

Commonality in Few: Few-Shot Multimodal Anomaly Detection via Hypergraph-Enhanced Memory

Yuxuan Lin¹, Hanjing Yan², Xuan Tong³, Yang Chang³, Huanzhen Wang³, Ziheng Zhou³,
Shuyong Gao¹, Yan Wang^{4*}, Wenqiang Zhang^{1,3*}

¹Shanghai Key Lab of Intelligent Information Processing, College of Computer Science and Artificial Intelligence, Fudan University

²School of Information Science and Engineering, East China University of Science and Technology

³College of Intelligent Robotics and Advanced Manufacturing, Fudan University

⁴School of Data Science and Engineering, East China Normal University

{yuxuanlin24, xtong23, ychang24, hzwang24, zhoush24}@m.fudan.edu.cn, {wqzhang, sy_gao}@fudan.edu.cn, Y21250025@mail.ecust.edu.cn, yanwang@dase.ecnu.edu.cn

Abstract

Few-shot multimodal industrial anomaly detection is a critical yet underexplored task, offering the ability to quickly adapt to complex industrial scenarios. In few-shot settings, insufficient training samples often fail to cover the diverse patterns present in test samples. This challenge can be mitigated by extracting structural commonality from a small number of training samples. In this paper, we propose a novel few-shot unsupervised multimodal industrial anomaly detection method based on structural commonality, **CIF** (Commonality In Few). To extract intra-class structural information, we employ hypergraphs, which are capable of modeling higher-order correlations, to capture the structural commonality within training samples, and use a memory bank to store this intra-class structural prior. Firstly, we design a semantic-aware hypergraph construction module tailored for single-semantic industrial images, from which we extract common structures to guide the construction of the memory bank. Secondly, we use a training-free hypergraph message passing module to update the visual features of test samples, reducing the distribution gap between test features and features in the memory bank. We further propose a hyperedge-guided memory search module, which utilizes structural information to assist the memory search process and reduce the false positive rate. Experimental results on the MVTec 3D-AD dataset and the Eyecandies dataset show that our method outperforms the state-of-the-art (SOTA) methods in few-shot settings.

Code — <https://github.com/Sunny5250/CIF>

Extended version — <https://arxiv.org/abs/2511.05966>

Introduction

Anomaly detection (AD) is a widely studied problem in computer vision, and industrial AD which this paper focuses on is essential for ensuring product quality and production line stability. Manual detection is inefficient and easily affected by human factors, making it unsuitable for large-scale, high-precision manufacturing. Traditional AD

*Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

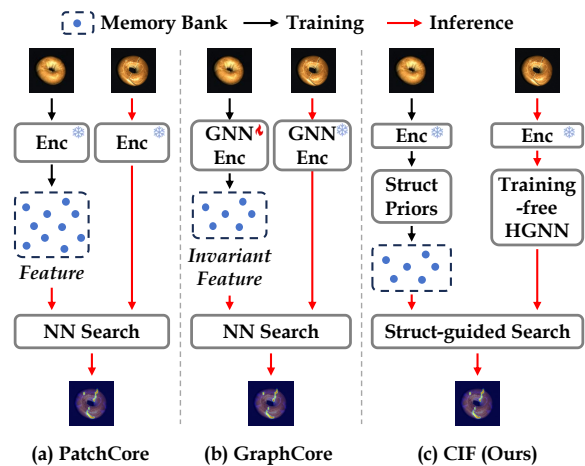


Figure 1: The Main Idea of CIF. Different from (a) PatchCore and (b) GraphCore, our HyperAD (c) extracts higher-order correlations among patch features through a hypergraph and uses training-free message passing to obtain patch features enriched with structured contextual information.

methods (Crosby 1994; Breunig et al. 2000; Tang et al. 2002; Liu, Ting, and Zhou 2008) based on data distribution struggle to detect diverse or subtle anomalies in complex industrial environments and lack robustness. Machine learning and deep learning-based AD methods are mainly categorized into feature embedding-based methods (Gudovskiy, Ishizaka, and Kozuka 2022; Rudolph et al. 2023; Roth et al. 2022; Wang et al. 2023) and reconstruction-based methods (Zavrtnik, Kristan, and Skočaj 2021b,a; Zhang et al. 2023; Chen et al. 2023). However, anomaly samples in real industrial environments are extremely scarce and costly to obtain. As a result, most existing methods are designed in unsupervised settings. By learning the feature distribution of normal samples, unsupervised methods detect anomalies as deviations from normal patterns.

Few-shot AD is used in scenarios with limited normal training samples. Compared to full-shot methods, it offers

stronger transferability and faster adaptability, making it well-suited for complex and dynamic industrial environments. MetaFormer (Wu et al. 2021) uses meta-learning by training on additional datasets and fine-tuning with few normal samples. AnomalyGPT (Gu et al. 2024) uses Large Vision-Language Models (LVLMs) for effective few-shot anomaly detection and localization, benefiting from their rich priors and generalization. Memory bank-based method (Defard et al. 2021; Roth et al. 2022; Lee, Lee, and Song 2022; Wang et al. 2023; Xie et al. 2023) constructs a memory bank of normal features, which are used to compare with test features to detect anomalies. However, in few-shot settings, the memory bank often fails to cover the normal patterns in test samples, resulting in high false positive rates. GraphCore (Xie et al. 2023) improves this by using Graph Neural Networks (GNNs) to extract isometric invariant features, which reduces redundancy in the memory bank, improving the coverage of memory bank features over test features.

For single-semantic industrial images, samples of the same class often share consistent structures. While graphs struggle to capture higher-order correlations between patches, hypergraphs can better model these correlations and extract shared structural information among samples, enhancing memory bank coverage in few-shot settings. In this work, we propose CIF (Commonality In Few), a hypergraph-based few-shot unsupervised multimodal industrial anomaly detection method. We construct hypergraphs from a few training samples to extract shared structures, guiding memory bank construction and compression. We also use training-free hypergraph message passing to update the test features and use hyperedge features to guide memory search, reducing the false positive rate. As shown in Figure 1, PatchCore (Roth et al. 2022) detects anomalies via patch feature memory sampling and nearest-neighbor search. GraphCore (Xie et al. 2023) uses GNNs to aggregate neighborhood information from patches and extract vision isometric invariant features. Our CIF uses hypergraphs to model higher-order correlations among patches, capture structural commonalities across samples, and guide memory search, finally reducing false positive rates.

In summary, the contributions of our work are as follows:

- We propose CIF, which explores the application of hypergraphs in few-shot unsupervised multimodal industrial anomaly detection, and outperforms the state-of-the-art methods on the MVTec 3D-AD and Eyecandies datasets in few-shot settings.
- We introduce Semantic-Aware Hypergraph Construction (SAHC), which constructs a hypergraph with more evenly distributed hyperedges for single-semantic industrial images.
- We design Structure-Guided Memory Sampling (SGMS), which uses intra-class structural commonality to guide the construction and compression of the memory bank.
- We use Bidirectional Training-Free Hypergraph Message Passing (Bi-TF-MP) to reduce the distribution gap between test sample node features and memory bank node features.
- We propose Hyperedge-Guided Memory Search (HGMS), which uses hyperedge features to guide the matching be-

tween test sample features and memory bank features, effectively reducing the false positive rate.

Related Work

Industrial Anomaly Detection

Industrial anomaly detection identifies anomaly patterns in industrial data. Traditional methods (Crosby 1994; Angiulli and Pizzuti 2002; Breunig et al. 2000; Shyu et al. 2003; Liu, Ting, and Zhou 2008; Schölkopf et al. 2001) rely on statistical or distance-based analysis. With advances in deep learning and computer vision, current deep learning-based methods are mainly unsupervised, covering both single-modal and multimodal settings (e.g., RGB images, 3D point clouds). They are generally categorized into feature embedding-based and reconstruction-based methods. Feature embedding-based methods use pretrained extractors to obtain features for detection and segmentation, including teacher-student architecture methods (Bergmann et al. 2020; Zhang et al. 2023; Rudolph et al. 2023), one-class classification methods (Li et al. 2021; Liu et al. 2023), distribution map methods (Yu et al. 2021; Gudovskiy, Ishizaka, and Kozuka 2022), and memory bank methods (Roth et al. 2022; Wang et al. 2023; Xie et al. 2023). Reconstruction-based methods (Zavrtanik, Kristan, and Skočaj 2021a; Zhang, Xu, and Zhou 2024; Chen et al. 2023) reconstruct test samples to normal patterns and detect anomalies via reconstruction differences. Few-shot anomaly detection learns compact, generalizable features from limited data, offering better flexibility and transferability than full-shot approaches. Meta learning-based methods (Wu et al. 2021; Huang et al. 2022) train generalizable models adaptable to new domains with minimal fine-tuning. Recently, large model-based methods (Gu et al. 2024; Lee and Choi 2024; Li et al. 2024; Zhu and Pang 2024) leverage the rich priors and strong perception abilities of large models, achieving powerful few-shot and transfer performance.

Hypergraph Learning

Hypergraphs are a generalized form of graphs that allow an edge (hyperedge) to connect multiple nodes, effectively representing high-order correlations and complex structures. Zhou et al. (Zhou, Huang, and Schölkopf 2006) first introduced hypergraph learning to minimize the label difference among correlational nodes. HGNN (Feng et al. 2019) and its extended version HGNN+ (Gao et al. 2022) designs message passing frameworks between hyperedges and nodes, constructing more general hypergraph neural networks to learn representations of high-order data structures. ViHGNN (Han et al. 2023) applies hypergraph neural networks to computer vision, employing a fuzzy clustering algorithm (fuzzy c-means) to construct hyperedges among image patches and capture high-order correlations, achieving high accuracy in tasks such as image classification and object detection. Based on ViHGNN, DVHGNN (Li et al. 2025) introduces dilated hyperedges to capture multi-scale sparse dependencies. While these hypergraph learning methods are generalizable across various tasks, they require adjustments to adapt to industrial data. Therefore, using hypergraphs for industrial anomaly detection remains a challenge.

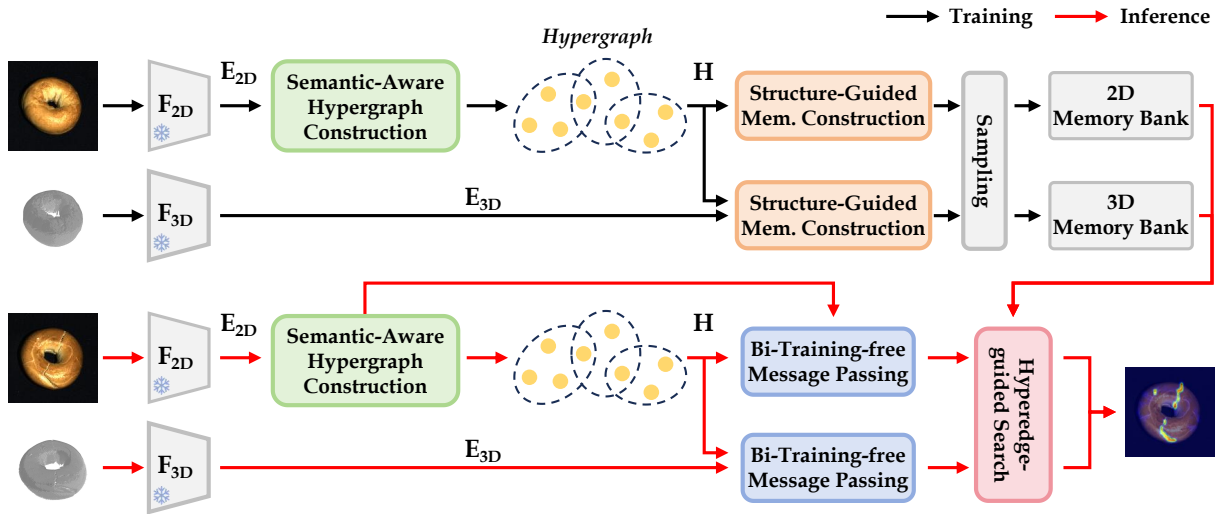


Figure 2: **The pipeline of CIF.** Our CIF contains four important parts: (1) Semantic-Aware Hypergraph Construction (SAHC), which constructs hypergraphs for single-semantic samples based on clustering. (2) Structure-Guided Memory Sampling (SGMS), which uses structural commonality in the hypergraph to guide the construction and compression of the memory bank. (3) Bidirectional Training-Free Hypergraph Message Passing (Bi-TF-MP), which performs bidirectional message passing between test samples and the memory bank to reduce their distribution gap. (4) Hyperedge-Guided Memory Search (HGMS), which uses a hyperedge-guided two-stage search to reduce the false positive rate of detection.

Method

Preliminary: Hypergraph Definition

In traditional graph theory, a graph $G = (V, E)$ represents binary correlations between entities, where V is the set of nodes and $E \subseteq V \times V$ is the set of edges connecting pairs of nodes. However, such binary correlations often struggle to capture the complex higher-order interactions among multiple entities. To address this issue, a hypergraph provides a structure capable of expressing higher-order correlations. As a generalization of a graph, a hypergraph is defined as $\mathcal{G} = (V, \mathcal{E})$, where V is the set of nodes (vertices) and \mathcal{E} is the set of hyperedges. Each hyperedge $e \in \mathcal{E}$ is a non-empty subset of nodes, i.e., $e \subseteq V$, with $|e| \geq 2$. Unlike an ordinary graph edge that connects exactly two nodes, a hyperedge can simultaneously connect an arbitrary number of nodes, enabling the modeling of higher-order correlations among multiple entities. A hypergraph is commonly represented by an incidence matrix $\mathbf{H} \in \mathbb{R}^{|V| \times |\mathcal{E}|}$, where $\mathbf{H}(v, e) = 1$ if vertex v belongs to hyperedge e , and 0 otherwise. Based on the incidence matrix, various hypergraph Laplacian operators can be constructed to support analysis and learning tasks.

Overview

Our CIF is a memory bank-based anomaly detection method that enhances memory construction and search through hypergraph-based structural representations. Unlike other memory bank-based methods, we use intra-class structural priors from training samples to guide the construction and compression of the memory bank. Moreover, we do not directly perform nearest-neighbor search to match the features of test samples and those in the memory bank. Instead, we use hyperedge features to match the structural information of the

test samples with that in the memory bank, and then perform nearest-neighbor search guided by this structural alignment. To ensure structural consistency between the memory bank and test samples in few-shot settings, we apply bidirectional training-free hypergraph message passing to update the features of test samples, thereby narrowing the distribution gap between them and the memory bank features. We use two pretrained feature extractors: DINO (Caron et al. 2021) for extracting 2D image features, and PointMAE (Pang et al. 2022) for extracting 3D point cloud features.

Semantic-Aware Hypergraph Construction

Traditional hard clustering algorithms (e.g., K-Means) cannot assign nodes to multiple hyperedges, while fuzzy clustering (e.g., Fuzzy C-Means (Bezdek, Ehrlich, and Full 1984)) results in uneven hyperedges when applied to single-semantic industrial images. We give a comparison of different hypergraph construction methods in extended version. We propose a **Semantic-Aware Hypergraph Construction (SAHC)** method, which first generates a set of centers (hyperedges) using K-Means. Then, we calculate the similarity between patch (node) features and the centers, and determine the assignment of nodes to hyperedges based on a predefined threshold.

For image data, we use a pretrained feature extractor to obtain patch features $X = [x_1, x_2, \dots, x_N]$, where each patch feature can be viewed as a node in the set $V = \{v_1, v_2, \dots, v_N\}$. To reduce the interference of background information during hypergraph construction, we extract the foreground mask of each sample using its 3D point cloud, and filter out the foreground nodes from V to obtain the foreground node set V_{fore} . As shown in Figure 3, we then perform K-means clustering on V_{fore} to obtain $|\mathcal{E}|$ cluster centers,




Figure 3: The pipeline of Hypergraph construction. We use a clustering algorithm to compute the cluster centers among node features, which serve as hyperedge centers. We then calculate the similarity between each hyperedge center and all foreground node features, and determine the hyperedge membership of each node using a predefined threshold, resulting in the hypergraph incidence matrix.




Figure 4: Comparison: Greedy Coreset Sampling (Top) vs. Our Sampling (Bottom) at different sampling rates (SR). Red points represent features before sampling, blue points represent features after sampling, and black crosses represent hyperedge features.

which serve as hyperedge centers, forming the hyperedge set \mathcal{E} . We calculate the cosine similarity between all node features in V_{fore} and the hyperedge centers, and apply min-max normalization. Using the similarity matrix, we determine the assignment of each node to one or more hyperedges based on a predefined threshold, resulting in an incidence matrix $\mathbf{H} \in \mathbb{R}^{|V| \times |\mathcal{E}|}$. Due to the needs of subsequent modules, we also compute a hard incidence matrix $\mathbf{H}_{\text{hard}} \in \mathbb{R}^{|V| \times 1}$, where each node is assigned only to one hyperedge with the highest similarity.

During the hypergraph construction process, we observed that hypergraphs built from RGB image features of industrial samples are more balanced and reliable, while those constructed from 3D point cloud features exhibit significant imbalance. Therefore, we uniformly use hypergraphs constructed from RGB image features to support the subsequent modules in both the 2D and 3D modals.

Structure-Guided Memory Sampling

For single-semantic industrial images, samples of the same class often exhibit consistent structural patterns. Such intra-class structural priors can be used to guide the matching between memory bank features and test sample features, and can alleviate the issue of sparse information in the memory bank. Using SAHC, we obtain the incidence matrices and hyperedge features of training samples, and then design a **Structure-Guided Memory Sampling (SGMS)** module to update and compress the memory bank. We fix the number of hyperedges per class and use the patch features, incidence matrix, and hyperedge features of the first training sample to initialize the memory bank for this class. The memory bank

update process consists of three stages: node assignment, hyperedge update, and memory bank sampling.

In the node assignment stage, we first compute the average patch features for each hyperedge in the memory bank to obtain the corresponding hyperedge features. For a newly input training sample, we compute its hyperedge features and calculate the pairwise distances between these and the hyperedge features in the memory bank. To facilitate node assignment, we use \mathbf{H}_{hard} as the reference for node-hyperedge membership. Based on this, we merge the nodes from each hyperedge in the training sample with the most similar hyperedge in the memory bank. In the hyperedge update stage, after assigning the node features of a new sample, we update the hyperedge features of the memory bank by computing the merged hyperedge features, preparing them for the next node assignment step. In the memory bank sampling stage, existing methods (Roth et al. 2022; Xie et al. 2023) perform greedy coreset sampling across all nodes in the memory bank. In contrast, we perform greedy coreset sampling within each hyperedge individually. If a hyperedge contains fewer than one node after sampling, we select a single node from this hyperedge such that the maximum distance from this node to all other nodes in the hyperedge is minimized. Figure 4 visually compares our sampling method with greedy coreset sampling.

Bidirectional Training-Free Hypergraph Message Passing

To mitigate the distribution gap between test sample node features and memory bank node features caused by the limited number of features in the memory bank, we leverage intra-class structural priors from the memory bank and the structural information of the test sample. We propose a **Bidirectional Training-Free hypergraph Message Passing (Bi-TF-MP)** module, which constructs a joint hypergraph between test sample node features and memory bank node features to explicitly connect them. We then perform L layers of training-free hypergraph message passing on this joint hypergraph, enabling information exchange between each node and its L -hop neighbors in one step.

We denote the test hypergraph as $\mathcal{G}^{\text{test}} = (X^{\text{test}} \in \mathbb{R}^{N \times \text{dim}}, \mathbf{H}^{\text{test}} \in \{0, 1\}^{N \times \mathcal{E}_{\text{num}}})$ and the memory bank hypergraph as $\mathcal{G}^{\text{mem}} = (X^{\text{mem}} \in \mathbb{R}^{A \times \text{dim}}, \mathbf{H}^{\text{mem}} \in \{0, 1\}^{A \times \mathcal{E}_{\text{num}}})$, where dim is the feature dimension of each node, \mathbf{H}^* is the incidence matrix, N is the number of nodes in the test sample, and A is the number of nodes in the memory bank.

We compute the cosine similarity between X^{test} and X^{mem} , and for each test node x_i^{test} , we select the top- k most similar memory bank nodes. We construct a hyperedge that includes x_i^{test} and the selected k memory bank nodes. In total, N such cross-domain hyperedges are constructed, resulting in an incidence matrix $\mathbf{H}^{t \rightarrow m} \in \{0, 1\}^{(A+N) \times N}$. Similarly, for each memory node x_i^{mem} , we select the top- k most similar test nodes, and construct a hyperedge that includes x_i^{mem} and the selected k test nodes. In total, A such cross-domain hyperedges are constructed, resulting in an incidence matrix $\mathbf{H}^{m \rightarrow t} \in \{0, 1\}^{(A+N) \times A}$. By concatenating $\mathbf{H}^{t \rightarrow m}$ and $\mathbf{H}^{m \rightarrow t}$ along the hyperedge dimension, we obtain the cross-domain incidence matrix $\mathbf{H}^{cross} \in \{0, 1\}^{(A+N) \times (A+N)}$.

To enable message passing across all three domains (test sample domain, memory bank domain, and cross-domain), we apply zero-padding to extend $\mathbf{H}^{test} \in \{0, 1\}^{N \times \mathcal{E}_{num}}$ to $\tilde{\mathbf{H}}^{test} \in \{0, 1\}^{(N+A) \times \mathcal{E}_{num}}$, and similarly extend $\mathbf{H}^{mem} \in \{0, 1\}^{A \times \mathcal{E}_{num}}$ to $\tilde{\mathbf{H}}^{mem} \in \{0, 1\}^{(N+A) \times \mathcal{E}_{num}}$. We then concatenate $\tilde{\mathbf{H}}^{test}$, $\tilde{\mathbf{H}}^{mem}$, and \mathbf{H}^{cross} along the hyperedge dimension to obtain the final joint incidence matrix \mathbf{H}^{joint} .

$$\mathbf{H}^{joint} = \begin{bmatrix} \underbrace{\tilde{\mathbf{H}}^{test}}_{test} & \underbrace{\tilde{\mathbf{H}}^{mem}}_{memory} & \underbrace{\mathbf{H}^{cross}}_{cross} \end{bmatrix} \in \{0, 1\}^{(N+A) \times (2\mathcal{E}_{num} + N + A)} \quad (1)$$

then we concatenate the test sample node features X^{test} and the memory bank node features X^{mem} to obtain the joint node features X^{joint} :

$$X^{joint} = \begin{bmatrix} \underbrace{X^{test}}_{test} & \underbrace{X^{mem}}_{memory} \end{bmatrix} \in \mathbb{R}^{(N+A) \times dim} \quad (2)$$

We follow TF-MP (Tang et al. 2024) to perform message passing on the joint hypergraph. For a hypergraph $\mathcal{G} = (X, \mathbf{H})$, we first compute the transition matrix \mathbf{A}_{trans} , and then compute the training-free message passing kernel \mathbf{S} as follows:

$$\mathbf{S} = (1 - \alpha)^L \mathbf{A}_{trans}^L + \alpha \sum_{l=0}^{L-1} (1 - \alpha)^l \mathbf{A}_{trans}^l, \quad (3)$$

$$\mathbf{A}_{trans} = \mathbf{D}_v^{-\frac{1}{2}} (\mathbf{H} \mathbf{D}_e^{-1} \mathbf{H}^\top + \mathbf{I}_N) \mathbf{D}_v^{-\frac{1}{2}}$$

where \mathbf{D}_e is the hyperedge degree matrix of \mathbf{H} , \mathbf{D}_v is the node degree matrix of $\tilde{\mathbf{W}} = \mathbf{H} \mathbf{D}_e^{-1} \mathbf{H}^\top + \mathbf{I}_N$, and \mathbf{I}_N is the identity matrix. Moreover, α is the retention coefficient, a larger α emphasizes the preservation of the own information of a node, while a smaller α increases the influence of neighboring nodes. L denotes the number of message passing layers, allowing any node to exchange information with its L -hop neighbors in the hypergraph.

We apply the training-free message passing kernel \mathbf{S} to the joint hypergraph $\mathcal{G}^{joint} = (X^{joint} \in \mathbb{R}^{(N+A) \times dim}, \mathbf{H}^{joint} \in \{0, 1\}^{(N+A) \times (2\mathcal{E}_{num} + N + A)})$ to obtain the updated node features of test samples, denoted as X_{new}^{test} .

$$\begin{aligned} X_{new}^{joint} &= (\mathbf{S} X^{joint})^\top \in \mathbb{R}^{(N+A) \times dim} \implies \\ X_{new}^{test} &= X_{new}^{joint}[:, :N] \in \mathbb{R}^{N \times dim} \end{aligned} \quad (4)$$

Hyperedge-Guided Memory Search

To match structurally consistent test sample nodes and memory bank nodes based on intra-class structural priors, we design a **Hyperedge-Guided Memory Search (HGMS)** module that introduces hyperedge features to assist in matching test sample features with those in the memory bank, enabling more targeted search based on intra-class structural priors. Specifically, we compute pairwise cosine similarity between the updated hyperedge features (obtained through Bi-TF-MP) of the test sample and the hyperedge features in the memory bank. For the i -th hyperedge in the test sample, we collect all nodes within this hyperedge to form a test node subset $X_{sub\ i}^{test}$. We then select the top- k most similar hyperedges from the memory bank and collect all memory nodes within these k hyperedges to form a memory bank node subset $\mathcal{M}_{sub\ i}$. A patch-level nearest neighbor search is then performed between $X_{sub\ i}^{test}$ and $\mathcal{M}_{sub\ i}$ to compute the patch-level anomaly score \mathcal{A} for the test sample:

$$\mathcal{A}_{ij} = \min_{m \in \mathcal{M}_{sub\ i}} \|X_{ij}^{test} - m\|_2 \quad (5)$$

where \mathcal{A}_{ij} and X_{ij}^{test} denote the anomaly score and patch feature for the j -th patch within the i -th hyperedge of the test sample, respectively. m represents a node feature from the memory bank node subset $\mathcal{M}_{sub\ i}$ corresponding to the i -th hyperedge in the test sample. Additionally, we retain the conventional patch-level search method that directly compares test sample node features with memory bank node features. The resulting scores are element-wise multiplied with \mathcal{A} to obtain the final anomaly scores.

Experiments

Dataset and Evaluation Metrics

Dataset MVTEC 3D-AD (Bergmann et al. 2021b) is a dataset designed for multimodal industrial anomaly detection, containing 10 categories of real-world objects with a total of 4147 samples. Both the training and validation sets consist of only normal samples, while the test set includes various types of anomalies such as scratches, dents, and contamination. Eyecandies (Bonfiglioli et al. 2022) is a synthetic dataset for multimodal anomaly detection, containing 10 categories of candies, with 10000 training images, 1000 validation images, and 4000 test images. Each sample includes RGB images and depth maps, rendered under six different lighting conditions.

Evaluation Metrics We use the area under the receiver operator curve on image-level (I-AUROC) and pixel-level (P-AUROC) to evaluate the detection and segmentation performance of the proposed method, respectively. Additionally, to reduce the impact of the anomaly area, we use per-region overlap (PRO) (Bergmann et al. 2021a) and compute the area under the PRO curve (AUPRO) to evaluate the segmentation performance. AUPRO represents the average overlap of the prediction with each connected component of ground truth.

MVTEC 3D-AD														
Setting	Type	Method	Bagel	Cable Gland	Carrot	Cookie	Dowel	Foam	Peach	Potato	Rope	Tire	Mean	
1-shot	TB	AST (Rudolph et al. 2023)	70.7/75.9	42.2/73.3	54.8/88.0	49.0/60.2	53.8/79.4	46.4/44.0	51.9/84.0	49.7/85.9	72.0/75.8	41.9/74.0	53.2/74.0	
		EasyNet (Chen et al. 2023)	61.4/79.6	21.2/75.1	75.9/69.8	56.5/85.8	62.8/49.4	65.7/69.0	63.0/88.1	94.6/71.8	47.7/75.4	60.1/75.5	60.1/75.5	
		ShapeGuided (Chu et al. 2023)	65.9/95.9	44.4/71.6	62.3/93.5	93.8/94.1	59.3/86.4	57.6/63.8	67.6/94.0	42.8/96.3	93.3/88.8	62.9/90.1	65.0/87.4	65.0/87.4
		M3DM (Wang et al. 2023)	87.8/95.3	64.1/81.5	78.0/97.2	92.7/90.3	64.2/81.6	65.3/82.0	75.5/94.0	79.8/94.8	85.8/95.2	45.4/89.8	73.9/90.2	73.9/90.2
		CFM (Costanzino et al. 2024)	43.5/92.6	56.7/79.4	76.4/96.9	95.4/92.8	53.2/85.2	71.4/87.8	63.0/94.9	64.2/96.1	91.8/96.6	59.2/91.4	67.5/91.4	67.5/91.4
		Patchcore+FPFH (Horwitz and Hoshen 2023)	62.2/92.8	53.4/76.8	54.0/96.7	55.9/92.8	54.7/84.6	63.3/71.9	49.6/95.9	60.7/96.1	88.8/90.8	56.6/84.1	59.9/88.3	59.9/88.3
TF	CIF (Ours)	78.9/85.2	68.7/79.5	72.2/95.6	81.2/86.2	62.9/82.0	72.8/70.3	83.9/93.3	60.4/94.2	83.5/88.1	55.8/86.4	72.0/86.1	72.0/86.1	
2-shot	TB	AST (Rudolph et al. 2023)	71.9/75.9	43.4/74.0	54.5/87.8	50.8/62.2	53.7/79.5	46.1/43.6	51.6/83.7	50.4/85.6	75.8/76.5	40.2/72.8	53.8/74.2	
		EasyNet (Chen et al. 2023)	47.6/77.8	76.1/62.9	52.6/92.6	60.2/59.1	31.7/58.8	52.3/57.2	71.9/21.1	76.1/15.2	61.2/43.1	51.2/47.7	58.1/49.3	
		ShapeGuided (Chu et al. 2023)	47.9/96.6	46.0/73.2	60.5/96.5	95.9/95.5	55.3/86.5	50.2/71.4	69.7/95.3	41.3/96.3	93.6/89.3	79.3/91.3	64.0/89.2	64.0/89.2
		M3DM (Wang et al. 2023)	91.8/95.5	57.0/82.9	79.8/97.2	94.5/88.0	61.4/87.0	79.5/79.6	79.2/95.1	75.1/94.2	92.8/95.5	54.1/91.0	74.8/90.6	74.8/90.6
		CFM (Costanzino et al. 2024)	82.6/95.3	65.0/80.3	80.3/97.9	97.4/93.1	53.3/85.6	68.0/89.1	70.4/96.1	74.7/96.3	92.1/96.9	64.0/94.1	74.8/92.5	74.8/92.5
		Patchcore+FPFH (Horwitz and Hoshen 2023)	63.2/94.6	47.7/76.5	55.4/96.7	64.5/93.4	58.3/85.0	61.1/67.4	55.0/96.2	56.6/96.6	88.2/91.0	64.4/88.3	61.4/88.6	61.4/88.6
TF	CIF (Ours)	85.3/87.1	62.9/79.5	74.0/95.8	72.2/87.1	64.6/82.1	79.5/75.3	77.6/93.7	66.9/94.9	86.1/88.7	62.7/87.7	73.2/87.2	73.2/87.2	
4-shot	TB	AST (Rudolph et al. 2023)	70.1/74.7	42.9/73.7	55.7/87.7	51.8/61.3	54.0/79.6	46.6/41.3	52.0/84.3	49.8/85.9	72.6/75.9	39.8/74.2	53.5/73.9	
		EasyNet (Chen et al. 2023)	67.5/70.7	36.3/13.8	54.7/86.9	69.1/72.0	72.4/39.3	50.2/53.6	74.3/86.2	63.1/90.2	44.9/15.5	53.7/66.3	58.6/56.5	
		ShapeGuided (Chu et al. 2023)	65.4/97.3	48.8/78.9	73.1/97.3	96.5/95.4	69.8/90.4	59.1/83.6	68.1/95.7	49.9/97.5	92.2/89.6	71.5/92.1	69.8/91.8	69.8/91.8
		M3DM (Wang et al. 2023)	98.6/96.1	68.5/87.2	83.7/97.3	93.8/90.5	59.6/86.5	87.8/86.6	85.3/96.4	70.5/94.6	89.3/95.5	56.1/93.1	79.3/92.4	79.3/92.4
		CFM (Costanzino et al. 2024)	92.8/96.5	64.3/84.4	87.8/98.0	98.4/93.5	64.1/89.7	77.8/92.3	85.2/97.2	75.3/96.8	93.3/97.1	61.5/94.8	80.1/94.0	80.1/94.0
		Patchcore+FPFH (Horwitz and Hoshen 2023)	52.3/95.7	54.0/79.7	59.2/97.2	62.9/95.1	59.4/87.5	58.2/75.3	61.6/96.5	70.5/97.3	91.8/91.2	73.5/88.0	64.3/90.4	64.3/90.4
TF	CIF (Ours)	91.8/92.9	70.0/83.3	77.5/96.9	83.4/86.4	70.2/84.7	79.4/84.1	85.6/95.2	75.0/95.5	89.8/89.3	53.4/88.2	77.6/89.6	77.6/89.6	
Eyecandies														
Setting	Type	Method	Can. C.	Cho. C.	Cho. P.	Conf.	Gum. B.	Haz. T.	Lic. S.	Lollipop.	Marsh.	Pep. C.	Mean	
1-shot	TB	M3DM (Wang et al. 2023)	36.2/86.8	66.9/82.5	73.1/70.6	84.8/94.2	71.3/74.9	50.2/54.8	57.4/71.0	59.9/84.2	60.3/89.8	81.0/90.0	64.1/79.9	
		CFM (Costanzino et al. 2024)	40.8/91.4	48.8/81.5	73.1/69.7	88.8/91.9	55.1/77.4	72.0/66.1	45.4/63.6	59.4/82.9	45.4/63.6	81.0/90.0	61.7/79.1	
		TF	CIF (Ours)	38.1/88.4	81.8/71.4	70.2/56.3	84.0/86.9	59.6/63.9	59.7/50.6	58.1/55.7	68.6/85.8	89.1/65.7	85.4/67.1	69.5/69.2
2-shot	TB	M3DM (Wang et al. 2023)	38.9/81.1	67.5/84.4	81.1/68.6	92.3/97.3	61.7/76.1	53.6/57.6	59.4/72.8	64.0/85.6	76.5/89.3	86.9/90.8	68.2/80.4	
		CFM (Costanzino et al. 2024)	38.2/92.4	63.2/84.3	76.5/73.2	87.4/92.5	61.1/79.5	72.5/70.3	51.8/64.9	64.9/83.2	79.4/86.3	67.8/81.5	68.3/80.8	
		TF	CIF (Ours)	41.8/87.0	79.0/76.9	78.6/57.9	94.1/85.5	71.3/61.0	70.9/52.4	55.7/59.9	66.4/86.8	84.5/71.8	93.6/74.3	73.6/71.3
4-shot	TB	M3DM (Wang et al. 2023)	42.4/82.4	74.9/84.9	78.7/72.1	91.4/96.8	70.2/80.2	53.4/61.7	80.2/81.6	67.8/88.6	86.6/94.6	90.6/92.1	73.6/83.5	
		CFM (Costanzino et al. 2024)	43.4/93.0	71.4/85.1	78.1/75.0	85.8/94.1	72.3/81.8	73.0/75.3	50.1/75.1	66.9/84.3	95.4/93.6	79.2/85.2	71.6/84.3	
		TF	CIF (Ours)	48.2/87.4	81.8/78.1	74.6/58.4	97.9/88.9	61.9/68.0	69.8/58.1	73.3/63.9	61.8/88.1	94.2/86.0	87.4/78.8	75.1/75.6

Table 1: I-AUROC/AUPRO scores for anomaly detection and localization of all categories of MVTEC 3D-AD and Eyecandies in few-shot settings (1-shot, 2-shot and 4-shot). **Bold** and underline denote the best-performing and second-best-performing methods, respectively. TB and TF represent Training-Based and Training-Free, respectively.

Modules			Performance	
SGMS	HGMS	Bi-TF-MP	I-AUROC	AUPRO
			68.6	76.1
✓			71.2	85.0
✓	✓		71.7	86.0
✓	✓	✓	72.0	86.1

Table 2: Module ablation study of CIF. The best is in **bold**.

Implementation Details

For hypergraph construction, we set the number of clusters $|\mathcal{E}|$ to 4 for MVTEC 3D-AD (8 for Eyecandies). For training-free hypergraph message passing, we set the number of message passing steps L to 1 and the retention coefficient α to 0.9. For memory bank construction, we set the sampling rate to 0.1. All experiments were conducted on a single NVIDIA A100 40GB GPU using PyTorch-1.13.1.

Comparison with SOTAs

We compare CIF with several multimodal anomaly detection methods on MVTEC 3D-AD and Eyecandies. Table 1 shows anomaly detection performance (I-AUROC scores) and localization performance (AUPRO scores) in few-shot settings. Existing methods are categorized into training-based and training-free approaches. (1) On MVTEC 3D-AD, compared with training-free methods, CIF achieves the highest I-AUROC in the 1-shot, 2-shot, and 4-shot settings, outperforming Patchcore+FPFH (Horwitz and Hoshen 2023) by 20.2%, 19.2%, and 20.6%, respectively. Compared with training-based methods, although CIF does not achieve the highest

I-AUROC, the performance gap with the best method is relatively small: only 2.6%, 4.5%, and 3.2% lower in the 1-shot, 2-shot, and 4-shot settings, respectively. (2) On Eyecandies, we compare CIF with the two best-performing training-based methods on MVTEC 3D-AD: M3DM (Wang et al. 2023) and CFM (Costanzino et al. 2024). CIF achieves the highest I-AUROC in the 1-shot, 2-shot, and 4-shot settings, outperforming the second-best method by 8.4%, 7.8%, and 2.0%, respectively. (3) On both datasets, using other methods as a baseline, CIF performs better in the 1-shot setting than in the 2-shot and 4-shot settings. This shows that CIF performs better when fewer training samples are available, indicating that the structural information extracted by hypergraphs is effective. (4) On both datasets, CIF does not achieve the highest AUPRO score in any setting, indicating that its ability in anomaly localization is relatively limited. Figure 5 presents anomaly segmentation visualizations of our method, M3DM, and CFM across all categories on MVTEC 3D-AD and Eyecandies (in multimodal, 1-shot setting). "PC" denotes point cloud, and "GT" denotes ground truth. Visually, our method effectively reduces the false positive rate.

Ablation Study

We conduct an ablation study in a multimodal, 1-shot setting on MVTEC 3D-AD. We analyze CIF in Table 2 with the following settings: 1) Without SGMS, HGMS, or Bi-TF-MP; 2) With SGMS only; 3) With SGMS and HGMS, but without Bi-TF-MP; 4) With SGMS, HGMS, and Bi-TF-MP (Ours). Comparing row 1 and row 2 in Table 2, we can find that adding SGMS greatly improves the results on both metrics (I-AUROC by 3.8% and AUPRO by 11.7%), indicating that

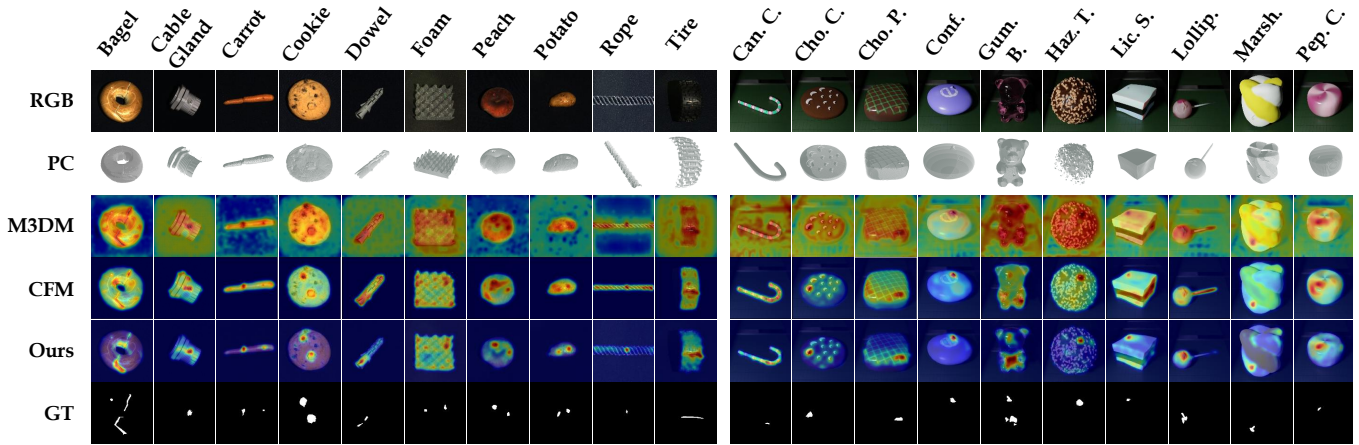


Figure 5: Visualization of anomaly scores for each category of MVTec 3D-AD and Eyecandies (in multimodal, 1-shot setting).




Figure 6: Visualization of the impact of parameters in SAHC and Bi-TF-MP on anomaly detection and localization performance.

the node features retained through SGMS are more effective. Comparing row 2 and row 3, we observe that HGMS further enhances detection and segmentation performance (I-AUROC by 0.7% and AUPRO by 1.2%). Our full model is shown in row 4 in Table 2, and compared with row 3, we have 0.4% I-AUROC and 0.1% AUPRO improvement, showing that Bi-TF-MP reduces the distribution gap between test sample nodes and memory bank nodes.

Analysis of SAHC Parameter

The number of hyperedges determines the ability of the hypergraph to capture the structural information. Too few hyperedges cause different structures to be treated as similar, while too many hyperedges force the same structure to be split apart. We explore the impact of different numbers of hyperedges in SAHC on anomaly detection and localization performance. As shown in Figure 6 (a), we conducted experiments in a multimodal, 1-shot setting on MVTec 3D-AD, and calculated the mean I-AUROC, P-AUROC, and AUPRO scores across all classes. We found that when the number of hyperedges is 4, I-AUROC reaches its highest; when the number of hyperedges is 8, P-AUROC and AUPRO reach their highest. This shows that fewer hyperedges help improve anomaly detection performance, while more hyperedges benefit anomaly localization performance. We give a class-specific analysis in extended version.

Analysis of Bi-TF-MP Parameter

During message passing, the retention coefficient α and the number of Bi-TF-MP layers L respectively determine the

level of retention of the own information of nodes and the range of information exchange with neighboring nodes. We conducted experiments in a multimodal, 1-shot setting on MVTec 3D-AD and calculated the mean I-AUROC and AUPRO scores across all classes. As shown in Figure 6 (b) and (c), we found that when $\alpha = 0.9$, both I-AUROC and AUPRO reached their highest. This indicates that slight message passing can alleviate the distribution gap between test sample nodes and memory bank nodes. Furthermore, as shown in Figure 6 (d), when $\alpha = 0.9$, anomaly detection and localization performance were best at $L = 1$, and the performance became stable as L increased. We give a detailed analysis of message passing parameters in extended version.

Conclusion

In this paper, we propose a hypergraph-based few-shot multimodal industrial anomaly detection method. We use hypergraphs to extract intra-class structural commonality of samples to guide the construction and search of the memory bank. We introduce semantic-aware hypergraph construction and structure-guided memory sampling to build and compress a memory bank that contains structural information. We then use bidirectional training-free hypergraph message passing to reduce the distribution gap between test samples and the memory bank. Finally, we design hyperedge-guided memory search to reduce the false positive rate in anomaly detection. Our method outperforms the state-of-the-art results on MVTec 3D-AD and Eyecandies in few-shot settings. We hope our work will be helpful for future research.

Acknowledgments

This work was supported by National Natural Science Foundation of China (No.62576109, 62072112, 62406075), National Key Research and Development Program of China (2023YFC3604802).

References

- Angiulli, F.; and Pizzuti, C. 2002. Fast outlier detection in high dimensional spaces. In *European conference on principles of data mining and knowledge discovery*, 15–27. Springer.
- Bergmann, P.; Batzner, K.; Fauser, M.; Sattlegger, D.; and Steger, C. 2021a. The MVTec anomaly detection dataset: a comprehensive real-world dataset for unsupervised anomaly detection. *International Journal of Computer Vision*, 129(4): 1038–1059.
- Bergmann, P.; Fauser, M.; Sattlegger, D.; and Steger, C. 2020. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4183–4192.
- Bergmann, P.; Jin, X.; Sattlegger, D.; and Steger, C. 2021b. The mvtec 3d-ad dataset for unsupervised 3d anomaly detection and localization. *arXiv preprint arXiv:2112.09045*.
- Bezdek, J. C.; Ehrlich, R.; and Full, W. 1984. FCM: The fuzzy c-means clustering algorithm. *Computers & geosciences*, 10(2-3): 191–203.
- Bonfiglioli, L.; Toschi, M.; Silvestri, D.; Fioraio, N.; and De Gregorio, D. 2022. The eyecandies dataset for unsupervised multimodal anomaly detection and localization. In *Proceedings of the Asian Conference on Computer Vision*, 3586–3602.
- Breunig, M. M.; Kriegel, H.-P.; Ng, R. T.; and Sander, J. 2000. LOF: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, 93–104.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9650–9660.
- Chen, R.; Xie, G.; Liu, J.; Wang, J.; Luo, Z.; Wang, J.; and Zheng, F. 2023. Easynet: An easy network for 3d industrial anomaly detection. In *Proceedings of the 31st ACM International Conference on Multimedia*, 7038–7046.
- Chu, Y.-M.; Liu, C.; Hsieh, T.-I.; Chen, H.-T.; and Liu, T.-L. 2023. Shape-guided dual-memory learning for 3D anomaly detection. In *International Conference on Machine Learning*, 6185–6194. PMLR.
- Costanzino, A.; Ramirez, P. Z.; Lisanti, G.; and Di Stefano, L. 2024. Multimodal industrial anomaly detection by cross-modal feature mapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17234–17243.
- Crosby, T. 1994. How to detect and handle outliers. Defard, T.; Setkov, A.; Loesch, A.; and Audigier, R. 2021. Padim: a patch distribution modeling framework for anomaly detection and localization. In *International Conference on Pattern Recognition*, 475–489. Springer.
- Feng, Y.; You, H.; Zhang, Z.; Ji, R.; and Gao, Y. 2019. Hypergraph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 3558–3565.
- Gao, Y.; Feng, Y.; Ji, S.; and Ji, R. 2022. HGNN+: General hypergraph neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 3181–3199.
- Gu, Z.; Zhu, B.; Zhu, G.; Chen, Y.; Tang, M.; and Wang, J. 2024. Anomalygpt: Detecting industrial anomalies using large vision-language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 1932–1940.
- Gudovskiy, D.; Ishizaka, S.; and Kozuka, K. 2022. Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 98–107.
- Han, Y.; Wang, P.; Kundu, S.; Ding, Y.; and Wang, Z. 2023. Vision hgnn: An image is more than a graph of nodes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19878–19888.
- Horwitz, E.; and Hoshen, Y. 2023. Back to the feature: classical 3d features are (almost) all you need for 3d anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2968–2977.
- Huang, C.; Guan, H.; Jiang, A.; Zhang, Y.; Spratling, M.; and Wang, Y.-F. 2022. Registration based few-shot anomaly detection. In *European Conference on Computer Vision*, 303–319. Springer.
- Lee, M.; and Choi, J. 2024. Text-Guided Variational Image Generation for Industrial Anomaly Detection and Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26519–26528.
- Lee, S.; Lee, S.; and Song, B. C. 2022. Cfa: Coupled-hypersphere-based feature adaptation for target-oriented anomaly localization. *IEEE Access*, 10: 78446–78454.
- Li, C.; Li, T.; Hu, X.; Luo, D.; and Jin, T. 2025. DVHGNN: Multi-Scale Dilated Vision HGNN for Efficient Vision Recognition. *arXiv preprint arXiv:2503.14867*.
- Li, C.-L.; Sohn, K.; Yoon, J.; and Pfister, T. 2021. Cut-paste: Self-supervised learning for anomaly detection and localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9664–9674.
- Li, X.; Zhang, Z.; Tan, X.; Chen, C.; Qu, Y.; Xie, Y.; and Ma, L. 2024. Promptad: Learning prompts with only normal samples for few-shot anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16838–16848.
- Liu, F. T.; Ting, K. M.; and Zhou, Z.-H. 2008. Isolation forest. In *2008 eighth IEEE international conference on data mining*, 413–422. IEEE.
- Liu, Z.; Zhou, Y.; Xu, Y.; and Wang, Z. 2023. Simplenet: A simple network for image anomaly detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20402–20411.

- Pang, Y.; Wang, W.; Tay, F. E.; Liu, W.; Tian, Y.; and Yuan, L. 2022. Masked autoencoders for point cloud self-supervised learning. In *European conference on computer vision*, 604–621. Springer.
- Roth, K.; Pemula, L.; Zepeda, J.; Schölkopf, B.; Brox, T.; and Gehler, P. 2022. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14318–14328.
- Rudolph, M.; Wehrbein, T.; Rosenhahn, B.; and Wandt, B. 2023. Asymmetric student-teacher networks for industrial anomaly detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2592–2602.
- 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*, 13(7): 1443–1471.
- Shyu, M.-L.; Chen, S.-C.; Sarinnapakorn, K.; and Chang, L. 2003. A novel anomaly detection scheme based on principal component classifier. In *Proceedings of the IEEE foundations and new directions of data mining workshop*, 172–179. IEEE Press.
- Tang, B.; Liu, Z.; Jiang, K.; Chen, S.; and Dong, X. 2024. Training-Free Message Passing for Learning on Hypergraphs. *arXiv preprint arXiv:2402.05569*.
- Tang, J.; Chen, Z.; Fu, A. W.-C.; and Cheung, D. W. 2002. Enhancing effectiveness of outlier detections for low density patterns. In *Advances in Knowledge Discovery and Data Mining: 6th Pacific-Asia Conference, PAKDD 2002 Taipei, Taiwan, May 6–8, 2002 Proceedings 6*, 535–548. Springer.
- Wang, Y.; Peng, J.; Zhang, J.; Yi, R.; Wang, Y.; and Wang, C. 2023. Multimodal Industrial Anomaly Detection via Hybrid Fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8032–8041.
- Wu, J.-C.; Chen, D.-J.; Fuh, C.-S.; and Liu, T.-L. 2021. Learning unsupervised metaformer for anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4369–4378.
- Xie, G.; Wang, J.; Liu, J.; Zheng, F.; and Jin, Y. 2023. Pushing the limits of fewshot anomaly detection in industry vision: Graphcore. *arXiv preprint arXiv:2301.12082*.
- Yu, J.; Zheng, Y.; Wang, X.; Li, W.; Wu, Y.; Zhao, R.; and Wu, L. 2021. Fastflow: Unsupervised anomaly detection and localization via 2d normalizing flows. *arXiv preprint arXiv:2111.07677*.
- Zavrtanik, V.; Kristan, M.; and Skočaj, D. 2021a. Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8330–8339.
- Zavrtanik, V.; Kristan, M.; and Skočaj, D. 2021b. Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition*, 112: 107706.
- Zhang, X.; Li, S.; Li, X.; Huang, P.; Shan, J.; and Chen, T. 2023. DeSTSeg: Segmentation Guided Denoising Student-Teacher for Anomaly Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3914–3923.
- Zhang, X.; Xu, M.; and Zhou, X. 2024. RealNet: A Feature Selection Network with Realistic Synthetic Anomaly for Anomaly Detection. *arXiv preprint arXiv:2403.05897*.
- Zhou, D.; Huang, J.; and Schölkopf, B. 2006. Learning with hypergraphs: Clustering, classification, and embedding. *Advances in neural information processing systems*, 19.
- Zhu, J.; and Pang, G. 2024. Toward generalist anomaly detection via in-context residual learning with few-shot sample prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17826–17836.