

# TailedCore: Few-Shot Sampling for Unsupervised Long-Tail Noisy Anomaly Detection

Jung, Yoon G.; Park, Jaewoo; Yoon, Jaeho; Peng, Kuan-Chuan; Kim, Wonchul; Teoh, Andrew B. J.; Camps, Octavia

TR2025-077 June 07, 2025

## Abstract

We aim to solve unsupervised anomaly detection in a practical challenging environment where the normal dataset is both contaminated with defective regions and its product class distribution is tailed but unknown. We observe that existing models suffer from tail-versus-noise trade-off where if a model is robust against pixel noise, then its performance deteriorates on tail class samples, and vice versa. To mitigate the issue, we handle the tail class and noise samples independently. To this end, we propose TailSampler, a novel class size predictor that estimates the class cardinality of samples based on a symmetric assumption on the class-wise distribution of embedding similarities. TailSampler can be utilized to sample the tail class samples exclusively, allowing to handle them separately. Based on these facets, we build a memory-based anomaly detection model TailedCore, whose memory both well captures tail class information and is noise-robust. We extensively validate the effectiveness of TailedCore on the unsupervised long-tail noisy anomaly detection setting, and show that TailedCore outperforms the state-of-the-art in most settings. Code is available in TailedCore.

*IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2025*



# TailedCore: Few-Shot Sampling for Unsupervised Long-Tail Noisy Anomaly Detection

Yoon Gyo Jung<sup>\*1</sup> Jaewoo Park<sup>\*2</sup> Jaeho Yoon<sup>\*3</sup> Kuan-Chuan Peng<sup>4</sup>  
Wonchul Kim<sup>2</sup> Andrew Beng Jin Teoh<sup>3</sup> Octavia Camps<sup>†1</sup>

<sup>1</sup>Northeastern University <sup>2</sup>AiV Co. <sup>3</sup>Yonsei University

<sup>4</sup>Mitsubishi Electric Research Laboratories

## Abstract

We aim to solve unsupervised anomaly detection in a practical challenging environment where the normal dataset is both contaminated with defective regions and its product class distribution is tailed but unknown. We observe that existing models suffer from tail-versus-noise trade-off where if a model is robust against pixel noise, then its performance deteriorates on tail class samples, and vice versa. To mitigate the issue, we handle the tail class and noise samples independently. To this end, we propose TailSampler, a novel class size predictor that estimates the class cardinality of samples based on a symmetric assumption on the class-wise distribution of embedding similarities. TailSampler can be utilized to sample the tail class samples exclusively, allowing to handle them separately. Based on these facets, we build a memory-based anomaly detection model TailedCore, whose memory both well captures tail class information and is noise-robust. We extensively validate the effectiveness of TailedCore on the unsupervised long-tail noisy anomaly detection setting, and show that TailedCore outperforms the state-of-the-art in most settings. Code is available in [TailedCore](#).

## 1. Introduction

In the complex landscape of anomaly detection, the challenge often lies in navigating practical scenarios that feature a diverse and imbalanced data distribution [6, 8, 11, 21]. We delve into a challenging yet highly realistic scenario in anomaly detection, where the training dataset, comprising multiple product classes [39], is beset by two significant complications: *contamination with noise* [20] and *presence of tailed class distributions* [25]. In this context, ‘tail’ (few-shot) classes have limited few-shot samples, unlike ‘head’ (many-shot) classes which are data-rich but might contain samples with defects. Crucially, in this unsupervised set-

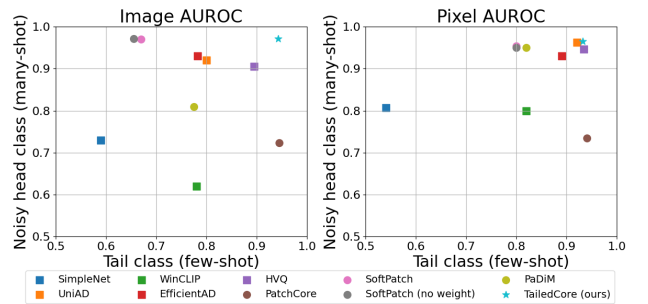


Figure 1. Tail class (x-axis) and noisy head class (y-axis) performance comparison. The tail-versus-noise trade-off is shown across memory-based anomaly detection models (circles), and is more indicative in anomaly classification task evaluated by image-level AUROC (left).

ting, the class information of these products remains unknown to the model trainer. This setting mirrors a common predicament in large-scale industrial applications, where companies are required to manage vast production lines encompassing a wide array of product types, each with varying production rates.

A critical issue arises in this scenario, which we term the “tail-versus-noise dilemma.” The dilemma underscores a fundamental trade-off where *models robust to class imbalance overfit to defect noises*, likely misclassifying many defects as normal. In contrast, existing *models that are robust against defect noise tend to underfit*, missing vital details in tail classes. Fig. 1 shows this tail-versus-noise trade-off which is prevalent through existing anomaly detection algorithms. This intrinsic characteristic comes from the fact that few-shot class features are equivalent to the noisy abnormal features in terms of their statistical occurrence. Consequently, existing models find it difficult to strike a balance between avoiding overfitting to noise and effectively capturing the nuances of less represented classes.

To resolve this challenge, we propose to handle tail classes independently from head classes by exclusively sampling them. To this end, we devise a novel class size predictor TailSampler that estimates the cardinality of the class of any training sample. For accurate estimation, TailSampler estimates the class sizes based on a reflective symmetry between inter-class

<sup>\*</sup> equal contribution <sup>†</sup> corresponding author

and intra-class similarity distributions in the embedding space. The estimated class sizes of the training samples then enable to sample tail class features exclusively from the training set.

By utilizing the aggregated few-shot class instances, we devise a memory-based anomaly detector TailedCore, which is built upon the noise-discriminated model by augmenting the tail class patch features sampled with TailSampler. As the memory of TailedCore is both clean from noisy defect patches and preserves the few-shot class information, it exhibits robustness against both noise contamination and class imbalance of the training set.

Our contributions are summarized as follows:

- We explore a practical and challenging anomaly detection scenario, where a single detection model is trained on multi-class product training samples that involve noise contamination and whose class distribution is long-tailed.
- We propose a memory-based anomaly detector TailedCore whose patch memory bank is both noise-free and augmented with the representative patches of few-shot class instances. To build TailedCore, we introduce TailSampler that can aggregate few-shot class instances exclusively. TailSampler is based on a novel class size predictor that estimates the class size of training samples based on the reflective symmetry between inter-class and intra-class class-wise embedding similarity distributions.
- We conduct extensive analyses and comparative studies on both TailedCore and few-shot sampler along with the proposed class size predictor on the proposed unsupervised long-tail noisy anomaly detection benchmark.

The codes for TailedCore along with the data generation will be open sourced. We use the term ‘few-shot’ and ‘tail’ interchangeably to denote the class with a small number of samples.

## 2. Related works

### 2.1. Anomaly detection in extreme scenarios

One vitally important aspect of anomaly detection [9, 34] is the environment where these models are tested. This exploration has led to the identification of several challenging scenarios, including multi-class settings with training data spanning multiple product types [26, 39], zero-shot and few-shot settings with minimal training images [15, 19, 23, 38, 43], and noisy settings featuring uncleaned training data with defective images [7, 20, 28].

**Multi-class anomaly detection** In contrast to single-product settings, anomaly detection models trained on multi-class datasets often encounter “shortcut learning” [39], leading to low anomaly scores for normal images. This issue is particularly pronounced in reconstruction-based models, which tend to reconstruct all samples, regardless of their normality or abnormality. To resolve it, UniAD [39] employs neighbor masked attention within a transformer model, focusing attention on local parts to prevent shortcut learning. HVQ [26] utilizes vector quantization to discretize the latent space, restricting the network’s recon-

struction fidelity on normal data distributions. UniFormally [22] increases sensitivity to irrelevant patterns by masking less relevant patches and employing top-k matchings with a pretrained foundation model. These multi-class anomaly detection methods do not consider either long-tail classes or contaminated training data setups, and do not work well in the combined “noisy long-tail” setup, which is shown in our experimental results.

**Zero/few-shot anomaly detection** The inception of few-shot learning benchmarks in anomaly detection marked a significant advancement, with methods leveraging template matching [18] and graph representations [38] for minimal training data scenarios. AnomalyCLIP [43], WinCLIP [19], and AnomalyGPT [17], on the other hand, detect anomalies in the few-shot classes based on the acquired knowledge of pretrained foundation models such as CLIP [33] and GPT [32]. However, these methods often require sophisticated prompt engineering or delicate prompt-tuning, and have yet to achieve optimal precision and recall in the single-class training setting with a sufficient amount of training samples. In contrast, our proposed TailedCore needs none of these process.

**Noise anomaly detection** Dealing with training datasets contaminated with defective samples presents another extreme learning scenario. SoftPatch [20] and InReaCh [28] focus on purifying memory coresets through outlier detection algorithms and distance association filtering, respectively. IGD [7] adopts a robust Gaussian modeling approach, addressing both contaminated training data and small dataset challenges. Yet, these methods do not fully consider scenarios involving both data contamination and class imbalance, unlike our proposed TailedCore.

### 2.2. Long-tail noisy learning

In the realm of image classification, several studies have focused on learning scenarios involving training sets with label contamination and class imbalance. To address these challenges, the existing literature relies on the robustness and unsupervised aspects of self-supervised learning [14, 27, 37, 41]. The learned embeddings of self-supervised frameworks are less impacted by label contamination and class imbalances. However, these works differ from ours in that they deal with label-level contamination rather than pixel-level defects. Furthermore, these methods require explicit supervised labels to remediate the contamination. In contrast, our work targets noise in pixels without any class label information.

### 2.3. Few-shot/outlier sampling

Both classical clustering methods and outlier detection algorithms can be applied to our task, as they facilitate the identification of few-shot tail class samples, either indirectly or directly. Unsupervised clustering algorithms [2, 13, 16] can estimate the cluster sizes but struggle to identify small sample clusters in the imbalanced data. DBSCAN [13], on the other hand, captures both clusters and outliers. This however can be problematic when attempting to identify clusters with only one

sample (1-shot) as it may be confused between few-shot samples and outliers. Unlike clustering methods, outlier detection algorithms [1, 5, 24, 35] are adept at identifying 1-shot samples which inherently resemble outliers. However, their effectiveness diminishes when a few-shot class contains more than one sample, as these no longer fit the typical outlier profile, potentially leading to incorrect identification of few-shot class samples. Both algorithm types, due to their inherent design, struggle to accurately predict class/cluster sizes in datasets that are both imbalanced and contaminated. Our proposed class size predictor specifically addresses this issue, aiming to reliably and accurately identify few-shot class samples and enable their exclusive sampling. Ultimately, the identified few-shot class samples are augmented to the noise-discriminated patch set, constituting the memory-based anomaly detection model TailedCore, which is robust against both class imbalance and training set noise.

### 3. Background

#### 3.1. Memory-based anomaly detection

To train a memory-based anomaly detector, a feature extractor  $f$  (e.g., WideResNet) extracts a feature map  $\phi_i \in \mathbb{R}^{C \times H \times W}$  from each training input image  $x_i$ . Then, a coreset selection algorithm  $\mathcal{S}_{core}$  is applied to the set  $P = \{\phi_i^{(h,w)}\}$  of patch features  $\phi_i^{(h,w)}$ , resulting in a memory bank

$$M_{PatchCore} := \mathcal{S}_{core}(P) \subset P \quad (1)$$

that contains representative patch features of  $P$ . For instance, PatchCore opts for greedy sampling strategy to obtain the coreset. In the inference stage, the anomaly score  $s^{(h,w)}$  of the testing image  $x$  is computed for each patch feature  $\phi^{(h,w)}$  by measuring its distance to the memory

$$score^{(h,w)} = dist(\phi^{(h,w)}, M), \quad (2)$$

where  $dist$  is the nearest neighbor distance in [34]. The pixel-level anomaly score is computed by up-scaling the score map  $score = [score^{(h,w)}] \in \mathbb{R}^{H,W}$ , and its image-level anomaly score is defined as its pixel maximum  $\max_{h,w} score^{(h,w)}$ .

#### 3.2. Noise discrimination of memory

The training set  $X = \{x_i\}$  may be contaminated with defect samples. The naively attained memory bank then contains abnormal patch features, and assigns a low anomaly score on defect images, resulting in poor performance. To resolve this, [20] applies a noise discrimination algorithm  $\mathcal{S}_{clean}$  on the patch feature set  $P$  prior to coreset sampling:

$$\mathcal{S}_{clean}(P) = \{\phi_i^{(h,w)} : outlier(\phi_i^{(h,w)}) < \tau_{0.15}\}, \quad (3)$$

aggregating only the clean normal patches with low outlier scores. SoftPatch [20] exploits local outlier factor (LOF) for outlier scoring, and selects the threshold  $\tau_{0.15}$  to remove only 15% of whole patches. In inference, SoftPatch utilizes the memory set  $M_{SoftPatch} = \mathcal{S}_{core}(\mathcal{S}_{clean}(P))$

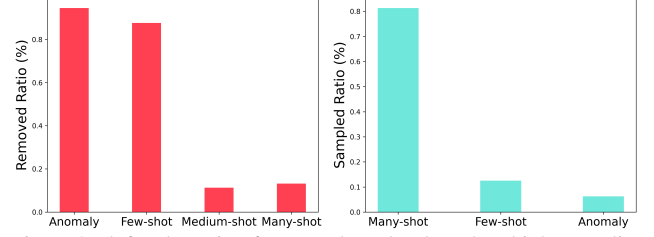


Figure 2. (left) The ratio of removed patches based on highest outlier scores by Eq. (3), which shows that most of few-shot class patches are lost. (right) The ratio of sampled patches by greedy coreset sampling from PatchCore which favors both few-shot and anomaly samples.

## 4. Problem and motivation

### 4.1. Problem

We consider a practical anomaly detection setting where the normal training data  $X = \{x_i\}$  is both multi-class with tail distribution and noise-contaminated. Particularly, each  $x_i$  is paired with the product class label  $y_i \in \mathcal{Y} = \{1, 2, \dots, |\mathcal{Y}|\}$  (e.g., hazelnut, screw, pill, etc.) and an anomaly label  $y_i^a \in \mathcal{Y}^a := \{0, 1\}$  with 1 indicating anomaly. However, this labeling information is *unknown* to the model trainer. The class distribution is long-tailed such that there are tail (i.e., few-shot) classes whose number of training samples is very low; i.e.,

$$|C_k| = |\{x_i : y_i^c = k\}| \leq K \quad (4)$$

with a very small  $K$  (e.g., 1, 4, or 20), where  $|\cdot|$  indicates the number of elements in the set. Moreover, some of the training samples in  $X$  are anomalies, containing defect blobs in their images. We assume the tail class samples are all normal since in the unsupervised setting its normality is determined by its majority. Overall, the training set  $X$  with the above descriptions constitutes *unsupervised tailed noisy anomaly detection* setting.

### 4.2. Motivation: Tail-versus-noise trade-off

An ideal model to handle a long-tail noisy dataset should selectively sample the few-shot samples while removing noisy samples. However, we observe that the existing models suffer from a trade-off between class-imbalance and noise robustness. For memory-based noise-discriminative models in particular (e.g., SoftPatch), the methods focus on sampling the majority patch features, which often ignores less dominating but important patches.

As shown in Fig. 2 left, the noise discrimination process removes both defect and few-shot class patches, losing most of the information in the tail distribution. On the other hand, PatchCore samples few-shot samples well due to greedy coreset sampling which captures maximally different features. This sampling principle, however, favors the noise patch features as well as shown in Fig. 2 (right) and the low performance on noisy head classes in Fig. 1. Previous methods, hence, face the tail-versus-noise dilemma described above.

To resolve this, we propose to handle the tail classes independently from the head classes that potentially contain defective local regions. This, however, needs exclusive sampling of tail classes. To achieve this, we use the globally average pooled embedding vector, which is invariant against local pixel variation and rather exclusively corresponds to the object-centric nature of image, *i.e.*, class.

## 5. Method

We aim to build a coreset-based anomaly detection model using patch features. We require its memory to aptly capture tail class information and also to be absent of noise patches as much as possible. Particularly, we augment the noise discriminated memory  $M_{clean}$  with the tail class memory  $M_{tail}$ , which contains the patches of few-shot class samples exclusively:

$$M_{TailedCore} = M_{clean} \cup M_{tail}. \quad (5)$$

The main challenge lies on sampling the tail class patches exclusively. To this end, we devise a few-shot sampler by estimating the class size of samples based on its embedding (*e.g.*, global average pooling layer of the encoder). Using the estimated class sizes of training samples, the few-shot sampler enables us to exclusively sample tail class samples and the patches thereof.

### 5.1. Few-shot sampler: TailSampler

TailSampler first estimates the class size of every training sample. Then, by automatically determining the maximum number of tail classes, denoted by  $y_{max}$ , TailSampler collects the tail class samples' patch features exclusively.

#### 5.1.1. Class size predictor

To sample patches exclusively from few-shot class, we identify samples from few-shot classes by estimating their "class sizes", namely the size of class each sample belongs to. To accurately measure the class size, we hypothesize as follows: In the embedding space, samples within an appropriate angle are majorly of the same class, and this angle is likely to be the half of maximum angle to another embedding. In addition, the neighborhoods of neighborhoods within the half angle give accurate class size estimate. Based on these hypotheses, we proceed as follows:

First, we adaptively set angle  $\alpha_i$  for each embedding  $e_i$  from training samples  $x_i$ . The coverage of angle  $\alpha_i$  decides the neighbors where a large value would assign many neighbors. We set the angle to be half of its maximum within the train set, and hypothesize that most of neighbors within the region is likely of the same class. Particularly, let  $H_i = \{e \in Z : \angle(e_i, e) \leq m_i/2\}$  denote the half-angle region where the maximum angle is  $m_i := \max_{e \in Z} \angle(e_i, e)$  and the set  $Z$  contains all train embeddings. Then, the adaptive angle  $\alpha_i$  is defined to contain  $p$ -th percentile of the half-angle region:

$$\alpha_i = \angle(e_i, e_{(p \cdot |H_i|)}), \quad (6)$$

where the index  $(j)$  of  $e_{(j)}$  is sorted in the increasing order of the angle  $\angle(e_i, e_{(j)})$ . We set  $p=0.85$  in all experiments unless specified otherwise.

After setting the adaptive angle  $\alpha_i$  for each train embedding  $e_i$ , we estimate its class size based on neighborhoods of neighborhoods. Let  $N_{\alpha}(e_i) = \{e \in Z : \angle(e_i, e) \leq \alpha\}$  denote the neighborhood of  $e_i$ , which is the set of all train embedding  $e$  within the angle  $\alpha$  of  $e_i$ . Then, the class size is estimated by the mode of the sizes of neighborhoods of the neighborhood:

$$\kappa_i = \text{mode}_{e \in N_{\alpha_i}(e_i)}(|N_{\alpha(e)}(e)|), \quad (7)$$

where  $\alpha(e)$  is the adaptive angle with respect to the embedding  $e$  belonging to the neighborhood  $N_{\alpha_i}(e_i)$  of embedding  $e_i$ . Using the neighborhoods of neighborhood and mode on it gives more robust estimation of the class size than mere count of the direct neighborhood of  $e_i$ .

The proposition in Supp. 9 shows that the angle of neighborhood given by Eq. (6) corresponds to the decision boundary that maximizes the inter-class separation under regularities.

#### 5.1.2. Estimation of maximum size of tail class

After estimating the class size of each training sample by  $\kappa_i$ , we determine the maximum size for the tail classes based on the elbow technique [36]. In particular, by utilizing the sample-wise class sizes  $\kappa_i$ , we first obtain rough estimation of the size of each class set  $\eta_y \approx |C_y|$ . Then, we estimate the maximum size of tail classes  $K_{max}$  by finding a point where the class size most abruptly changes. This optimal shift point is found by the elbow technique. Concretely, we first estimate the size of each class  $y$ ,  $\eta_y \approx |C_y|$ , inductively by

$$\begin{aligned} \eta_1 &= \text{round}\left(\frac{1}{\kappa_{(1)}} \sum_{i=1}^{\kappa_{(1)}} \kappa_{(i)}\right), \\ \eta_{y+1} &= \text{round}\left(\frac{1}{\kappa_{(\eta_y+1)}} \sum_{i=\eta_y+1}^{\min(\kappa_{(\eta_y+1)}, |X|)} \kappa_{(i)}\right), \end{aligned} \quad (8)$$

where the index  $(i)$  of  $\kappa_{(i)}$  is sorted in the increasing order  $\kappa_{(1)} \leq \dots \leq \kappa_{(|X|)}$ .

After we acquire the set  $\{(y, \eta_y) : y \in (1, \dots, |\mathcal{Y}|)\}$  through the inductive process above, we define the maximum size of tail classes as the elbow point of the estimated sizes of classes  $\eta_y$  since at this point the class size increases abruptly. The elbow is computed by linearly connecting  $(1, \eta_1)$  and  $(|\mathcal{Y}|, \eta_{|\mathcal{Y}|})$ , and finding the maximum orthogonal distance onto the line from the set of points  $\{(y, \eta_y) : y \in (1, \dots, |\mathcal{Y}|)\}$  as shown below:

$$K_{max} := \eta_{y_{max}} =: \text{elbow}(\eta_1, \dots, \eta_{|\mathcal{Y}|}) \quad (9)$$

where

$$y_{max} = \underset{y \in \{1, \dots, |\mathcal{Y}|\}}{\text{argmax}} \frac{|ay - \eta_y + b|}{\sqrt{a^2 + 1}} \quad (10)$$

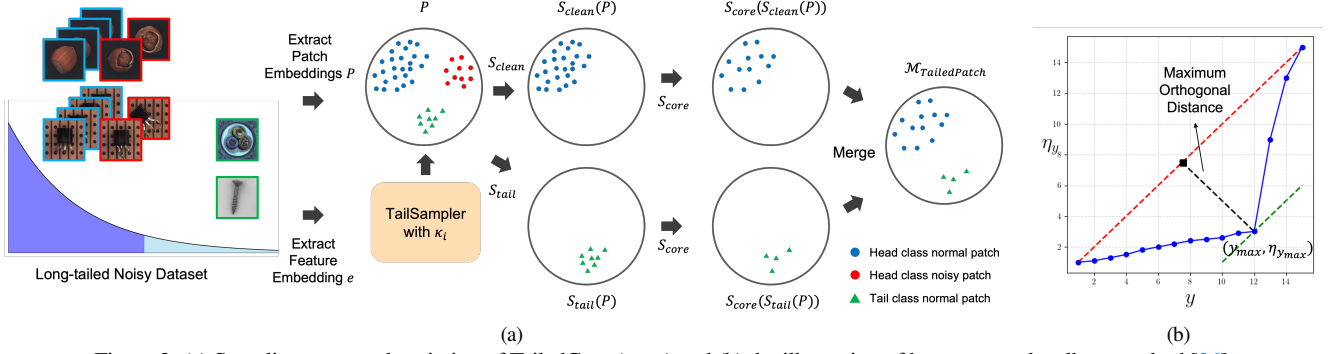


Figure 3. (a) Sampling process description of TailedCore (ours) and (b) the illustration of how we use the elbow method [36].

and where the line  $\eta_y = ay + b$  intersects  $(1, \eta_1)$  and  $(|\mathcal{Y}|, \eta_{|\mathcal{Y}|})$  with  $a = (\eta_{|\mathcal{Y}|} - \eta_1) / (|\mathcal{Y}| - 1)$  and  $b = \eta_1 - m$ . Fig. 3 (b) illustrates Eq. (10). We choose the elbow point  $(y_{max}, \eta_{y_{max}})$  as the maximum size of tail classes  $K_{max}$ .

The maximum size of tail classes  $K_{max}$  allows us to sample the patch features of tail classes exclusively:

$$S_{tail}(P) = \{\phi_i^{(h,w)} : \kappa_i \leq K_{max}\}. \quad (11)$$

with  $\kappa_i$  given by Eq. (7).

## 5.2. Proposed method: TailedCore

Our anomaly detection model is built upon the memory bank constructed as below. First, we sample the base memory bank as in [20] with noise-discrimination described in Eq. (3). Then, the base memory is *augmented with TailSampler*  $S_{tail}$  (Eq. (11)). Overall, TailedCore’s memory bank is defined by  $M_{TailedCore} = S_{core}(S_{clean}(P)) \cup S_{core}(S_{tail}(P))$ , where  $S_{core}$  is the greedy coresets sampling in Eq. (1),  $S_{clean}$  is the noise discrimination process that filters out defect patches based on their outlier scores (Eq. (3)), and the few-shot sampler  $S_{tail}$  augments the tail class patches exclusively on the noise-discriminated base memory bank. The whole process is shown in Fig 3. As TailedCore’s patch memory bank is both noise-clean and contains the representative features of few-shot classes, it is robust against both class imbalance and normal sample contamination. In inference, the anomaly score is obtained by Eq. (2) as in the conventional memory-based detector.

## 6. Experiments

Section 6 consists of dataset description, comparison to the SOTA models, and the detailed ablation study and analysis of TailedCore’s components.

### 6.1. Datasets for unsupervised long-tail noisy anomaly detection

We devise the long-tail noisy anomaly detection datasets by modifying the class-wise sample distributions of the widely used anomaly detection datasets MVTEC-AD [3] and VisA [44]. We consider three different types of long-tail distributions. Particularly, datasets with step-shaped tail distribution are devised

tail type	Pareto			step ( $K=4$ )			step ( $K=1$ )		
	$C_t$	$C_h$	all	$C_t$	$C_h$	all	$C_t$	$C_h$	all
PaDiM [9] ICPR’21	82.45	80.95	82.06	77.47	81.28	79.19	71.54	81.75	75.63
HVQ [26] NeurIPS’23	83.46	80.23	82.99	82.01	85.50	83.56	74.15	90.15	80.55
WinCLIP [19] CVPR’23	89.35	90.11	90.37	91.60	88.21	90.37	91.80	88.23	90.37
AnomalyCLIP [43] ICLR’24	90.93	90.98	91.48	91.82	90.83	91.48	91.21	91.90	91.48
PatchCore [34] CVPR’22	93.33	87.59	89.18	92.19	71.18	83.83	86.36	70.48	80.01
SoftPatch [20] NeurIPS’22	84.68	86.95	87.71	67.65	97.54	79.64	60.66	97.49	75.40
<b>TailedCore (ours)</b>	<b>96.55</b>	<b>95.24</b>	<b>96.12</b>	<b>95.82</b>	<b>95.34</b>	<b>95.71</b>	<b>93.54</b>	<b>95.77</b>	<b>94.43</b>

Table 1. Anomaly classification on MVTEC-AD with image-level AUROC (%). We report the mean over 5 random seeds for each measurement. Notations:  $C_h$  /  $C_t$ : head / tail classes.

tail type	Pareto			step ( $K=4$ )			step ( $K=1$ )		
	$C_t$	$C_h$	all	$C_t$	$C_h$	all	$C_t$	$C_h$	all
PaDiM [9] ICPR’21	70.70	83.35	78.64	60.65	88.93	72.43	55.98	86.75	68.80
HVQ [26] NeurIPS’23	73.47	84.03	68.25	68.25	89.30	77.02	61.57	80.40	69.42
WinCLIP [19] CVPR’23	73.25	76.92	75.47	75.98	74.76	75.47	78.80	70.80	75.47
AnomalyCLIP [43] ICLR’24	81.96	82.48	82.05	82.28	81.74	82.05	<b>83.26</b>	80.34	82.05
PatchCore [34] CVPR’22	86.11	85.73	85.59	83.53	67.51	76.85	79.33	68.56	74.84
SoftPatch [20] NeurIPS’22	78.04	92.16	86.56	59.70	<b>95.97</b>	74.81	52.61	<b>94.17</b>	69.92
<b>TailedCore (ours)</b>	<b>87.55</b>	<b>93.06</b>	<b>90.85</b>	<b>85.16</b>	<b>95.91</b>	<b>89.64</b>	<b>82.97</b>	<b>94.11</b>	<b>87.61</b>

Table 2. Anomaly classification on VisA with image-level AUROC (%). The format and evaluation protocol are the same as Tab. 1.

tail type	Pareto			step ( $K=4$ )			step ( $K=1$ )		
	$C_t$	$C_h$	all	$C_t$	$C_h$	all	$C_t$	$C_h$	all
PaDiM [9] ICPR’21	90.11	92.66	91.43	82.53	95.29	87.67	78.80	95.54	85.50
HVQ [26] NeurIPS’23	93.63	86.85	90.55	90.73	92.58	91.53	86.36	95.20	89.90
WinCLIP [19] CVPR’23	82.03	84.06	82.29	80.60	84.63	82.29	80.16	85.48	82.29
AnomalyCLIP [43] ICLR’24	91.24	91.69	91.08	89.96	92.66	91.08	89.34	93.68	91.08
PatchCore [34] CVPR’22	93.56	87.98	89.93	93.54	72.09	85.19	92.02	71.35	83.75
SoftPatch [20] NeurIPS’22	92.19	93.83	93.41	80.98	<b>96.49</b>	87.24	70.34	<b>96.89</b>	80.99
<b>TailedCore (ours)</b>	<b>96.08</b>	<b>95.01</b>	<b>95.29</b>	<b>95.56</b>	93.20	<b>94.74</b>	<b>94.19</b>	93.70	<b>93.99</b>

Table 3. Anomaly segmentation on MVTEC-AD with pixel-level AUROC (%). We report the mean over 5 random seeds for each measurement. Notations:  $C_h$  /  $C_t$ : head / tail classes.

by making the particular classes to have only a small number of samples; namely  $K=1$ , and  $K=4$ , imposing extreme class imbalance on the datasets. We name these two datasets by ‘step ( $K=1$ )’ and ‘step ( $K=4$ )’. On the other hand, another realistic type of tail-distribution is considered by making the class distribution to follow the Pareto distribution with the shape parameter 0.6. For the step-tailed distribution datasets, we impose 40% of classes to be head classes, containing the original number of training samples. Thus, in MVTEC-AD, 6 of 15 classes are head classes. For the Pareto tailed distribution, we regard as few-shot (*i.e.*, tail classes) the classes whose number of samples

tail type class type	Pareto			step ( $K=4$ )			step ( $K=1$ )		
	$C_t$	$C_h$	all	$C_t$	$C_h$	all	$C_t$	$C_h$	all
PaDiM [9] ICPR'21	89.02	95.10	82.81	83.90	97.36	89.51	82.57	96.57	88.40
HVQ [26] NeurIPS'23	95.27	<b>97.60</b>	96.71	93.88	<b>98.34</b>	95.74	90.58	95.51	92.63
WinCLIP [19] CVPR'23	71.94	73.97	73.19	74.60	71.21	73.19	73.81	72.32	73.19
AnomalyCLIP [43] ICLR'24	95.60	95.46	95.51	95.54	95.48	95.51	<b>96.16</b>	94.60	95.51
PatchCore [34] CVPR'22	96.84	87.99	91.13	95.39	62.96	81.88	94.11	65.30	82.10
SoftPatch [20] NeurIPS'22	93.20	96.74	95.27	83.95	97.10	89.43	80.73	96.82	87.43
TailedCore (ours)	<b>97.98</b>	<b>97.25</b>	<b>97.48</b>	<b>96.80</b>	97.02	<b>96.89</b>	<b>96.12</b>	<b>97.39</b>	<b>96.65</b>

Table 4. Anomaly segmentation on VisA with pixel-level AUROC (%). The format and evaluation protocol are the same as Tab. 3.

is less than 20 [25]. To impose noisy condition on the datasets, the anomaly samples (*i.e.*, the samples with defective regions) are added to the training samples such that the noisy samples constitute 10% of the training set following the overlap protocol in SoftPatch [20] for fair comparison. In all cases, the testing sets of datasets remain the same as the original MVTecAD and VisA. A more detailed configuration is given in Supp. 11.

## 6.2. Comparison to the state-of-the-art (SOTA)

We compare TailedCore to the SOTA anomaly detection methods in the unsupervised long-tail noisy settings. We carefully select our baselines. HVQ excels in multi-class learning. PaDiM and SoftPatch are robust against noises as PaDiM is based on Gaussian modeling of the data, while SoftPatch explicitly removes the defect patches by outlier scoring. Using foundation models, WinClip and AnomalyClip are strong for zero-shot learning. PatchCore are near-perfect in the single-class anomaly detection settings. We show the full experimental results in Supp. 12. All models are trained in the multi-class setting; namely, one model for one dataset with all classes rather than one model per class. The detection performance, however, is measured for each product class under image-level and pixel-level AUROC, and averaged for reporting and comparison. For comparison, we consider models that are specifically robust against noise, strong on multi-class training, and foundation models that can exhibit zero/few-shot learning capability.

Tabs. 1, 2, 3, and 4 show the superiority of TailedCore across all types of long-tail distributions. As anticipated, all baseline models encounter the tail-versus-noise trade-off, resulting in suboptimal performance either on tail classes with limited samples or on head classes with noisy training samples. These trends become more pronounced in highly unbalanced datasets, where the few-shot classes contain only 1 or 4 samples.

## 6.3. Ablation study and analysis

### 6.3.1. Tail (few-shot) sampling

We evaluate how accurately our proposed algorithm  $S_{tail}$  can sample tail (few-shot) class samples exclusively without including noise samples. First of all, we measure the ratio of mis-included tail class samples TailSampler which is the ratio of not sampled tail class samples to total tail class samples. A single mis-inclusion can vitally damage anomaly detection performance especially in case of the one-shot setting  $K=1$  as the model wouldn't be able to utilize the only sample of the class.

Additionally, the exclusive aspect of few-shot sampling is crucial; the few-shot sampled set shouldn't include anomaly and head class samples as possible. Included anomaly samples would harm the head class performance and head class sampled by few-shot sampler might cause class imbalance issue as well as the possibility of sampling anomaly from head classes. Hence, we measure the ratios of included anomalies and head-class samples along with AUROC which is measured on the outlier scores derived by a given algorithm and ground truth binary labels that indicate whether a given sample is tail class or head class. For TailSampler, the outlier score of each train sample is given as the negative of predicted class size. The included anomaly ratio is defined as the ratio of number of anomaly samples to the number of whole anomaly samples. The ratio of included head-class sample is defined as the number of included head class to the number of included tail class samples.

For extensive comparison, we compare TailSampler with widely used clustering algorithms [2, 13, 16], outlier detectors [1, 5, 24, 35], and density estimators [10, 30]. These methods provide clustering size and/or outlier score/prediction. In cases when only outlier scores are available, we apply the elbow method to induce deterministic prediction of whether a given sample is few-shot or not. We use default settings provided by Sklearn [31] and PyOD [42]. For all methods, we utilize the embedding vectors of WideResNet-50 [40]. Experiments with different architectures and details are provided in Supp. 10.

The combined ratio, which is the sum of missing tail samples and included anomaly samples, indicates that TailSampler is most effective for exclusively sampling few-shot class samples. The most critical factor is to not miss few-shot class samples. On this ratio, TailSampler achieves the lowest rate except for the algorithms that trivially sample all few-shot class samples. We find that most of the classical algorithms for clustering and outlier detection are not usable as they either miss too many few-shot samples or includes a large number of anomaly samples. We note that the class size predictor in our few-shot samples resembles DBSCAN and Affinity Propagation, respectively, in terms of using neighborhoods and voting mechanism. None of these methods however have instance-specific parameters that are adaptively adjusted across different samples.

### 6.3.2. Class size prediction

We measure how accurately our class size predictor estimates the class size of each sample by measuring the class size prediction error. It is measured for each train sample and weight-averaged through the whole train samples:

$$error = \frac{1}{|\mathcal{X}|} \sum_{i=1}^{|\mathcal{X}|} w_i |\kappa_i - |C_{k(i)}|| \quad (12)$$

where  $k(i)$  indicates the index such that  $x_i \in C_{k(i)}$  (namely,  $C_{k(i)}$  is the class of  $x_i$ ). The weight is chosen by  $w_i = 1/|C_{k(i)}|$ , which is to give more penalty on the few-shot class samples more than the head class samples.

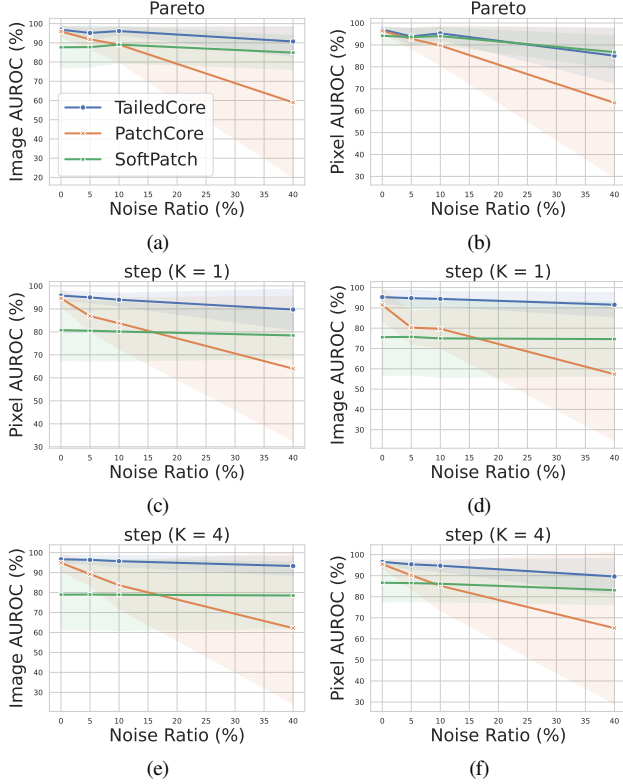


Figure 4. The plot of noise ratio of train data versus anomaly classification (image-level AUROC) and segmentation (pixel-level AUROC) performance. In all cases, TailedCore outperforms both PatchCore and SoftPatch. Dataset is MVTecAD with different tail distributions. The lines show the mean and shades show standard deviation of multiple runs of experiments with different seeds.

Tab. 5 shows that most of the clustering based methods vastly fail on this task, and the estimation by our class size predictor outperforms them significantly. We believe this is partly due to suboptimal hyperparameters used for these algorithms. Due to their classical nature, however, it is unclear how to tune them without the availability of the validation set in our scenario.

### 6.3.3. Abalation study on class size predictor

We conduct detailed ablation study on the proposed class size prediction in Eq. (7). One hyperparameter used in TailSampler  $S_{tail}$  of TailedCore is the percentiling proportion  $p$  to consider only the majority of nearby samples within the neighborhood. We analyze how this parameter can affect the anomaly detection performance of TailedCore. Tab. 5 indicates that using the nearest 85% embeddings (*i.e.*,  $p = 85\%$ ) is the best choice. Using all of the available neighborhoods (*i.e.*,  $p = 100\%$ ) would include different class samples in the neighborhood, deteriorating class estimation. On the other hand, counting only with the half of samples (*i.e.*,  $p = 50\%$ ) in the neighborhood may lose too much information.

Another component in the class size predictor that requires analysis is its voting mechanism used in Eq. (7). One may

method	combined ratio of missing $C_t$ & included AS ( $\downarrow$ )	ratio of missing $C_t$ samples ( $\downarrow$ )	ratio of included AS ( $\downarrow$ )	ratio of included $C_h$ samples ( $\downarrow$ )	class size prediction error ( $\downarrow$ )	AUROC (%) ( $\uparrow$ )
LOF	89.58 $\pm$ 44.42	31.22 $\pm$ 34.88	58.36 $\pm$ 9.54	92.50 $\pm$ 7.87	1.00 $\pm$ 0.00	77.08 $\pm$ 25.86
IF	18.15 $\pm$ 15.91	7.25 $\pm$ 10.05	10.90 $\pm$ 5.86	69.02 $\pm$ 25.96	1.00 $\pm$ 0.00	98.96 $\pm$ 1.27
OCSVM	49.95 $\pm$ 9.76	0.78 $\pm$ 2.19	49.17 $\pm$ 7.57	95.06 $\pm$ 6.08	1.00 $\pm$ 0.00	96.72 $\pm$ 3.59
DBSCAN	35.36 $\pm$ 49.55	30.64 $\pm$ 44.38	<b>4.72 <math>\pm</math> 5.17</b>	<b>25.57 <math>\pm</math> 27.24</b>	0.35 $\pm$ 0.44	84.46 $\pm$ 22.18
KMeans	94.67 $\pm$ 44.95	47.22 $\pm$ 35.49	47.45 $\pm$ 9.46	97.40 $\pm$ 3.99	4.27 $\pm$ 1.40	52.58 $\pm$ 20.46
GMM	100.00 $\pm$ 0.00	<b>0.00 <math>\pm</math> 0.00</b>	100.00 $\pm$ 0.00	97.58 $\pm$ 2.86	12.50 $\pm$ 1.53	50.00 $\pm$ 0.00
KDE	29.81 $\pm$ 24.15	2.14 $\pm$ 6.26	27.67 $\pm$ 17.89	87.47 $\pm$ 10.35	1.00 $\pm$ 0.00	97.98 $\pm$ 2.19
AP	27.81 $\pm$ 32.59	12.98 $\pm$ 24.01	14.83 $\pm$ 8.58	77.45 $\pm$ 28.21	0.70 $\pm$ 0.22	93.03 $\pm$ 12.8
KNN	43.49 $\pm$ 31.25	14.46 $\pm$ 24.07	29.03 $\pm$ 7.18	83.12 $\pm$ 18.28	1.00 $\pm$ 0.00	94.51 $\pm$ 10.34
TailSampler (ours)	<b>5.10 <math>\pm</math> 5.07</b>	<b>0.12 <math>\pm</math> 0.65</b>	<b>4.98 <math>\pm</math> 4.42</b>	<b>31.96 <math>\pm</math> 26.05</b>	<b>0.19 <math>\pm</math> 0.05</b>	<b>99.78 <math>\pm</math> 0.30</b>
Ablation on the class size predictor:						
TailSampler ( $p=100$ )	8.48 $\pm$ 10.88	1.07 $\pm$ 5.23	7.41 $\pm$ 5.65	42.64 $\pm$ 31.86	0.51 $\pm$ 0.04	99.76 $\pm$ 0.29
TailSampler ( $p=50$ )	7.37 $\pm$ 10.35	3.46 $\pm$ 6.80	3.91 $\pm$ 3.55	27.90 $\pm$ 27.36	0.12 $\pm$ 0.12	97.36 $\pm$ 5.65
TailSampler (top-1)	14.47 $\pm$ 7.99	0.20 $\pm$ 0.78	14.27 $\pm$ 7.21	73.97 $\pm$ 27.95	0.85 $\pm$ 0.04	99.73 $\pm$ 0.34
TailSampler (average)	6.48 $\pm$ 10.11	3.65 $\pm$ 6.99	2.83 $\pm$ 3.12	22.16 $\pm$ 24.33	0.21 $\pm$ 0.07	99.35 $\pm$ 2.59

Table 5. The analysis on the proposed TailSampler and its class size predictor. The results are obtained from both MVTecAD and VisA and averaged across all different types of tail distributions and random seeds (overall 60 different seeds) with the standard deviations indicated on the right. The best and second best performances are marked in **bold** and underline, respectively. Acronyms/notations: IF: Isolation Forest; AP: Affinity Propagation; AS: anomaly samples;  $C_h$  /  $C_t$ : head / tail classes.

task	tail type	Pareto			step ( $K=4$ )			step ( $K=1$ )		
		$C_t$	$C_h$	all	$C_t$	$C_h$	all	$C_t$	$C_h$	all
anomaly classification	w/o FSA	82.22	91.22	88.13	60.01	<b>97.60</b>	75.05	66.28	<b>97.85</b>	78.94
	w/ FSA	<b>96.55</b>	<b>95.24</b>	<b>96.13</b>	<b>93.54</b>	95.78	<b>94.44</b>	<b>95.82</b>	95.34	<b>95.72</b>
anomaly segmentation	w/o FSA	91.60	94.94	93.54	68.99	<b>96.85</b>	80.14	79.39	<b>96.35</b>	86.23
	w/ FSA	<b>96.09</b>	<b>95.01</b>	<b>95.30</b>	<b>94.19</b>	93.71	<b>94.00</b>	<b>95.56</b>	93.20	<b>94.75</b>

Table 6. Abalation on the few-shot augmentation (FSA) of TailedCore. The image-level and pixel-level AUROC metrics (%) are reported on the MVTecAD dataset with different types of tail distributions for anomaly classification and segmentation tasks, respectively. Notations:  $C_h$  /  $C_t$ : head / tail classes.

simply use the nearest embedding’s neighborhood size instead, or the class size of  $x_i$  can be estimated by averaging instead of majority vote. As indicated in Tab. 5, both the nearest embedding (*i.e.*, top-1) and averaging are inferior to majority voting. The former is inevitably less robust, while averaging can still be sensitive to the outlier values of neighborhood size.

### 6.3.4. Ablation on TailedCore

We analyze the impact of augmentation of few-shot class patches by TailSampler in TailedCore. The TailedCore model without these augmented patches has the same memory bank as SoftPatch, but its inference scoring is simpler. Tab. 6 indicates that without the augmented patches from TailSampler, the memory-based model cannot detect anomalies in few-shot class samples. This shows the necessity of tail class augmentation by TailSampler. We note some performance decrease in the head class detection performance since few-shot augmentation may involve few anomaly samples from the head classes that behave like few-shot class instances. This issue intrinsically is due to the weakness of encoder embedding, and the proposed TailSampler based on such embeddings can induce incorrectly sampled few-shot class instances, which we discuss in Sec. 7.

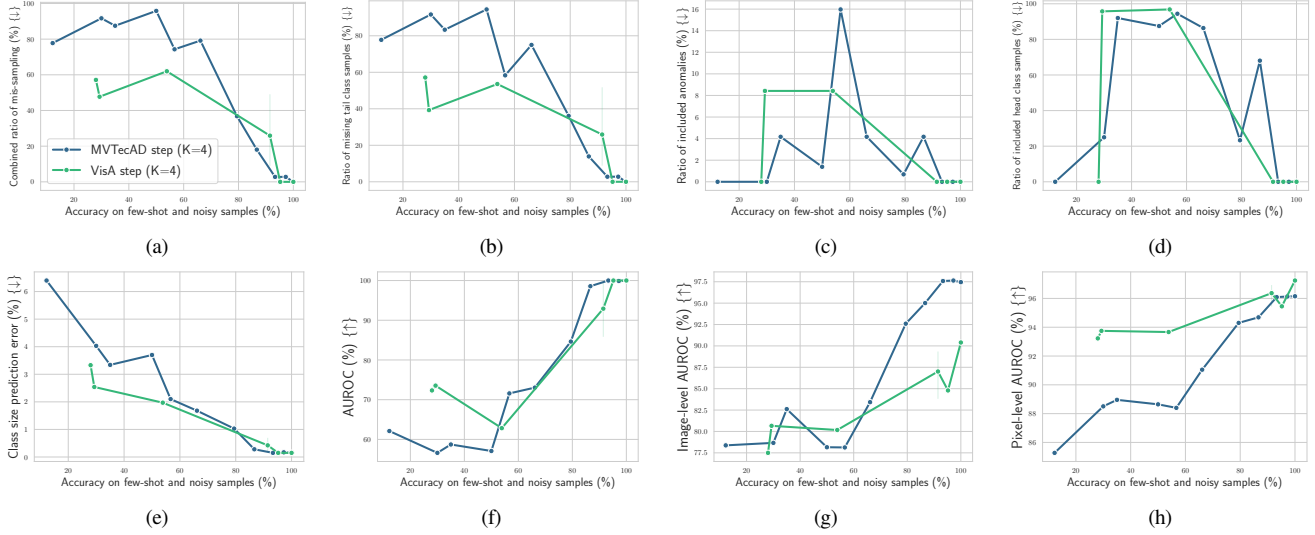


Figure 5. Classification accuracy between tail classes and noisy samples versus metrics relevant to class size prediction and few-shot sampling on MVTeCAD and VisA datasets with step-like tailed distribution and  $K = 4$ . The correlation is strong for (a) mis-sampling ratio, (b) ratio of missing few-shot samples, (e) class size prediction error, and (f) AUROC for predicting whether a sample is few-shot or not. We show that improving the discriminative aspect of network’s embeddings improves few-shot sampling and prediction error of class size predictor, which in turn improves (g) anomaly classification (image-level AUROC) and (h) anomaly segmentation (pixel-level AUROC) performance of TailedCore.

## 6.4. Noise ratio

We analyze the impact of noise ratio. Training data with a greater number of defect samples increase the chance of sampling noise instances for the few-shot sampler. We analyze this trend in Fig. 4, but the impact is not severe. In all cases, TailedCore either outperforms or is on par with both PatchCore and SoftPatch.

## 6.5. Relation between the discriminative aspect of embedding and TailedCore

TailedCore utilizes the embedding features  $e_i$  of training samples  $x_i$  for few-shot sampling and augmentation thereof. Therefore, TailedCore’s detection performance highly relies on the discriminative aspect of the embedding feature. Poorly designed embedding features will not discriminate different class samples in the embedding space. In this case, TailSampler cannot sample tail class samples and it may overly many head class samples. TailedCore’s memory bank obtained so will neither be class-balanced nor noise-free, making the inference of TailedCore poor.

In contrast, a highly discriminative embedding features will make TailedCore’s class size prediction accurate, thereby enabling it to sample the few-shot class samples exclusively without aggregating any anomaly and head class samples. In which case, TailedCore is both noise-free and class-balanced.

We test this behavior by Oracle labeling and training the embedding features with class labels. Fig. 5 shows that improving discriminative quality of embedding (measured by its classification accuracy on noise and tail class samples) accordingly improves the relevant metrics. Particularly, more discriminative embeddings make more accurate class size

prediction, improving TailSampler along with its anomaly classification and segmentation performances.

## 7. Limitations

The class size prediction in TailSampler can fail if the reflective-symmetric assumption on the inter-class and intra-class similarities breaks down. Particularly, this case can occur if the utilized encoder embedding is either poor or not aligned with the label space of classes. In some cases, the geometric aspect of defect samples highly resembles that of few-shot class instances in the embedding space. We believe that this is due to imperfect nature of encoder embedding, which is often more sensitive to spurious features than the notion of class and object-centric aspect [29]. This issue can be improved by enhancing the encoder network.

## 8. Conclusion

In this study, we introduced TailedCore, a novel memory-based anomaly detection model designed to effectively address the challenges of unsupervised anomaly detection in noisy and long-tail class distributions. By developing a unique class size predictor and a tailored memory bank, we successfully navigated the tail-versus-noise dilemma, enhancing the model’s performance in identifying anomalies within a contaminated and imbalanced dataset. Our comprehensive evaluations demonstrated TailedCore’s superior ability to mitigate noise contamination and class imbalance, underscoring its significance in advancing anomaly detection technologies. This work not only showcases a significant stride in handling real-world anomaly detection scenarios but also lays a robust foundation for future research in this critical field.

**Acknowledgments** Kuan-Chuan Peng was exclusively supported by Mitsubishi Electric Research Laboratories. This work was also supported by the Technology Innovation Program (20023364) funded by the Ministry of Trade, Industry & Energy (MOTIE, Korea), NSF grant 2038493, ONR grant N00014-21-1-2431, NIH grant R01CA240771 from NCI, and U.S. Department of Homeland Security grant 22STESE00001-03-02. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

## References

- [1] Fabrizio Angiulli and Clara Pizzuti. Fast outlier detection in high dimensional spaces. In *European conference on principles of data mining and knowledge discovery*, pages 15–27. Springer, 2002. 3, 6
- [2] David Arthur, Sergei Vassilvitskii, et al. k-means++: The advantages of careful seeding. In *Soda*, pages 1027–1035, 2007. 2, 6
- [3] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad—a comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9592–9600, 2019. 5
- [4] Paul Bergmann, Kilian Batzner, Michael Fauser, David Sattlegger, and Carsten Steger. Beyond dents and scratches: Logical constraints in unsupervised anomaly detection and localization. *International Journal of Computer Vision*, 130(4):947–969, 2022. 3
- [5] Markus M Breunig, Hans-Peter Kriegel, Raymond T Ng, and Jörg Sander. Lof: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, pages 93–104, 2000. 3, 6
- [6] Yunkang Cao, Xiaohao Xu, Jiangning Zhang, Yuqi Cheng, Xiaonan Huang, Guansong Pang, and Weiming Shen. A survey on visual anomaly detection: Challenge, approach, and prospect. *arXiv preprint arXiv:2401.16402*, 2024. 1
- [7] Yuanhong Chen, Yu Tian, Guansong Pang, and Gustavo Carneiro. Deep one-class classification via interpolated gaussian descriptor. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 383–392, 2022. 2
- [8] Zhuohang Chen, Jinglong Chen, Yong Feng, Shen Liu, Tianci Zhang, Kaiyu Zhang, and Wenrong Xiao. Imbalance fault diagnosis under long-tailed distribution: Challenges, solutions and prospects. *Knowledge-Based Systems*, 258:110008, 2022. 1
- [9] Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier. Padim: a patch distribution modeling framework for anomaly detection and localization. In *International Conference on Pattern Recognition*, pages 475–489. Springer, 2021. 2, 5, 6
- [10] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*, 39(1):1–22, 1977. 6
- [11] Jan Diers and Christian Pigorsch. A survey of methods for automated quality control based on images. *International Journal of Computer Vision*, pages 1–29, 2023. 1
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [13] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, pages 226–231, 1996. 2, 6
- [14] Chaowei Fang, Lechao Cheng, Huiyan Qi, and Dingwen Zhang. Combating noisy labels in long-tailed image classification. *arXiv preprint arXiv:2209.00273*, 2022. 2
- [15] Zheng Fang, Xiaoyang Wang, Haocheng Li, Jiejie Liu, Qiugui Hu, and Jimin Xiao. Fastrecon: Few-shot industrial anomaly detection via fast feature reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17481–17490, 2023. 2
- [16] Brendan J Frey and Delbert Dueck. Clustering by passing messages between data points. *science*, 315(5814):972–976, 2007. 2, 6
- [17] Zhaopeng Gu, Bingke Zhu, Guibo Zhu, Yingying Chen, Ming Tang, and Jinqiao Wang. Anomalygpt: Detecting industrial anomalies using large vision-language models. *arXiv preprint arXiv:2308.15366*, 2023. 2
- [18] Chaoqin Huang, Haoyan Guan, Aofan Jiang, Ya Zhang, Michael Spratling, and Yan-Feng Wang. Registration based few-shot anomaly detection. In *European Conference on Computer Vision*, pages 303–319. Springer, 2022. 2
- [19] Jongheon Jeong, Yang Zou, Taewan Kim, Dongqing Zhang, Avinash Ravichandran, and Onkar Dabeer. Winclip: Zero-/few-shot anomaly classification and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19606–19616, 2023. 2, 5, 6
- [20] Xi Jiang, Jianlin Liu, Jinbao Wang, Qiang Nie, Kai Wu, Yong Liu, Chengjie Wang, and Feng Zheng. Softpatch: Unsupervised anomaly detection with noisy data. *Advances in Neural Information Processing Systems*, 35:15433–15445, 2022. 1, 2, 3, 5, 6
- [21] Justin M Johnson and Taghi M Khoshgoftaar. Survey on deep learning with class imbalance. *Journal of Big Data*, 6(1):1–54, 2019. 1
- [22] Yujin Lee, Harin Lim, Seoyoon Jang, and Hyunsoo Yoon. Uniformly: Towards task-agnostic unified framework for visual anomaly detection. *Available at SSRN 4586132*. 2
- [23] Xurui Li, Ziming Huang, Feng Xue, and Yu Zhou. Musc: Zero-shot industrial anomaly classification and segmentation with mutual scoring of the unlabeled images. *arXiv preprint arXiv:2401.16753*, 2024. 2
- [24] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 eighth ieee international conference on data mining*, pages 413–422. IEEE, 2008. 3, 6
- [25] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2537–2546, 2019. 1, 6
- [26] Ruiying Lu, YuJie Wu, Long Tian, Dongsheng Wang, Bo Chen, Xiyang Liu, and Ruimin Hu. Hierarchical vector quantized transformer for multi-class unsupervised anomaly detection. *Advances in Neural Information Processing Systems*, 36, 2023. 2, 5, 6

- [27] Yang Lu, Yiliang Zhang, Bo Han, Yiu-ming Cheung, and Hanzi Wang. Label-noise learning with intrinsically long-tailed data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1369–1378, 2023. [2](#)
- [28] Declan McIntosh and Alexandra Branzan Albu. Inter-realization channels: Unsupervised anomaly detection beyond one-class classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6285–6295, 2023. [2](#)
- [29] Yifei Ming, Hang Yin, and Yixuan Li. On the impact of spurious correlation for out-of-distribution detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10051–10059, 2022. [8](#)
- [30] Emanuel Parzen. On estimation of a probability density function and mode. *The annals of mathematical statistics*, 33(3):1065–1076, 1962. [6](#)
- [31] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. [6](#)
- [32] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019. [2](#)
- [33] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [2](#)
- [34] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14318–14328, 2022. [2](#), [3](#), [5](#), [6](#)
- [35] Bernhard Schölkopf, Robert C Williamson, Alex Smola, John Shawe-Taylor, and John Platt. Support vector method for novelty detection. *Advances in neural information processing systems*, 12, 1999. [3](#), [6](#)
- [36] Robert L Thorndike. Who belongs in the family? *Psychometrika*, 18(4):267–276, 1953. [4](#), [5](#)
- [37] Tong Wei, Jiang-Xin Shi, Wei-Wei Tu, and Yu-Feng Li. Robust long-tailed learning under label noise. *arXiv preprint arXiv:2108.11569*, 2021. [2](#)
- [38] Guoyang Xie, Jinbao Wang, Jiaqi Liu, Feng Zheng, and Yaochu Jin. Pushing the limits of fewshot anomaly detection in industry vision: Graphcore. *arXiv preprint arXiv:2301.12082*, 2023. [2](#)
- [39] Zhiyuan You, Lei Cui, Yujun Shen, Kai Yang, Xin Lu, Yu Zheng, and Xinyi Le. A unified model for multi-class anomaly detection. *Advances in Neural Information Processing Systems*, 35:4571–4584, 2022. [1](#), [2](#)
- [40] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. [6](#)
- [41] Manyi Zhang, Xuyang Zhao, Jun Yao, Chun Yuan, and Weiran Huang. When noisy labels meet long tail dilemmas: A representation calibration method. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15890–15900, 2023. [2](#)
- [42] Yue Zhao, Zain Nasrullah, and Zheng Li. Pyod: A python toolbox for scalable outlier detection. *Journal of Machine Learning Research*, 20(96):1–7, 2019. [6](#)
- [43] Qihang Zhou, Guansong Pang, Yu Tian, Shibo He, and Jiming Chen. AnomalyCLIP: Object-agnostic prompt learning for zero-shot anomaly detection. In *International Conference on Learning Representations*, 2024. [2](#), [5](#), [6](#)
- [44] Yang Zou, Jongheon Jeong, Latha Pemula, Dongqing Zhang, and Onkar Dabeer. Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. *arXiv preprint arXiv:2207.14315*, 2022. [5](#)