	×	0.781	0.566	0.972	0.795	0.885
✓	✓	0.831	0.604	0.996	0.991	0.963

Table 3: Effects of *w/o granular-ball computing* and *w/o pruning* regarding VUS-PR.

\mathcal{L}_{rec}	\mathcal{L}_{gb}	SMD	IOPS	UCR	YAHOO	WSD
✓	×	0.780	0.545	0.955	0.869	0.936
×	✓	0.715	0.506	0.939	0.701	0.945
✓	✓	0.831	0.604	0.996	0.991	0.963

Table 4: Effects of \mathcal{L}_{rec} and \mathcal{L}_{gb} regarding VUS-PR .

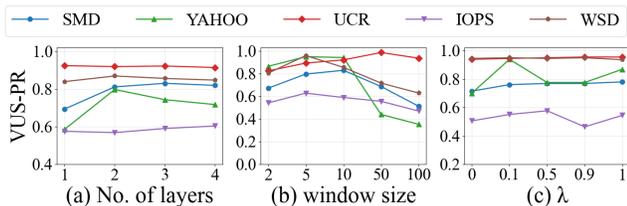


Figure 4: Sensitivity analysis of GBOC on VUS-PR to the number of LSTM encoder layers, input window size and loss weight λ across multiple datasets.

cally adjusts to local densities, enabling superior generalization to complex and irregular data distributions.

Effects of \mathcal{L}_{rec} and \mathcal{L}_{gb} . We evaluated: (1) *Reconstruction loss \mathcal{L}_{rec}* : Preserves input structure in latent representations. (2) *Granular-ball optimization loss \mathcal{L}_{gb}* : Aligns embeddings with granular-ball centers. As shown in Table 4, removing \mathcal{L}_{gb} leads to a drop in VUS-PR values due to disorganized embeddings and reduced discriminative ability, while excluding \mathcal{L}_{rec} compromises semantic coherence. These effects are particularly evident on the noisy YAHOO dataset, underscoring the importance of both losses.

4.5 Hyperparameter Sensitivity Analysis

In this section, sensitivity analysis regarding VUS-PR on three hyperparameters are performed. They are (i) number of LSTM encoder layers, (ii) size of sliding window and (iii) loss weight λ in Equation 8. More results are provided in Appendix E.3.

As shown in Figure 4 (left), the number of layers is a relatively insensitive hyperparameter. In most cases, a two-layer LSTM network achieves satisfactory performance.

In Figure 4 (middle), the window size demonstrates varying sensitivity across different datasets. For example, a small window size performs better on YAHOO dataset, while a window size of 50 yields better results for UCR dataset.

By varying λ from 0 to 1 in Figure 4 (right), we observe that selecting an appropriate value for λ achieves better results compared to $\lambda = 0$ or $\lambda = 1$, demonstrating the effectiveness of this trade-off parameter.

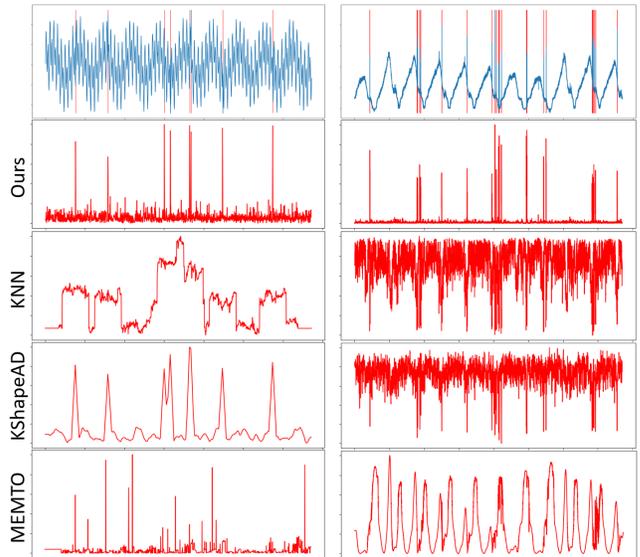


Figure 5: Visualization of anomaly scores on YAHOO (left) and WSD (right). Ground-truth anomalies are shaded in red.

4.6 Visualization Analysis

To assess GBOC’s robustness and interpretability, we visualize its anomaly scores on YAHOO and WSD, and compare them with three baselines: KNN, KShapeAD, and MEMTO. As shown in Figure 5, GBOC produces smooth, well-localized scores that align closely with the ground truth. In contrast, KShapeAD and KNN often yield fluctuating scores in normal regions with noise or periodic variations, causing false positives. MEMTO, while more stable, tends to blur anomaly boundaries and mislabel extended normal regions. GBOC suppresses spurious spikes caused by noise, thanks to its compact, high-density granular-ball memory, which filters noise at the representation level. These results confirm GBOC’s quantitative advantages and highlight its interpretability and consistency across diverse scenarios. See Appendix F.3 for more visualization results of anomaly scores.

5 Conclusion

In this work, we proposed GBOC, a novel one-class network for time series anomaly detection based on adaptive granular-ball vector data description (GVDD). By organizing latent representations into compact, high-density granular-balls, GBOC captures local structure effectively and provides a flexible, interpretable one-class data description. Additionally, we introduced a granular-ball representation optimization mechanism. Experiments on diverse benchmarks show that GBOC achieves competitive or superior performance, particularly excelling under noisy and nonstationary conditions. Ablation studies validated the contributions of its components, while case studies highlight its adaptability across anomaly scenarios. In conclusion, GBOC achieves strong detection performance with flexible distribution description, offering a promising direction for unsupervised time series analysis.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant Nos. 62221005, 62450043, 62222601, 62176033 and 62576056, as well as by Ant Group, and Chongqing Ant Consumer Finance Co.

References

- Abdulaal, A.; Liu, Z.; and Lancewicki, T. 2021. Practical approach to asynchronous multivariate time series anomaly detection and localization. In *Proceedings of the 27th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2485–2494.
- Aggarwal, C. C. 2017. An introduction to outlier analysis. *Springer New York*.
- Audibert, J.; Michiardi, P.; Guyard, F.; Marti, S.; and Zuluaga, M. A. 2020. USAD: Unsupervised anomaly detection on multivariate time series. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 3395–3404.
- Blázquez-García, A.; Conde, A.; Mori, U.; and Lozano, J. A. 2021. A review on outlier/anomaly detection in time series data. *ACM computing surveys (CSUR)*, 1–33.
- Boniol, P.; Paparrizos, J.; Palpanas, T.; and Franklin, M. J. 2021. SAND: Streaming subsequence anomaly detection. *Proceedings of the VLDB Endowment*, 1717–1729.
- Chen, L. 1982. Topological structure in visual perception. *Science*, 218(4573): 699–700.
- Cheng, D.; Li, Y.; Xia, S.; Wang, G.; Huang, J.; and Zhang, S. 2023. A fast granular-ball-based density peaks clustering algorithm for large-scale data. *IEEE Transactions on Neural Networks and Learning Systems*.
- Cheng, S.; Su, X.; Chen, B.; Chen, H.; Peng, D.; and Yuan, Z. 2025. GBMOD: A granular-ball mean-shift outlier detector. *Pattern Recognition*, 111115.
- Davis, J.; and Goadrich, M. 2006. The relationship between precision-recall and ROC curves. In *International Conference on Machine Learning*, 233–240.
- Goldberger, A. L.; Amaral, L. A.; Glass, L.; Hausdorff, J. M.; Ivanov, P. C.; Mark, R. G.; Mietus, J. E.; Moody, G. B.; Peng, C.-K.; and Stanley, H. E. 2000. PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. *Circulation*, e215–e220.
- Greenwald, S. D.; Patil, R. S.; and Mark, R. G. 1990. *Improved detection and classification of arrhythmias in noise-corrupted electrocardiograms using contextual information*. Ph.D. thesis, Massachusetts Institute of Technology.
- Hu, L.; Chen, F.; Zhao, S.; Duan, S.; et al. 2025. GRICP: Granular-ball iterative closest point with multikernel correntropy for point cloud fine registration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1710–1718.
- Huet, A.; Navarro, J. M.; and Rossi, D. 2022. Local evaluation of time series anomaly detection algorithms. In *Proceedings of the 28th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 635–645.
- Hundman, K.; Constantinou, V.; Laporte, C.; Colwell, I.; and Soderstrom, T. 2018. Detecting spacecraft anomalies using lstms and nonparametric dynamic thresholding. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 387–395.
- Laptev, N.; Amizadeh, S.; and Billawala, Y. 2015. S5-A labeled anomaly detection dataset, version 1.0 (16M)(2015). URL <https://webscope.sandbox.yahoo.com/catalog.php>.
- Li, Y.; Ouyang, X.; Pan, C.; Zhang, J.; Zhao, S.; Xia, S.; Yang, X.; Wang, G.; and Li, T. 2025. Multi-granularity open intent classification via adaptive granular-ball decision boundary. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 24512–24520.
- Liu, F. T.; Ting, K. M.; and Zhou, Z.-H. 2008. Isolation forest. In *IEEE International Conference on Data Mining*, 413–422.
- Liu, L.; Liu, D.; and Peng, Y. 2016. Detection and identification of sensor anomaly for aerospace applications. In *2016 Annual Reliability and Maintainability Symposium (RAMS)*, 1–6.
- Liu, Q.; and Paparrizos, J. 2024. The elephant in the room: Towards a reliable time-series anomaly detection benchmark. *Advances in Neural Information Processing Systems*, 108231–108261.
- Malhotra, P.; Vig, L.; Shroff, G. M.; and Agarwal, P. 2015. Long short term memory networks for anomaly detection in time series. In *The European Symposium on Artificial Neural Networks*.
- Mathur, A. P.; and Tippenhauer, N. O. 2016. 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)*, 31–36.
- Montgomery, D. C. 2020. *Introduction to statistical quality control*. John Wiley & Sons.
- Munir, M.; Siddiqui, S. A.; Dengel, A.; and Ahmed, S. 2018. DeepAnT: A deep learning approach for unsupervised anomaly detection in time series. *IEEE Access*, 1991–2005.
- Paparrizos, J.; Boniol, P.; Palpanas, T.; Tsay, R. S.; Elmore, A.; and Franklin, M. J. 2022. Volume under the surface: a new accuracy evaluation measure for time-series anomaly detection. *Proceedings of the VLDB Endowment*, 2774–2787.
- Paparrizos, J.; and Gravano, L. 2015. k-Shape: Efficient and accurate clustering of time series. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, 1855–1870.
- Paparrizos, J.; and Gravano, L. 2017. Fast and accurate time-series clustering. *ACM Transactions on Database Systems (TODS)*, 1–49.
- Park, D.; Hoshi, Y.; and Kemp, C. C. 2018. A multimodal anomaly detector for robot-assisted feeding using an lstm-based variational autoencoder. *IEEE Robotics and Automation Letters*, 1544–1551.
- Ramaswamy, S.; Rastogi, R.; and Shim, K. 2000. Efficient algorithms for mining outliers from large data sets. In *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, 427–438.

- Rassam, M. A.; Maarof, M. A.; and Zainal, A. 2018. A distributed anomaly detection model for wireless sensor networks based on the one-class principal component classifier. *International Journal of Sensor Networks*, 200–214.
- Ruff, L.; Vandermeulen, R.; Goernitz, N.; Deecke, L.; Siddiqui, S. A.; Binder, A.; Müller, E.; and Kloft, M. 2018. Deep one-class classification. In *International Conference on Machine Learning*, 4393–4402.
- Schmidl, S.; Wenig, P.; and Papenbrock, T. 2022. Anomaly detection in time series: A comprehensive evaluation. *Proceedings of the VLDB Endowment*, 1779–1797.
- Schölkopf, B.; Platt, J. C.; Shawe-Taylor, J.; Smola, A. J.; and Williamson, R. C. 2001. Estimating the support of a high-dimensional distribution. *Neural Computation*, 1443–1471.
- Shen, L.; Li, Z.; and Kwok, J. 2020. Timeseries anomaly detection using temporal hierarchical one-class network. *Advances in Neural Information Processing Systems*, 13016–13026.
- Song, J.; Kim, K.; Oh, J.; and Cho, S. 2023. MEMTO: Memory-guided transformer for multivariate time series anomaly detection. *Advances in Neural Information Processing Systems*, 57947–57963.
- Su, P.; Huang, S.; Ma, W.; Xiong, D.; and Lv, J. 2025a. Multi-view granular-ball contrastive clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 20637–20645.
- Su, X.; Wang, X.; Peng, D.; Song, X.; Zheng, H.; and Yuan, Z. 2025b. Identifying outliers via local granular-ball density. *IEEE Transactions on Neural Networks and Learning Systems*.
- Su, Y.; Zhao, Y.; Niu, C.; Liu, R.; Sun, W.; and Pei, D. 2019. Robust anomaly detection for multivariate time series through stochastic recurrent neural network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2828–2837.
- Tax, D. M.; and Duin, R. P. 2004. Support vector data description. *Machine Learning*, 45–66.
- Tuli, S.; Casale, G.; and Jennings, N. R. 2022. TranAD: Deep transformer networks for anomaly detection in multivariate time series data. *Proceedings of the VLDB Endowment*, 1201–1214.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- Wang, G. 2017. DGCC: Data-driven granular cognitive computing. *Granular Computing*, 343–355.
- Wang, J.; Kuang, Q.; and Duan, S. 2015. A new online anomaly learning and detection for large-scale service of internet of thing. *Personal and Ubiquitous Computing*, 1021–1031.
- Wu, H.; Hu, T.; Liu, Y.; Zhou, H.; Wang, J.; and Long, M. 2023. TimesNet: Temporal 2d-variation modeling for general time series analysis. In *International Conference on Learning Representations*.
- Wu, R.; and Keogh, E. J. 2021. Current time series anomaly detection benchmarks are flawed and are creating the illusion of progress. *IEEE Transactions on Knowledge and Data Engineering*, 2421–2429.
- Xia, S.; Liu, Y.; Ding, X.; Wang, G.; Yu, H.; and Luo, Y. 2019. Granular ball computing classifiers for efficient, scalable and robust learning. *Information Sciences*, 136–152.
- Xia, S.; Peng, D.; Meng, D.; Zhang, C.; Wang, G.; Giem, E.; Wei, W.; and Chen, Z. 2020. A fast adaptive k-means with no bounds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Xia, S.; Shi, B.; Wang, Y.; Xie, J.; Wang, G.; and Gao, X. 2025a. GBCT: Efficient and adaptive clustering via granular-ball computing for complex data. *IEEE Transactions on Neural Networks and Learning Systems*, 12159–12172.
- Xia, S.; Wang, Y.; Shen, L.; and Wang, G. 2025b. Granular-ball-induced multiple kernel k-means. In *International Joint Conference on Artificial Intelligence*, 6633–6641.
- Xie, J.; Hua, C.; Xia, S.; Cheng, Y.; Wang, G.; and Gao, X. 2024a. W-GBC: An adaptive weighted clustering method based on granular-ball structure. In *IEEE International Conference on Data Engineering*, 914–925.
- Xie, J.; Kong, W.; Xia, S.; Wang, G.; and Gao, X. 2023. An efficient spectral clustering algorithm based on granular-ball. *IEEE Transactions on Knowledge and Data Engineering*, 9743–9753.
- Xie, Q.; Zhang, Q.; Xia, S.; Zhao, F.; Wu, C.; Wang, G.; and Ding, W. 2024b. GBG++: A fast and stable granular ball generation method for classification. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2022–2036.
- Xu, J.; Wu, H.; Wang, J.; and Long, M. 2022. Anomaly transformer: Time series anomaly detection with association discrepancy. In *International Conference on Learning Representations*.
- Yeh, C.-C. M.; Zhu, Y.; Ulanova, L.; Begum, N.; Ding, Y.; Dau, H. A.; Silva, D. F.; Mueen, A.; and Keogh, E. 2016. Matrix profile I: all pairs similarity joins for time series: a unifying view that includes motifs, discords and shapelets. In *IEEE International Conference on Data Mining*, 1317–1322.
- Zhang, S.; Zhong, Z.; Li, D.; Fan, Q.; Sun, Y.; Zhu, M.; Zhang, Y.; Pei, D.; Sun, J.; Liu, Y.; et al. 2022. Efficient kpi anomaly detection through transfer learning for large-scale web services. *IEEE Journal on Selected Areas in Communications*, 2440–2455.