

ArcAD: Anomaly-Rectified Calibration for Cold-Start Supervised Anomaly Detection

Ningning Han^{1,6*}, Lei Fan^{2,6*}, Jia Guo³, Yunkang Cao⁴, Xiu Su⁵, Feng Cao⁶,
Donglin Di^{1,6}, and Tonghua Su^{1,7}

¹ Faculty of Computing, Harbin Institute of Technology, Harbin, China
25B903152@stu.hit.edu.cn, thsu@hit.edu.cn

² University of New South Wales, Sydney, Australia
lei.fan1@unsw.edu.au

³ School of Biomedical Engineering, Tsinghua University, Beijing, China

⁴ School of Artificial Intelligence and Robotics, Hunan University, Changsha, China

⁵ Big data institute, Central South University, Sydney, Australia

⁶ DZ-Matrix, Beijing, China

⁷ Chongqing Research Institute of HIT, Chongqing, China

Abstract. The deployment of Industrial Anomaly Detection (IAD) in real-world manufacturing frequently encounters a challenging cold-start bottleneck, in which limited normal samples fail to represent the full normal distribution and only a few anomalies are available. Under such a regime, existing methods struggle to form compact normal boundaries and fail to effectively exploit supervised signals from rare defects. To address this challenge, we propose Anomaly-Rectified Cold-start AD (ArcAD), a plug-and-play calibration framework for reconstruction-based IAD baselines. ArcAD follows a push-pull learning paradigm to construct a compact and discriminative normal boundary under data scarcity. On the one hand, ArcAD projects limited normal samples onto a hypersphere and pulls them into multiple compact clusters to maximize coverage of the normal manifold. On the other hand, it synthesizes pseudo-anomalies on the hypersphere and leverages real anomalies to push the boundary inward and sharpen anomaly discrimination. Extensive experiments on MVTec-AD, VisA, Real-IAD, and MANTA demonstrate that ArcAD significantly outperforms state-of-the-art supervised and unsupervised methods in both single-class and multi-class settings under cold-start conditions. Code is available at: <https://github.com/LGC-AD/ArcAD>.

Keywords: Anomaly Detection · Hypersphere · Cold Start

1 Introduction

Industrial Anomaly Detection (IAD) seeks to achieve zero-defect manufacturing by identifying rare defects, such as scratches, cracks, or structural deformations [7, 9, 18, 25]. The deployment of IAD systems in real-world applications

* Equal contribution.

Corresponding authors: Tonghua Su and Lei Fan.




Fig. 1: Left: Decision boundaries in cold-start scenarios. Unsupervised methods (a) yield loose manifolds and supervised methods (b) overfit to known anomalies, whereas ArcAD (c) forms compact and unified boundaries. **Right:** Paradigm comparison on data efficiency, anomaly utilization, and generalization.

frequently encounters a severe *cold-start* bottleneck. Unlike standard unsupervised settings [12, 41] that assume abundant normal data, cold-start scenarios during the initial ramp-up phase of a new production line present a distinct data distribution. During this stage, the available normal samples fail to cover the full range of normal patterns [39, 40]. Meanwhile, only a small number of anomalies are collected during early system deployment.

Recent supervised methods [6, 49, 54] show promising performance when trained with annotated normal and abnormal samples. However, in cold-start scenarios, these methods are prone to overfitting when trained on only a few anomalies, leading to poor generalization to unseen defect patterns (see Fig. 1). Conversely, unsupervised approaches, particularly reconstruction-based methods [5, 14, 15], have established dominance in both single-class and multi-class settings [50], primarily due to their strong ability to generalize to diverse and unseen anomaly patterns. These methods typically learn the distribution of normal data and identify instances that deviate from the learned normal manifold as anomalies. While abundant data allows models to implicitly learn a continuous and robust normal manifold, data scarcity in cold-start scenarios leads to a fragmented and loosely bounded latent space. Moreover, by strictly relying on normal data, these unsupervised paradigms inherently underutilize the valuable guidance provided by the few available anomaly samples. Therefore, a pivotal question emerges: *How can we leverage limited normal samples and rare anomalies to construct a compact normal boundary?*

To address this challenge, we propose Anomaly-Rectified Cold-start AD (ArcAD), a plug-and-play framework designed to enhance reconstruction-based models. ArcAD constructs a compact and discriminative normal boundary from two complementary perspectives: explicitly organizing normal samples into compact clusters and calibrating the boundary using anomaly signals. First, we project patch features onto a hypersphere, separating directional information from magnitude variations to establish a bounded geometric representation [4]. These hyperspherical features are then modeled using the von Mises-Fisher (vMF) distribution [29]. To ensure that limited normal data adequately covers the latent manifold, we introduce Sinkhorn-based Prototype Modeling (SPM). By formulating feature clustering as an optimal transport problem, SPM mitigates biased

feature aggregation and partitions normal embeddings into compact and uniform distributed clusters.

Beyond modeling the normal manifold, ArcAD further leverages anomaly signals to explicitly rectify the normal boundary. To alleviate the scarcity of anomalies, we design a prototype-restricted anomaly synthesis strategy, which generates synthetic anomalies directly on the hypersphere by filtering candidate samples against normal prototypes. These synthesized samples, together with the few genuine defects, drive Defect-Guided Calibration (DGC). This module employs a contrastive objective that pulls anomalies together while pushing them away from their nearest normal prototypes. This process explicitly reinforces the compactness and discriminative power of the learned normal manifold. We summarize our contributions as follows:

- We propose ArcAD, a generic plug-and-play calibration framework for reconstruction based models under cold-start scenarios.
- We introduce Sinkhorn-based Prototype Modeling (SPM) to organize limited normal samples into compact and uniform clusters on the hypersphere.
- We propose Defect-Guided Calibration (DGC), which introduces a prototype-restricted synthesis strategy to generate pseudo-anomalies, thereby leveraging both real and synthesized defects to explicitly rectify the normal boundary in the latent space.
- Extensive experiments on four datasets [1, 8, 43, 55] demonstrate that ArcAD consistently enhances state-of-the-art baselines such as ReContrast [13], RD4AD [5], and Dinomaly [14]. Specifically, under the challenging multi-class setting of the Real-IAD dataset [43], ArcAD achieves Image-level AUROC gains of +2.2%, +8.9%, and +3.7% for these models, respectively.

2 Related Work

Industrial Anomaly Detection. Depending on the availability of training data, existing IAD approaches can be broadly categorized into unsupervised and supervised AD methods. Unsupervised AD typically follows a one-class classification paradigm, including reconstruction-based [13–15, 52] and embedding-based methods [30, 39, 48]. These approaches learn the distribution of normal samples and detect anomalies as deviations during inference. However, under cold-start scenarios, limited normal samples often lead to loosely defined latent boundaries, significantly degrading detection performance. Supervised AD [22, 35, 49] instead leverages the prior knowledge derived from observed anomalies to mitigate false positives. For instance, DRA [6] learns representations of observed anomalies, pseudo-anomalies, and residual anomalies to enhance anomaly discrimination. Similarly, DPDL [44] confines normal samples within a compact feature space and employs the Schrödinger bridge to guide anomalous samples towards normal distribution. However, these methods tend to overfit to the limited observed anomalies, which restricts their generalization to unseen defect patterns.

Hypersphere Learning. Constraining representations onto a unit hypersphere has emerged as an effective strategy for enhancing feature uniformity and embedding alignment [3, 10, 45]. To statistically characterize these hyperspherical representations, the von Mises-Fisher (vMF) distribution [29] provides a principled probabilistic foundation and has been widely adopted in clustering [11, 37], semantic segmentation [17], and face recognition [20, 47]. In addition, hyperspherical representations have also been explored in generative modeling to constrain latent spaces and improve generation fidelity [19, 27]. More recently, hyperspherical modeling has demonstrated advantages in out-of-distribution detection [23, 28, 31] by naturally providing a bounded latent space.

Motivated by these geometric properties, we introduce ArcAD, which models features on a unit hypersphere to organize limited normal samples into a compact manifold while leveraging rare anomalies to rectify the normal boundary.

3 Methodology

3.1 Overview

Cold-Start Setting. We consider industrial anomaly detection under a realistic *cold-start* setting, where only limited normal samples and a few genuine defect samples are available for training. Let $\mathcal{X} = \mathcal{X}_N \cup \mathcal{X}_A$ denote the training set, where $\mathcal{X}_N = \{x_i\}$ represents normal images and $\mathcal{X}_A = \{x_j\}$ represents a small set of genuine defect samples, with $|\mathcal{X}_N| \gg |\mathcal{X}_A|$ (*e.g.*, 9:1). The overall training data is limited to reflect early-stage deployment, while normal samples strictly outnumber anomalies. For evaluation, we define the test set as $\mathcal{X}_{test} = \mathcal{X}_{test,N} \cup \mathcal{X}_{test,A}$. The objective is to learn a compact and discriminative boundary for normal data under the cold-start regime.

Framework Overview. ArcAD is built upon reconstruction-based IAD frameworks [5, 13, 14] with an encoder $\mathcal{E}(\cdot)$, a bottleneck $\mathcal{B}(\cdot)$, and a decoder $\mathcal{D}(\cdot)$. Given an input image $x_i \in \mathcal{X}_N$, the encoder and bottleneck map it to a latent representation $z_i = \mathcal{B}(\mathcal{E}(x_i))$. The decoder then reconstructs representations from the latent space. The optimization objective can be formulated as:

$$\mathcal{L}_{recon} = d(\mathcal{E}(x_i), \mathcal{D}(\mathcal{B}(\mathcal{E}(x_i))))), \quad (1)$$

where $d(\cdot, \cdot)$ denotes a distance function (*e.g.*, mean squared error or cosine distance) measuring the discrepancy between the original and reconstructed features. While reconstruction-based loss implicitly captures the normal manifold, under cold-start conditions, the limited normal data can lead to a biased and loosely bounded latent structure, yielding an imprecise normal boundary.

To mitigate this issue, we propose Anomaly-Rectified Cold-start AD (ArcAD), a plug-and-play calibration framework that operates on the latent space without changing the reconstruction backbone (see Fig. 2). ArcAD consists of two core components: Sinkhorn-based Prototype Modeling (SPM), which organizes limited normal samples to form a compact and well-covered normal structure, and Defect-Guided Calibration (DGC), which refines the normal boundary by leveraging both synthesized pseudo-anomalies and real anomalies.

prototype, we learn multiple prototypes to capture diverse normal patterns. Under cold-start conditions, limited normal samples may lead to a biased feature space [36, 39], where prototypes collapse into a few dense regions. To encourage uniform prototype learning, we formulate the feature-to-prototype assignment as a balanced optimal transport problem.

Specifically, we initialize K prototypes $\mathcal{U} = \{\boldsymbol{\mu}_k\}_{k=1}^K$ on the hypersphere using K -means with cosine distance. Given a batch with N_{batch} normalized patch features, we seek a soft assignment matrix $Q \in \mathbb{R}^{N_{\text{batch}} \times K}$, where Q_{ik} denotes the assignment weight between feature $\mathbf{z}_i$ and prototype $\boldsymbol{\mu}_k$. The assignment is optimized to favor high feature-prototype similarity while satisfying an equipartition constraint. Concretely, each feature is fully assigned across prototypes (row sums equal 1), and each prototype receives equal total mass (column sums equal N_{batch}/K).

To compute Q , we solve this optimal transport problem using the iterative Sinkhorn-Knopp algorithm [2]. We compute the pairwise cosine similarity matrix $S \in \mathbb{R}^{N_{\text{batch}} \times K}$ with $S_{ik} = \mathbf{z}_i^\top \boldsymbol{\mu}_k$, and initialize the exponential similarity matrix as $M = \exp(S/\epsilon)$, where ϵ is a temperature parameter. The algorithm then derives an entropically regularized solution formulated as:

$$Q = \text{diag}(\mathbf{u})M\text{diag}(\mathbf{v}), \quad (3)$$

where $\mathbf{u} \in \mathbb{R}^{N_{\text{batch}}}$ and $\mathbf{v} \in \mathbb{R}^K$ are non-negative scaling vectors. The assignment and prototype update follow an Expectation-Maximization (EM)-style procedure [32]. In the E-step, we infer Q with fixed network parameters and prototypes. In the M-step, we update prototypes $\boldsymbol{\mu}_k$ using the exponential moving average (EMA) [16] to ensure stable evolution on the hypersphere.

With the optimal transport plan Q acting as the soft assignment weight, we formulate $\mathcal{L}_{spm}$ as a cross-entropy objective to align the network predictions:

$$\mathcal{L}_{spm} = -\frac{1}{N_{\text{batch}}} \sum_{i=1}^{N_{\text{batch}}} \sum_{k=1}^K Q_{ik} \log \frac{\exp(\mathbf{z}_i^\top \boldsymbol{\mu}_k / \tau)}{\sum_{j=1}^K \exp(\mathbf{z}_i^\top \boldsymbol{\mu}_j / \tau)}, \quad (4)$$

where $\tau = 1/\kappa$ is a temperature hyperparameter. By optimizing $\mathcal{L}_{spm}$, we enforce the normal representations to form compact clusters within the latent space. It is crucial to note that the Sinkhorn mechanism strictly enforces uniformity to ensure no single prototype dominates the mass assignments, effectively preventing biased learning toward high-density regions. Concurrently, the comprehensive coverage of the normal manifold is achieved by the large capacity of the multi-cluster prototypes (K). Together, these properties establish a highly robust and discriminative normal boundary.

3.3 Defect-Guided Calibration (DGC)

Building upon the normal prototypes learned by SPM, we next use anomaly signals to explicitly calibrate the normal boundary. In particular, real anomalies

serve as negative samples that are pushed away from the normal prototypes, enforcing a margin between anomalous samples and the normal distribution.

In the cold-start regime, however, only a few real anomalies are available, which provide insufficient support to shape a stable anomalous region. To alleviate this scarcity, we introduce a prototype-restricted anomaly synthesis strategy to expand the negative set near the boundary. Unlike image-level augmentation (*e.g.*, Cutpaste [21] or DRAEM [51]) or feature-level augmentation in unbounded Euclidean space (*e.g.*, SimpleNet [26]), we perform perturbations directly on the hypersphere so that synthesized samples remain within a bounded geometry and avoid drifting into trivial outliers.

Prototype-Restricted Anomaly Synthesis. Given a normal feature $\mathbf{z} \in \mathbb{S}^{D-1}$ as an anchor, we generate a candidate set of perturbed features by injecting Gaussian noise $\boldsymbol{\xi} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ and re-projecting the resulting features onto the unit hypersphere:

$$\mathbf{a}'_{cand} = (\mathbf{z} + \boldsymbol{\xi}) / \|\mathbf{z} + \boldsymbol{\xi}\|_2. \quad (5)$$

However, blindly using all candidates may corrupt the normal manifold, as some candidates can remain close to normal regions. To enforce a discriminative margin, we filter the candidates using the learned normal prototypes $\mathcal{U} = \{\boldsymbol{\mu}_k\}_{k=1}^K$ as geometric constraints. Specifically, we select the candidate that minimizes its maximum similarity to the prototypes:

$$\mathbf{a}_{syn} = \arg \min_{\mathbf{a}'_{cand} \in \mathcal{A}'_{cand}} \max_{\boldsymbol{\mu}_k \in \mathcal{U}} (\mathbf{a}'_{cand})^\top \boldsymbol{\mu}_k. \quad (6)$$

In practice, for each normal anchor feature, we generate $N_{cand} = 5$ independent noise perturbations to form the candidate set $\mathcal{A}'_{cand}$, from which exactly one pseudo-anomaly is selected.

Contrastive Calibration. We calibrate the boundary with two complementary objectives. First, for real anomaly features, we penalize the maximum similarity to normal prototypes, pushing anomalies away from the normal distribution. Second, we align each real anomaly with synthesized pseudo-anomalies, encouraging a coherent anomalous region near the boundary.

Formally, given an abnormal image $\mathbf{x}_j \in \mathcal{X}_A$, we extract its patch-level features as a set $\mathcal{A} = \{\mathbf{a}_i\}_{i=1}^{N_p} \subset \mathbb{R}^D$. Guided by its down-sampled mask, we isolate and ℓ_2 -normalize anomalous patch embeddings onto the unit hypersphere, forming the anomaly set:

$$\hat{\mathcal{A}} = \{\mathbf{a}_i / \|\mathbf{a}_i\|_2 \mid \mathbf{a}_i \in \mathcal{A}, m_i = 1\} \quad (7)$$

where $m_i \in \{0, 1\}$ is the i -th element of the ground-truth mask $M \in \{0, 1\}^{N_p}$.

Let $\mathcal{A}_{syn}$ denote the set of all synthesized pseudo-anomalies generated from the normal anchors within a batch. Integrating the two objectives, the defect-guided calibration loss $\mathcal{L}_{dgc}$ is formulated as:

$$\mathcal{L}_{dgc} = \mathbb{E}_{\mathbf{a} \in \hat{\mathcal{A}}} \left(\max_{\boldsymbol{\mu}_k \in \mathcal{U}} (\mathbf{a}^\top \boldsymbol{\mu}_k) + \left(1 - \frac{1}{|\mathcal{A}_{syn}|} \sum_{\mathbf{a}_{syn} \in \mathcal{A}_{syn}} \mathbf{a}^\top \mathbf{a}_{syn} \right) \right), \quad (8)$$

where $|\mathcal{A}_{syn}|$ denotes the size of the synthesized anomaly set. Consequently, these anomalies provide a repulsive signal that sharpens the normal boundary and leads to a more compact fit of the normal distribution.

3.4 Optimization

Training. We introduce a lightweight discriminator $f_{disc}(\cdot)$ (e.g., a 2-layer MLP) to act as a global structural regularizer rather than merely a probability predictor. Combined with the SPM and DGC modules, it further enforces a clear decision margin in the latent space while maintaining the compactness of the normal hyperspherical manifold. To achieve this while mitigating the class imbalance inherent in the cold-start setting, the discriminator is optimized using the Binary Focal Loss [24]:

$$\mathcal{L}_{cls} = \mathbb{E}_{\mathbf{e} \in \mathcal{Z} \cup \mathcal{A}_{syn} \cup \hat{\mathcal{A}}} [\mathcal{L}_{FL}(f_{disc}(\mathbf{e}), y)], \quad (9)$$

where $\mathcal{Z} = \{\mathbf{z}_i\}$ denotes the set of normal features, and $y \in \{0, 1\}$ is the corresponding binary label (0 for normal features, and 1 for both real and synthetic anomalies).

During the training phase, the bottleneck network is optimized using the following comprehensive objective function:

$$\mathcal{L}_{total} = \mathcal{L}_{recon} + \lambda_1 \mathcal{L}_{spm} + \lambda_2 \mathcal{L}_{dgc} + \lambda_3 \mathcal{L}_{cls}, \quad (10)$$

where λ_1, λ_2 , and λ_3 are hyperparameters balancing the contribution of each module. The decoder is updated using only the reconstruction loss ($\mathcal{L}_{recon}$) derived from normal samples, while the parameters of the discriminator are optimized solely via the classification loss ($\mathcal{L}_{cls}$).

Inference. ArcAD does not alter the standard reconstruction-based anomaly scoring mechanism. Instead, it regularizes the latent space during training. At inference time, anomaly scores are computed directly from the reconstruction-based score map [5, 14].

4 Experiments

4.1 Experimental Setup

Datasets. We evaluate the proposed ArcAD on four standard industrial datasets. *MVTec-AD* [1] features 15 categories, comprising 3,629 normal training images and a test set of 1,725 images. *VisA* [55] consists of 12 distinct categories. The official data split provides 8,659 normal samples for training and 2,162 for testing. Furthermore, we utilize the large-scale *Real-IAD* [43] dataset, which covers 30 industrial categories. There are 36,465 normal images for training and a test set of 114,585 images. *MANTA* [8] contains 38 object categories. We utilize the tiny version of this dataset, which comprises 107,415 images in the training set and 44,585 images in the testing set.




Fig. 3: Statistics of standard and cold-start setting datasets (see details in *Supp*).

In our cold-start setting, 30% of the standard training set is utilized as normal training data, which is then combined with a subset of anomalies sampled from the original test set. Typically, anomalies account for 10% of this newly constructed training set. Due to the severe scarcity of anomalous samples, we set this ratio to 5% for VisA and MANTA. Unlike standard unsupervised datasets, the cold-start training set contains fewer normal samples and introduces a few anomalies. The test set incorporates the normal samples removed from the training set and contains reduced anomalies. Unsupervised methods are trained solely on the normal subset of the training data, whereas supervised methods and our approach utilize the full cold-start training set. All baselines are evaluated on this identical cold-start test set. Detailed statistics are presented in Fig. 3.

Evaluation Metrics. Consistent with established protocols [14, 52], we evaluate performance using three primary metrics. For image-level anomaly detection, we report the Area Under the Receiver Operating Characteristic curve (I-AUROC). For pixel-level localization, we assess performance via P-AUROC and P-F₁-max. Additional metrics are provided in the *Supp*. All reported scores are averaged across all classes.

Implementation Details. We integrate our proposed framework into three reconstruction based methods: Dinomaly [14], RD4AD [5], and ReContrast [13]. Data preprocessing follows the standard input and output configurations of each respective framework. All models are trained for 10,000 iterations across all evaluated datasets. We optimize the networks using the AdamW optimizer. The hyperparameters are configured with a learning rate of 2×10^{-3} , momentum parameters $\beta = (0.9, 0.999)$, and weight decay of 1×10^{-4} . The balancing hyperparameters λ_1, λ_2 , and λ_3 in our total loss function are set to 0.1. All baseline methods are reimplemented under the same cold-start setting.

4.2 Main Results

Multi-class Performance. We conducted experiments under the widely applied multi-class setting, where all classes of samples are mixed for training. We com-

Table 1: Performance under multi-class setting (%). Each cell reports I-AUROC / P-AUROC / P-F₁-max. Detailed per-class performances are presented in the *Supp.*

	Method	MVTec-AD [1]	VisA [55]	Real-IAD [43]	MANTA [8]
Unsup.	SimpleNet [26]	79.4/84.5/30.7	85.1/92.8/37.6	63.6/66.3/2.5	60.3/72.2/7.3
	DeSTSeg [53]	94.5/97.3/60.9	86.6/96.0/39.5	83.5/92.7/39.8	75.5/85.0/20.0
	MambaAD [15]	95.5/96.7/56.8	94.2/97.6/39.1	85.3/97.7/38.2	85.2/91.8/29.8
	ReContrast [13]	93.2/97.6/57.8	93.2/98.0/46.0	85.9/97.8/40.8	86.4/93.4/30.5
	RD4AD [5]	91.4/96.2/44.1	90.3/97.5/42.2	77.8/96.9/29.7	76.8/91.8/23.5
	UniNet [46]	92.7/96.4/49.7	90.9/98.0/44.4	76.5/96.8/32.9	68.4/89.9/15.4
	Dinomaly [14]	99.6/98.3/68.4	98.4/98.9/53.9	88.8/98.8/46.7	90.3/93.8/39.8
	Dinomaly* [14]	99.6/98.4/69.2	98.7/98.7/55.7	89.3/98.8/47.1	91.1/95.1/53.1
Sup.	SDNet [42]	70.1/49.4/6.9	79.0/54.9/1.4	74.7/63.3/0.9	78.5/66.9/4.2
	SDNet [42] (Dinov2)	59.7/60.2/3.9	77.4/42.1/0.6	55.4/71.8/0.8	60.7/70.6/5.1
	DevNet [34]	57.7/61.5/3.7	50.5/51.6/3.9	50.0/59.2/7.1	51.3/49.7/6.1
	DRA [6]	85.2/63.9/13.1	79.3/49.9/1.7	76.0/60.3/4.9	77.4/34.8/1.9
	ArcAD [†]	99.7/99.2/68.9	98.9/99.0/54.9	92.5/99.0/49.8	93.3/95.5/48.5

[†] Baseline built upon Dinomaly. * Ref: performance under standard setting from [12].

pared the proposed ArcAD based on Dinomaly, with both unsupervised [5, 13–15, 26, 46, 53] and supervised methods [6, 34, 42]. Experimental results are presented in Table 1. Overall, the proposed ArcAD consistently outperforms other state-of-the-art (SOTA) methods across all datasets and metrics.


Compared to unsupervised learning baselines, our approach yields significant improvements across the board. On both the MVTec-AD and VisA, ArcAD achieves the best performance with 99.7%/99.2%/68.9% and 98.9%/99.0%/54.9% (I-AUROC/P-AUROC/P-F₁-max), surpassing all SOTA methods. This performance gap widens significantly on more complex, large-scale datasets with over 30 categories. Specifically, on the Real-IAD dataset, ArcAD achieves an I-AUROC of 92.5% and a P-F₁-max of 49.8%, exceeding Dinomaly by notable margins of 3.7% and 3.1%. Similarly, on the MANTA dataset, ArcAD obtains an I-AUROC of 93.3% and a P-F₁-max of 48.5%, outperforming the best Dinomaly by 3.0% and 8.7%. Compared to unsupervised methods, our approach leverages a few anomaly samples as guidance to explicitly constrain and calibrate the normal manifold by SPM and DGC, thereby yielding a highly discriminative boundary.

Furthermore, we compared ArcAD against previous supervised learning approaches. Additionally, we replaced the backbone of SDNet with the identical Dinov2 [33] model used in our baseline. Specifically, we achieve an I-AUROC of 99.7% on MVTec-AD (+14.5% over DRA), 98.9% on VisA (+19.6% over DRA), 92.5% on Real-IAD (+16.5% over DRA) and 93.3% on MANTA (+14.8% over SDNet). This is primarily because conventional supervised models are highly prone to overfitting the scarce anomalous samples available in cold start scenarios. Compared to the other supervised methods, ArcAD leverages the limited anomaly data to rectify the normal distribution, achieving significantly more accurate anomaly detection and superior generalization capabilities.

Table 2: Performance under single-class setting (%). Each cell reports I-AUROC / P-AUROC / P-F₁-max.

	Method	MVTec-AD [1]	VisA [55]	Real-IAD [43]	MANTA [8]
Unsup.	SimpleNet [26]	99.1/98.9/58.8	95.6/98.7/46.4	88.3/96.3/30.8	87.5/95.6/43.4
	ReContrast [13]	98.4/98.3/62.8	93.2/96.9/41.7	90.7/98.9/44.0	92.8/96.2/35.0
	RD4AD [5]	96.6/98.1/60.0	94.5/98.3/45.5	90.0/98.8/45.8	89.7/94.7/32.8
	Dinomaly [14]	99.7/99.0/66.8	98.6/98.9/52.6	92.0/98.9/47.1	94.3/96.2/43.7
Sup.	DRA [6]	76.3/67.8/9.8	89.1/57.7/8.1	81.1/82.4/15.1	90.6/48.6/4.3
	ArcAD [†]	100.0/99.3/67.5	99.1/99.1/52.7	95.4/99.4/48.5	96.0/97.2/48.0

[†] Baseline built upon Dinomaly.

**Fig. 4:** Performance of baseline methods (Dinomaly, ReContrast, and RD4AD) equipped with the ArcAD across four datasets under the multi-class setting.

Single-class Performance. We evaluated ArcAD under the single-class setting across four datasets, as summarized in Table 2. Overall, ArcAD achieves the best overall anomaly detection performance. We achieve 100.0%/99.3%/67.5% (I-AUROC/P-AUROC/P-F₁-max) on MVTEC-AD, outperforming the best competing method by (+0.3%/+0.3%/+0.7%). On VisA dataset, our approach obtains 99.1%/99.1%/52.7%, exceeding Dinomaly by (+0.5%/+0.2%/+0.1%). The advantages are even more pronounced on challenging large-scale datasets. On Real-IAD, ArcAD reaches 95.4%/99.4%/48.5%, surpassing the top competitor by notable margins of (+3.4%/+0.5%/+1.4%). Similarly, on the MANTA dataset, ArcAD obtains 96.0%/97.2%/48.0%, yielding significant improvements of (+1.7%/+1.0%/+4.3%) over the best Dinomaly. These results demonstrate that ArcAD remains highly effective under the single-class setting.

Generality of ArcAD. To validate the generality of our proposed ArcAD, we integrated it into three mainstream reconstruction-based anomaly detection base-

Table 3: Performance comparison (I-AUROC / P-AUROC / P-F₁-max) under different anomaly ratios on Real-IAD dataset. Subscripts denote the performance improvement compared to the cold-start performance.

Method	Official	Cold Start	+ArcAD 3%	+ArcAD 5%	+ArcAD 10%
ReContrast	86.4/97.8/38.2*	85.9/97.8/40.8	88.1 _{+2.2} /98.0 _{+0.2} /44.9 _{+4.1}	88.2 _{+2.3} /98.2 _{+0.4} /45.2 _{+4.4}	88.1 _{+2.2} /98.2 _{+0.4} /45.1 _{+4.3}
RD4AD	82.4/97.3/32.7*	77.8/96.9/29.7	83.0 _{+5.2} /97.1 _{+0.2} /33.0 _{+3.3}	84.4 _{+6.6} /97.0 _{+0.1} /33.1 _{+3.4}	86.7 _{+8.9} /97.1 _{+0.2} /30.3 _{+0.6}
Dinomaly	89.3/98.8/47.7*	88.8/98.8/46.7	89.7 _{+0.9} /98.8 _{+0.0} /48.6 _{+1.9}	90.8 _{+2.0} /98.8 _{+0.0} /49.4 _{+2.7}	92.5 _{+3.7} /99.0 _{+0.2} /49.8 _{+3.1}

* *Ref: performance under standard setting from [14].*

lines (Dinomaly, ReContrast, and RD4AD) and evaluated them across four datasets under the multi-class setting. As illustrated in Fig. 4, after incorporating the ArcAD framework, all baseline models exhibit improvements across all metrics on the four datasets. Specifically, integrating ArcAD into the RD4AD framework yields I-AUROC scores of 96.4% (+5.0%), 93.8% (+3.5%), 86.7% (+8.9%), and 88.0% (+11.2%) on MVTec-AD, VisA, Real-IAD, and MANTA, respectively. This demonstrates that ArcAD does not overfit to specific methods but rather fundamentally enhances the discriminative capability of the reconstruction-based models in cold-start scenarios.

Effect of Anomaly Ratio. To evaluate the sensitivity of ArcAD to varying anomaly ratios, we conducted experiments with different anomaly ratios (3%, 5%, and 10%) on the Real-IAD dataset under the multi-class setting. As reported in Table 3, the integration of ArcAD consistently yields performance improvements across all three baseline models. Specifically, at a 10% anomaly ratio, ArcAD boosts the cold-start I-AUROC of Dinomaly, ReContrast, and RD4AD to 92.5% (+3.7%), 88.1% (+2.2%), and 86.7% (+8.9%), respectively. Remarkably, even at a mere 3% anomaly ratio, our method still improves the I-AUROC of Dinomaly, ReContrast, and RD4AD by 0.9%, 2.2%, and 5.2%, respectively. The ArcAD serves as a highly generalizable enhancement, consistently elevating the performance regardless of the anomaly ratio.

4.3 Ablation Study

Overall Ablation. To investigate the individual contribution of each module, we conducted an incremental ablation study on the Real-IAD dataset under the multi-class setting. The results are shown in Table 4. First, removing the SPM module ($\mathcal{L}_{spm}$) decreases I-AUROC, P-AUROC, and P-F₁-max by 0.8%, 0.1%, and 0.5%, respectively, verifying the necessity of constraining feature uniformity and compactness. Second, removing the DGC module ($\mathcal{L}_{dgc}$) drops the metrics by 2.3%, 0.2%, and 1.7%, demonstrating that anomaly-guided distribution calibration is crucial for rectifying the normal boundary.

Furthermore, eliminating the discriminator loss ($\mathcal{L}_{cls}$) drops performance by 2.3%, 0.1%, and 0.7%, confirming its effectiveness in boundary refinement. Notably, adding only the discriminator to the baseline actually degrades performance (*e.g.*, I-AUROC drops from 88.8% to 87.8%), proving its effectiveness intrinsically relies on synergy with the SPM and DGC.

Table 4: Ablation of individual components.

$\mathcal{L}_{recon}$	$\mathcal{L}_{spm}$	$\mathcal{L}_{cls}$	$\mathcal{L}_{dgc}$	I-AUROC	P-AUROC	P-F ₁ -max
✓				88.8	98.8	46.7
✓	✓		✓	90.2	98.9	49.1
✓		✓	✓	91.7	98.9	49.3
✓	✓	✓		90.2	98.8	48.1
✓		✓		87.8	98.5	45.5
✓	✓	✓	✓	92.5	99.0	49.8

Clustering Strategy. To investigate the effectiveness of the SPM, we replace the proposed clustering strategy with Euclidean distance-based K -means and cosine distance-based K -means, respectively. As shown in Table 5, the Sinkhorn-based clustering achieves the highest performance of 92.5%/99.0%/49.8%, outperforming K -means (Euclidean) by margins of 20.1%/6.2%/24.0% and surpassing K -means (Spherical) by 1.6%/0.1%/3.3%.

In particular, the substantial improvement over K -means (Euclidean) demonstrates the necessity of optimizing on the hypersphere to prevent mode collapse. Compared to the biased cluster assignments often produced by spherical K -means, the proposed SPM enforces an equipartition constraint, thereby achieving a more uniform and compact manifold.

Table 5: Ablation of prototyping strategies.

Prototyping Method	I-AUROC	P-AUROC	P-F ₁ -max
K -means (Euclidean)	72.4	92.8	25.8
K -means (Spherical)	90.9	98.9	46.5
Sinkhorn (Ours)	92.5	99.0	49.8

Data Augmentation. To evaluate the effectiveness of our proposed anomaly synthesis strategy, we compared it against variants. As reported in Table 6, our prototype-restricted anomaly synthesis strategy achieves the best performance of 92.5%/99.0%/49.8%. It improves upon the baseline without any anomaly synthesis by 0.9%/0.2%/0.7%, and surpasses the augmentation strategies applied in SimpleNet and DRAEM by 3.6%/0.4%/2.0% and 4.1%/0.3%/2.8%, respectively. Notably, compared to the strategy without filter (w/o Filter), which applies to all candidate anomalies, ArcAD yields a significant gain of 2.3%/0.5%/6.4%. This improvement demonstrates that utilizes all candidate anomalies is sub-optimal. By selecting pseudo-anomalies far from normal prototypes, our strat-

Table 6: Ablation of anomaly synthesis strategies.

Synthesis Strategy	I-AUROC	P-AUROC	P-F ₁ -max
w/o Anomaly synthesis	91.6	98.8	49.1
w/o Filter	90.2	98.5	43.4
SimpleNet [26]	88.9	98.6	47.8
DRAEM [51]	88.4	98.7	47.0
Ours	92.5	99.0	49.8




Fig. 6: Left: Qualitative results of anomaly localization for multi-class anomaly detection. Test images with anomalous regions are highlighted in red contours. **Right:** Feature visualization of ArcAD before and after optimization on the MANTA dataset.

egy yields hard anomalies, enabling accurate boundary delineation.

Number of Prototypes K . To evaluate the impact of the prototype quantity in the SPM module, we conducted an ablation study by varying the number of prototypes K across $\{100, 300, 500, 800, 1000\}$. The model reaches its highest I-AUROC (92.5%) at $K = 500$, with P-AUROC staying relatively stable above 99.0%. Although the highest P- F_1 -max (50.2%) occurs at $K = 800$, we ultimately select $K = 500$ as it provides the best balance between image-level and pixel-level detection performance.




Fig. 5: Ablation of the number of prototypes K in the SPM.

4.4 Visualization

Pixel-level Localization. We qualitatively visualize the localization results of our proposed ArcAD (based on Dinomally), as illustrated in Fig. 6 (left). Compared to the original Dinomally and RD4AD baselines, incorporating ArcAD significantly reduces both false positives and false negatives, thereby yielding superior segmentation performance, especially on challenging small-scale anomalies. By enforcing a compact and uniform normal feature distribution while calibrated by anomalies, the model establishes a highly discriminative decision boundary.

Feature Visualization. We visualize the learned representations on a 3D spherical manifold using t-SNE with cosine similarity applied to the ℓ_2 -normalized normal patch features extracted from the bottleneck. Specifically, we randomly select 20 clusters from the MANTA dataset under the multi-class setting to compare the feature distributions before and after optimization. As illustrated in

Fig. 6 (right), before optimization, the patch features are randomly and sparsely scattered across the hypersphere. In contrast, after applying our proposed optimization, the representations are effectively constrained, forming highly uniform and compact clusters.

5 Conclusion

We propose ArcAD, a novel plug-and-play calibration framework that effectively overcomes the cold-start bottleneck in IAD. By integrating Sinkhorn-based prototype modeling with defect-guided calibration, ArcAD explicitly rectifies the normal manifold. Extensive experiments across four datasets demonstrate that ArcAD achieves SOTA performance in both single-class and multi-class settings.

Limitations. Although our approach introduces additional computational cost during the training phase, it does not compromise inference efficiency due to its direct reliance on standard reconstruction discrepancy. Future work will explore more lightweight prototype optimization and calibration strategies to further accelerate the training process.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Grant No. 62277011), CAAI-CANN Open Fund, developed on OpenI Community, and Project of Chongqing MEITC (Grant No. YJX-2025001001009).

References

1. Bergmann, P., Fauser, M., Sattlegger, D., Steger, C.: Mvtec ad—a comprehensive real-world dataset for unsupervised anomaly detection. In: Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition. pp. 9592–9600 (2019)
2. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems* **26** (2013)
3. Davidson, T.R., Falorsi, L., et al.: Hyperspherical variational auto-encoders. In: 34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018. pp. 856–865. Association For Uncertainty in Artificial Intelligence (AUAI) (2018)
4. Defard, T., Setkov, A., Loesch, A., Audigier, R.: Padim: a patch distribution modeling framework for anomaly detection and localization. In: International conference on pattern recognition. pp. 475–489. Springer (2021)
5. Deng, H., Li, X.: Anomaly detection via reverse distillation from one-class embedding. In: Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition. pp. 9737–9746 (2022)
6. Ding, C., Pang, G., Shen, C.: Catching both gray and black swans: Open-set supervised anomaly detection. In: Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition. pp. 7388–7398 (2022)

7. Fan, L., Fan, D., Ding, Y., Wu, Y., Di, D., Pagnucco, M., Song, Y.: Grainbrain: Multiview identification and stratification of defective grain kernels. *IEEE Transactions on Industrial Informatics* **21**(5), 3627–3635 (2025)
8. Fan, L., Fan, D., Hu, Z., Ding, Y., Di, D., Yi, K., Pagnucco, M., Song, Y.: Manta: A large-scale multi-view and visual-text anomaly detection dataset for tiny objects. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 25518–25527 (2025)
9. Fan, L., Huang, J., Di, D., Su, A., Song, T., Pagnucco, M., Song, Y.: Salvaging the overlooked: Leveraging class-aware contrastive learning for multi-class anomaly detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 21419–21428 (2025)
10. Fu, L., Huang, S., Lai, Y., Liao, T., Zhang, C., Chen, C.: Beyond federated prototype learning: Learnable semantic anchors with hyperspherical contrast for domain-skewed data. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 39, pp. 16648–16656 (2025)
11. Gopal, S., Yang, Y.: Von mises-fisher clustering models. In: *International Conference on Machine Learning*. pp. 154–162. PMLR (2014)
12. Guo, J., Lu, S., Fan, L., Li, Z., Di, D., Song, Y., Zhang, W., Zhu, W., Yan, H., Chen, F., et al.: One dinomaly2 detect them all: A unified framework for full-spectrum unsupervised anomaly detection. *arXiv preprint arXiv:2510.17611* (2025)
13. Guo, J., Lu, S., Jia, L., Zhang, W., Li, H.: Recontrast: Domain-specific anomaly detection via contrastive reconstruction. *Advances in Neural Information Processing Systems* **36**, 10721–10740 (2023)
14. Guo, J., Lu, S., Zhang, W., Chen, F., Li, H., Liao, H.: Dinomaly: The less is more philosophy in multi-class unsupervised anomaly detection. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 20405–20415 (2025)
15. He, H., Bai, Y., Zhang, J., He, Q., Chen, H., Gan, Z., Wang, C., Li, X., Tian, G., Xie, L.: Mambaad: Exploring state space models for multi-class unsupervised anomaly detection. *Advances in Neural Information Processing Systems* **37**, 71162–71187 (2024)
16. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*. pp. 9729–9738 (2020)
17. Hwang, J.J., Yu, S.X., et al.: Segsort: Segmentation by discriminative sorting of segments. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 7334–7344 (2019)
18. Jeong, J., Zou, Y., Kim, T., Zhang, D., Ravichandran, A., Dabeer, O.: Winclip: Zero-/few-shot anomaly classification and segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 19606–19616 (2023)
19. Ke, G., Xue, H.: Hyperspherical latents improve continuous-token autoregressive generation. *arXiv preprint arXiv:2509.24335* (2025)
20. Kobayashi, T.: T-vmf similarity for regularizing intra-class feature distribution. In: *Computer Vision and Pattern Recognition*. pp. 6616–6625 (2021)
21. Li, C.L., Sohn, K., Yoon, J., Pfister, T.: Cutpaste: Self-supervised learning for anomaly detection and localization. In: *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*. pp. 9664–9674 (2021)
22. Li, C., Shi, Y., Hu, J., Zhu, X.X., Mou, L.: Scale-aware contrastive reverse distillation for unsupervised medical anomaly detection. In: *The Thirteenth International Conference on Learning Representations*

23. Li, J.L., Chang, M.C., Chen, W.C.: Learning multi-manifold embedding for out-of-distribution detection. In: European Conference on Computer Vision. pp. 417–433. Springer (2024)
24. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2980–2988 (2017)
25. Ling, H., Guo, J., Tao, Z., Cao, Y., Di, D., Xu, H., Su, X., Song, Y., Fan, L.: Adnet: A large-scale and extensible multi-domain benchmark for anomaly detection across 380 real-world categories. arXiv preprint arXiv:2511.20169 (2025)
26. Liu, Z., Zhou, Y., Xu, Y., Wang, Z.: Simplenet: A simple network for image anomaly detection and localization. In: Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition. pp. 20402–20411 (2023)
27. Loshchilov, I., Hsieh, C.P., Sun, S., Ginsburg, B.: ngpt: Normalized transformer with representation learning on the hypersphere. In: The Thirteenth International Conference on Learning Representations
28. Lu, H., Gong, D., et al.: Learning with mixture of prototypes for out-of-distribution detection. arXiv preprint arXiv:2402.02653 (2024)
29. Mardia, K.V., Jupp, P.E.: Directional statistics. John Wiley & Sons (2009)
30. McIntosh, D., Albu, A.B.: Inter-realization channels: Unsupervised anomaly detection beyond one-class classification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6285–6295 (2023)
31. Ming, Y., Sun, Y., et al.: How to exploit hyperspherical embeddings for out-of-distribution detection? arXiv preprint arXiv:2203.04450 (2022)
32. Moon, T.K.: The expectation-maximization algorithm. IEEE Signal processing magazine **13**(6), 47–60 (1996)
33. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. Transactions on Machine Learning Research (2024)
34. Pang, G., Ding, C., Shen, C., Hengel, A.v.d.: Explainable deep few-shot anomaly detection with deviation networks. arXiv preprint arXiv:2108.00462 (2021)
35. Pang, G., Shen, C., Van Den Hengel, A.: Deep anomaly detection with deviation networks. In: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 353–362 (2019)
36. Park, H., Noh, J., Ham, B.: Learning memory-guided normality for anomaly detection. In: Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition. pp. 14372–14381 (2020)
37. Qin, X., Cunningham, P., et al.: Online trans-dimensional von mises-fisher mixture models for user profiles. Journal of Machine Learning Research **17**(200), 1–51 (2016)
38. Rippel, O., Mertens, P., König, E., Merhof, D.: Gaussian anomaly detection by modeling the distribution of normal data in pretrained deep features. IEEE Transactions on Instrumentation and Measurement **70**, 1–13 (2021)
39. Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., Gehler, P.: Towards total recall in industrial anomaly detection. In: Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition. pp. 14318–14328 (2022)
40. Rudolph, M., Wandt, B., Rosenhahn, B.: Same same but different: Semi-supervised defect detection with normalizing flows. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 1907–1916 (2021)

41. Ruff, L., Kauffmann, J.R., Vandermeulen, R.A., Montavon, G., Samek, W., Kloft, M., Dietterich, T.G., Müller, K.R.: A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE* **109**(5), 756–795 (2021)
42. Tabernik, D., Šela, S., Skvarč, J., Skočaj, D.: Segmentation-based deep-learning approach for surface-defect detection. *Journal of Intelligent Manufacturing* **31**(3), 759–776 (2020)
43. Wang, C., Zhu, W., Gao, B.B., Gan, Z., Zhang, J., Gu, Z., Qian, S., Chen, M., Ma, L.: Real-iad: A real-world multi-view dataset for benchmarking versatile industrial anomaly detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 22883–22892 (2024)
44. Wang, F., Zhang, T., Wang, Y., Qiu, Y., Liu, X., Guo, X., Cui, Z.: Distribution prototype diffusion learning for open-set supervised anomaly detection. In: *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*. pp. 20416–20426 (2025)
45. Wang, T., Isola, P.: Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In: *International conference on machine learning*. pp. 9929–9939. PMLR (2020)
46. Wei, S., Jiang, J., Xu, X.: Uninet: A contrastive learning-guided unified framework with feature selection for anomaly detection. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 9994–10003 (2025)
47. Xu, J., Li, S., Deng, A., et al.: Probabilistic knowledge distillation of face ensembles. In: *Computer Vision and Pattern Recognition*. pp. 3489–3498 (2023)
48. Yao, X., Li, R., Qian, Z., Wang, L., Zhang, C.: Hierarchical gaussian mixture normalizing flow modeling for unified anomaly detection. In: *European Conference on Computer Vision*. pp. 92–108. Springer (2024)
49. Yao, X., Li, R., Zhang, J., Sun, J., Zhang, C.: Explicit boundary guided semi-push-pull contrastive learning for supervised anomaly detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 24490–24499 (2023)
50. You, Z., Cui, L., Shen, Y., Yang, K., Lu, X., Zheng, Y., Le, X.: A unified model for multi-class anomaly detection. *Advances in Neural Information Processing Systems* **35**, 4571–4584 (2022)
51. Zavrtnik, V., Kristan, M., Skočaj, D.: Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8330–8339 (2021)
52. Zhang, J., Chen, X., Wang, Y., Wang, C., Liu, Y., Li, X., Yang, M.H., Tao, D.: Exploring plain vit reconstruction for multi-class unsupervised anomaly detection. *arXiv preprint arXiv:2312.07495* (2023)
53. Zhang, X., Li, S., Li, X., Huang, P., Shan, J., Chen, T.: Destseg: Segmentation guided denoising student-teacher for anomaly detection. In: *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*. pp. 3914–3923 (2023)
54. Zhu, J., Ding, C., Tian, Y., Pang, G.: Anomaly heterogeneity learning for open-set supervised anomaly detection. In: *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*. pp. 17616–17626 (2024)
55. Zou, Y., Jeong, J., Pemula, L., Zhang, D., Dabeer, O.: Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. In: *European Conference on Computer Vision*. pp. 392–408. Springer (2022)