

## Towards Accurate Unified Anomaly Segmentation

Wenxin Ma<sup>1,2</sup>

Qingsong Yao<sup>3</sup>  
 Zihang Jiang<sup>1,2\*</sup>

Xiang Zhang<sup>4</sup>  
 S. Kevin Zhou<sup>1,2,5,6\*</sup>

Zhelong Huang<sup>1,2</sup>

<sup>1</sup>School of Biomedical Engineering, Division of Life Sciences and Medicine, University of Science and Technology of China (USTC), Hefei Anhui, 230026, China

<sup>2</sup>Center for Medical Imaging, Robotics, Analytic Computing & Learning (MIRACLE), Suzhou Institute for Advance Research, USTC, Suzhou Jiangsu, 215123, China

<sup>3</sup>Stanford University, Palo Alto, California, 94025, United State

<sup>4</sup>School of Medicine, Shanghai University, Shanghai 200444, China

<sup>5</sup>Key Laboratory of Precision and Intelligent Chemistry, USTC, Hefei Anhui, 230026, China

<sup>6</sup>Key Laboratory of Intelligent Information Processing of Chinese Academy of Sciences (CAS), Institute of Computing Technology, CAS

{wxma, zhelonghuang}@mail.ustc.edu.cn jzh0103@ustc.edu.cn

rookiefcb@shu.edu.cn {qingsongyao98, s.kevin.zhou}@gmail.com

### Abstract

Unsupervised anomaly detection (UAD) from images strives to model normal data distributions, creating discriminative representations to distinguish and precisely localize anomalies. Despite recent advancements in the efficient and unified one-for-all scheme, challenges persist in accurately segmenting anomalies for further monitoring. Moreover, this problem is obscured by the widely-used AUROC metric under imbalanced UAD settings. This motivates us to emphasize the significance of precise segmentation of anomaly pixels using pAP and DSC as metrics. To address the unsolved segmentation task, we introduce the Unified Anomaly Segmentation (UniAS). UniAS presents a multi-level hybrid pipeline that progressively enhances normal information from coarse to fine, incorporating a novel multi-granularity gated CNN (MGG-CNN) into Transformer layers to explicitly aggregate local details from different granularities. UniAS achieves state-of-the-art anomaly segmentation performance, attaining 65.12/59.33 and 40.06/32.50 in pAP/DSC on the MVTEC-AD and VisA datasets, respectively, surpassing previous methods significantly. The codes are shared at <https://github.com/Mwxinnn/UniAS>.

### 1. Introduction

Unsupervised Anomaly Detection (UAD) in images involves modeling the normal data distribution [33, 42, 53]

\*Corresponding author.

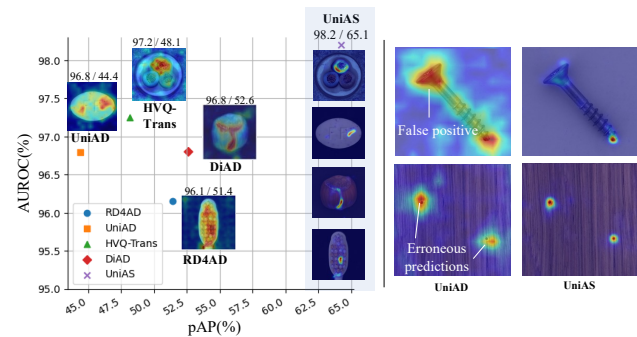


Figure 1. Prediction examples (Left) of the MVTEc dataset [1] and failed cases (Right) of the SOTA model UniAD [45]. UniAD has high AUROC but relatively poor segmentation performance, improved by our UniAS.

to detect rare and diverse unexpected signals within visual data, such as lesions in medical images [13, 41] and defects in industrial images [1]. While individual solutions per class consume significant computational resources, the newly proposed One-for-All scheme [15, 27, 45, 52] leverages a unified model to capture the complex joint distribution of various normal samples across diverse object classes without the need for fine-tuning. Moreover, extensive studies prove that, compared to traditional One-for-One schemes [29, 43], One-for-All UAD methods are faster, more memory-efficient, and show greater promise for generalization [44].

Most recent state-of-the-art (SOTA) UAD frameworks [15, 27, 44, 45] integrate transformer structures to learn to reconstruct compact representations of normal sam-

ples that are independent of abnormal samples. This makes it challenging to reconstruct defect-corrupted representations. Accordingly, the difference between reconstructed normal and abnormal features indicates anomalies. These frameworks achieve significant progress as assessed by the widely used metric, the area under the receiver operating characteristic curve (AUROC).

Despite significant advancements in image-level anomaly detection, recent observations [3, 28] indicate that even the top-performing models still struggle to precisely segment anomalies, often resulting in false-positive (FP) predictions (as shown in Fig. 1). Due to the class imbalance caused by numerous small anomalies in AD settings, AUROC tends to be a less effective indicator of model performance. High AUROC scores reported in the literature can possibly overlook predicted false-positive (FP) pixels, especially when true negative (TN) background pixels constitute the majority [28]. Nonetheless, precisely segmenting abnormal pixels is significant for quantifying the degree of anomaly and facilitating subsequent anomaly monitoring and modification [51].

In this paper, we delve deeper into the unresolved One-for-All anomaly segmentation challenge. We first demonstrate the inadequacy of solely relying on AUROC as the metric for evaluating anomaly segmentation and involve metrics such as pixel-wise Average Precision (pAP) and Dice Similarity Coefficient (DSC), which are widely used to evaluate segmentation performance. Accordingly, we find that the original cutting-edge feature-reconstruction-based One-for-All UAD methods show unsatisfactory performance on these segmentation metrics. We argue that the problem stems from their downsampling operations on the CNN-extracted features, which are then reconstructed at a coarse scale [27, 44, 45], compromising the intricate low-level information crucial for precise anomaly segmentation.

Motivated by this observation, we introduce our Unified Anomaly Segmentation (UniAS) approach, which considers multi-level features to enhance segmentation performance. UniAS establishes a multi-level hybrid pipeline to gradually reconstruct the extracted features in a level-by-level manner. Within each level, we propose a hybrid transformer-CNN module that leverages the transformer’s global receptive field to capture high-level features of normal data while simultaneously utilizing our Multi-Granularity Gated CNN (MGG-CNN) to recover intricate local details of normal samples. When anomalies are present, transformers can semantically detect anomalous pixels, with MGG-CNNs helping delineate the granularity-relevant boundaries of anomalies. By aggregating anomaly signals at different levels and granularities from coarse to fine, UniAS leverages both high and low-level anomaly maps to suppress semantic false-positive regions and manage to precisely segment anomalies eventually.

We conduct extensive experiments to demonstrate the effectiveness of our multi-level hybrid pipeline. Comparing UniAS with various cutting-edge models, our model shows significant improvement in precisely segmenting anomalies, as shown in Fig. 1 (Left). Our model achieves the highest pAP with excellent AUROC, and FP predictions are significantly suppressed by our design, as shown in Fig. 1 (Right).

In summary, we carefully analyze the ineffectiveness of AUROC under imbalanced situations and propose our UniAS to improve the performance of one-for-all anomaly segmentation. Our primary contributions are:

1. We explore the **limitation of AUROC** as the sole metric and discover the underlying FP problem in anomaly segmentation. Accordingly, we suggest greater emphasis on additional metrics including **pAP and DSC** for the evaluation of anomaly segmentation task.
2. We propose our new model, UniAS, which follows a **multi-level hybrid pipeline**, leveraging the benefits of the transformer’s global properties and the CNN’s locality to foster a holistic understanding of features across multiple levels and granularities, facilitating precise anomaly segmentation.
3. Extensive experimental results on widely used MVTEC-AD and VisA datasets demonstrate the effectiveness of UniAS, setting new SOTA to 65.12/59.33 and 40.06/32.50 in pAP and DSC, respectively, surpassing previous models by a large margin.

## 2. Related Work

### 2.1. Unsupervised Anomaly Detection

Current methods can primarily be divided into the following three categories: 1) **Augmentation-based methods** [22, 25, 47–49] typically involve the insertion of artificially generated anomalous patterns into normal signals, transforming the task into a supervised task. High-quality synthetic anomalies must be ensured. 2) **Boundary-based discriminative methods** compactly map normal features into a higher-dimensional space that separates them from abnormal components. To enhance discriminative capabilities, specialized designs [10, 11, 14, 19, 20, 25, 29, 30, 37, 38, 51] need to be developed, which can be computationally expensive. 3) **Reconstruction-based methods** train a model on normal samples exclusively [2, 7, 15, 43, 44], which is reconstructed in either feature or image space, with the assumption that the reconstruction error is notably higher for anomalous inputs. This approach is suitable for One-for-All settings [15, 27, 44, 45], where the distribution of either normal data or anomalies is highly complex, making it difficult to delineate a classification boundary or synthesize artificial anomaly samples to aid training. However, the reconstruction model may have “identical shortcut” issue where

abnormal inputs are also well reconstructed [27,45,46]. Additionally, normal images may contain details that are challenging to reconstruct, resulting in significant false positive predictions. The majority of current methods utilize downsampling techniques to highlight differences in high-dimensional semantics, neglecting lower-level features that are crucial for anomaly segmentation.

## 2.2. Transformers in Pixel-Level Tasks

The Transformer architecture has proven effective in various pixel-level tasks, including object detection and segmentation [5, 8, 9, 18]. By reformulating the tasks as a target-querying problem, they utilize learnable queries to encapsulate semantic information within features extracted by CNNs, ultimately used to match targets as bounding boxes [5] or segmentation masks [8, 9]. Specifically, DETR and MaskFormer process features at the highest level, while Mask2Former [8] sequentially processes features at each scale, achieving widespread success in segmentation [21].

For reconstruction-based One-for-All UAD, recent works [27, 45] have endowed the queries with a new role as a memory matrix to help reconstruct normal samples. UniAD [45] compresses the features to the same resolution and uses a MaskFormer-like structure. However, this approach fails to capitalize on the advantages of multi-level processing as Mask2Former. Our UniAS maximizes the benefits of a Mask2Former-like structure, with one-for-all AD-targeted modifications and achieving SOTA segmentation performance in UAD.

## 3. Our Approach

### 3.1. Overview

#### 3.1.1 Problem Formulation

One-for-All anomaly segmentation is defined as follows: Consider a set of anomaly-free data  $X^{train} = \{x_1, x_2, \dots, x_n\}$  across various categories, and a test set  $X^{test} = \{x'_1, x'_2, \dots, x'_n\}$  comprising both normal and abnormal samples from all the above categories. The general objective is to train a unified model capable of accurately identifying various anomalous pixels in anomaly samples across different categories within the test set. For reconstruction-based UAD, the training process takes anomaly-free data  $X^{train}$  as input, learning to replicate the distribution of the data or its features, while during inference, the model outputs pixel-wise reconstruction error as anomaly map  $A$ . This map is taken as the segmentation prediction accordingly.

#### 3.1.2 Observation: The limitation of AUROC

The Area Under the Receiver Operating Characteristic Curve (AUROC) between the anomaly map  $A$  and the seg-




Figure 2. (Left) An example of MVTEC-AD with a red line delimiting anomalous GT [1], along with toy anomaly maps, segmentation predictions, and the corresponding metrics. Bad predictions can have high AUROC numbers. (Right) A real prediction example and corresponding metrics, showing the limitation of AUROC.

mentation mask has been the most widely used metric for AD in pixel level [7, 10, 11, 14, 15, 22, 23, 25, 27, 29, 40, 44–46, 49, 51], which is calculated as Eq. (1). Here, TP stands for correctly predicted anomaly pixels, while TN stands for correctly predicted normal pixels. FP and FN are false-positive and false-negative pixels, respectively.

$$AUROC = \int_0^1 TPR d(FPR), \quad (1)$$

$$TPR = \frac{TP}{TP + FN}, \quad FPR = \frac{FP}{FP + TN}.$$

The performances reported in the literature are converging towards 100% regarding pixel-level AUROC, giving the impression that the pixel-level AD task has been solved [3]. However, anomaly samples typically appear rarely with few anomalous pixels [3, 28] (as shown in Fig. 2 (Left)). The imbalance problem can negatively impact the AUROC’s effectiveness, leading to failures in extreme scenarios [31]. As shown in Eq. (1), when detecting small anomaly signals, TN tends to be very large, resulting in a low FPR regardless of FP predictions. At the same time, the upsampling operation in recent works [27, 45] ensures a high TPR with erroneous predictions (high FP). These two factors contribute to the inflated AUROC numbers. We generate toy predictions as an example to demonstrate this serious problem of AUROC in Fig. 2 (Left). AUROC tends to be very high even with unsatisfactory segmentation results.

For real situations, we calculate the percentage of anomalous pixels in MVTEC and VisA, denoted as Anomaly Rate (AR), to reference the degree of imbalance in the datasets. A lower AR indicates that anomalous samples mainly consist of small regions, which can possibly lead to the discussed problem. The average AR in both datasets is relatively low (13.17% and 3.00%, respectively). As shown in Fig. 2 (Right). UniAD’s prediction has a high AUROC value with noticeable FP, while other metrics focusing on TP, such as pAP and DSC, can better reflect the precision of the result by eliminating the TP-TN imbalance problem.

In summarize, our observation indicates that high performance in the AUROC metric alone does not necessarily correlate with effective anomaly segmentation.

### 3.1.3 Motivation and Overall Design

**Motivation:** As demonstrated, pixel-level segmentation remains an unsolved task. We attribute this to downsampling in recent methods and the consequent loss of low-level features, which are crucial for capturing intricate details in anomaly segmentation. This motivates us to propose UniAS, aiming to model and aggregate precise multi-scale features to enhance pixel-level segmentation performance.

**The Overall Architecture** of our UniAS is depicted in Fig. 3. Initially, multi-scale features are extracted by a CNN backbone (see Sec. 3.2) which serves as the inputs and targets of reconstruction. To preserve information from different levels, we propose a **multi-level pipeline** to reconstruct features hierarchically. To precisely model the distribution of normal samples, we use a **hybrid architecture** combining Transformer and Multi-Granularity Gated CNN (MGG-CNN) together, leveraging their global and local benefits. In particular, MGG-CNN is designed to amalgamate details from diverse granularities. The whole process is aided by a Sample-Aware Reweighted (SAR) Query to enhance One-for-All ability (see Sec. 3.4). Finally, the anomaly maps (Sec. 3.3.2) are produced by comparing reconstructed features with extracted ones from various levels, then pixelwisely multiplied together to form the final prediction.

## 3.2. Feature Extraction and Filtering

For a normal input image  $x_i \in \mathbb{R}^{H \times W \times 3}$ , we use an ImageNet-pretrained CNN with  $N$  pyramidal structures  $\{\Phi_1, \Phi_2, \dots, \Phi_N\}$  to extract  $K$  multi-scale features, denoted by  $\mathcal{Z}^{ori} = \{\mathbf{z}_1^o, \dots, \mathbf{z}_K^o | \mathbf{z}_i \in \mathbb{R}^{H_i \times W_i \times C_i}\}$ , where  $H_i$ ,  $W_i$ , and  $C_i$  represent the height, width, and channel size of the  $i$ -th feature map. Additionally, as demonstrated to be effective in [25, 29], we apply a Gaussian filter  $W_G$  to aggregate the neighbor information and concatenate the residual to preserve details:

$$\begin{aligned} \mathcal{Z}^{ori} &= \{\mathbf{z}_1^o, \dots, \mathbf{z}_K^o | \mathbf{z}_i^o = f_{\Phi_i}(x) \odot W_G\}, \\ \mathcal{Z}^{cat} &= \{\mathbf{z}_1^c, \dots, \mathbf{z}_K^c | \mathbf{z}_i^c = cat(\mathbf{z}_i^o, \mathbf{z}_i^o - f_{\Phi_i}(x))\}. \end{aligned} \quad (2)$$

## 3.3. Multi-level Hybrid Reconstruction Pipeline

After extracting  $K$  levels of feature, UniAS learns to model and reconstruct the features at each level. Each reconstruction block consists of a transformer layer for **global** modeling and a MGG-CNN for **local** refinement. The transformer layer introduces normal signals to MGG-CNN, aided by a SAR query which acts as a memory matrix at the beginning (see details in Sec. 3.4).

### 3.3.1 Normal Data Modeling

**Patchify process:** Given the varying resolutions of the extracted features and computational constraints, we ensure

the number of patches at each level is consistent. This is achieved by utilizing various patch embedding layers  $\theta_1, \dots, \theta_k$  to divide the features  $\mathcal{Z}^{input}$  into patches with different sizes. Specifically, for the  $i$ -th feature  $\mathbf{z}_i \in \mathcal{Z}^{input}$  with resolution  $H_i \times W_i$ , the patch size is determined as  $H_i/H_K \times W_i/W_K$ , where  $H_K$  and  $W_K$  are the height and width of the feature at the highest  $K$ -th level, resulting in  $N = H_K \times W_K$  patches. These  $N$  embeddings are denoted as:

$$\mathbf{h}_i = f_{\theta_i}(\mathbf{z}_i), \quad \mathbf{h}_i \in \mathbb{R}^{N \times C'}, \quad i = 1, 2, \dots, K \quad (3)$$

**Global Modeling:** Transformer’s global receptive field allows it to capture long-range dependencies and context across the entire input, enabling more comprehensive and accurate modeling of complex patterns and relationships [5, 8, 9, 36]. To take advantage of this, we use transformer layers similar to Mask2Former [8], represented as:

$$\begin{aligned} \mathbf{d}_{CA} &= \text{MHCA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) + \mathbf{Q}, \\ \mathbf{d}_{SA} &= \text{MHSA}(\mathbf{d}_{CA}) + \mathbf{d}_{CA}, \\ \mathbf{d} &= \text{FFN}(\mathbf{d}_{SA}) + \mathbf{d}_{SA}, \end{aligned} \quad (4)$$

where  $\text{MHCA}(\cdot)$ ,  $\text{MHSA}(\cdot)$ , and  $\text{FFN}(\cdot)$  are the standard multi-head cross attention, multi-head self-attention and feed-forward network in a vanilla Transformer [12]. It takes the SAR query  $\mathbf{q}_0$  as the initial  $\mathbf{Q}$ , which is then updated by the output of  $(i-1)$ -th layer  $\mathbf{q}_i$ , and it takes feature patches  $\mathbf{h}_{K-i}$  as  $\mathbf{K}$  and  $\mathbf{V}$  for calculating attention, from high level to low level.

**Local Modeling:** The CNN’s local receptive field and inherent inductive bias excel at capturing fine-grained details and local patterns, making it highly effective for local feature extraction. This capability complements the transformer’s global modeling strength, creating a synergistic effect where the CNN handles local refinements while the transformer captures broader context [4, 6, 24, 35].

Considering the diverse granularities of anomalies across the dataset, we propose Multi-Granularity Gated CNN (MGG-CNN), by employing multiple branches with different receptive fields for better multi-granularity modeling. As depicted in Fig. 3 (left), these branches consist of one  $\text{Conv}1 \times 1$  branch accompanied by 3  $\text{Conv}3 \times 3$  branches, where the  $3 \times 3$  convolutional block is stacked for 1, 2 and 3 times, respectively. This design creates an MGG-CNN with a dynamic combination of receptive fields spanning 1, 3, 5, and 7. Furthermore, a GELU layer [16] acts as a gate function to activate the feature of each branch to encourage distinct information at each scale and introduce more non-linearity. Each branch’s output is added together. Our MGG-CNN is expressed as:

$$\mathbf{q}_i = \text{Gate}(\text{Conv}1(\mathbf{d}_i)) + \sum_{j=1}^3 \text{Gate}(\text{Conv}3^{j-1}(\mathbf{d}_i)), \quad (5)$$



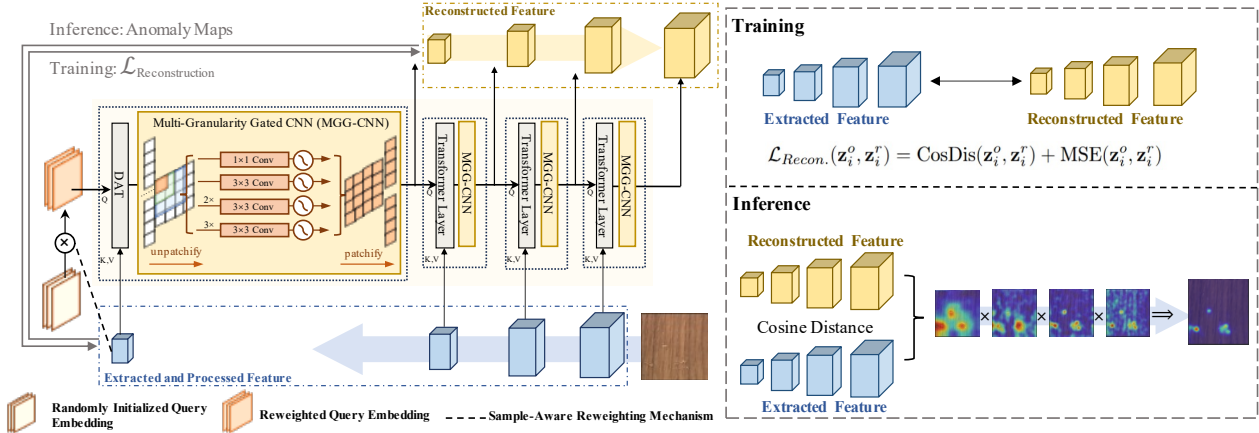


Figure 3. **(Left) The overview of UniAS.** After extracting features, our multi-level hybrid decoder, composed of Transformer and MGG-CNN hybrid blocks, hierarchically reconstructs normal features from coarse to fine. Transformers and MGG-CNNs play complementary roles in global and local modeling. The SAR Query is incorporated to facilitate One-for-All segmentation. **(Right) Illustration of the training loss and anomaly map generation during inference.** By multiplying anomaly maps together, information from separate levels is aggregated, facilitating accurate anomaly segmentation.

where each  $\mathbf{d}_i$  is the output of last transformer layer, and  $\mathbf{q}_i$  is the input for the next level transformer layer for multi-level modeling.

### 3.3.2 Reconstruction and Anomaly Maps

Following each MGG-CNN, an unpatchify operation and a transposed convolution layer are applied, resulting in the final reconstructed features  $\mathcal{Z}^{rec} = \{\mathbf{z}_1^r, \mathbf{z}_2^r, \dots, \mathbf{z}_K^r\}$  with the same dimensions as extracted feature  $\mathbf{z}_i^{ori}$ . As shown in Fig. 3, during training, the loss is calculated as the sum of Cosine Distance (CosDis) and Mean Squared Error (MSE), following [32]. During inference, we employ cosine distance to measure the dissimilarity between the reconstructed features and extracted features at each level, resulting in  $K$  anomaly maps. These maps are then upsampled to the same size  $H \times W$  as input images. The final prediction  $A$  is generated by multiplying these maps together, as Eq. (6). Based on these designs, UniAS can accurately segment anomalies in a coarse-to-fine fashion.

$$A = \prod_{i=1}^K \text{upsample}\left(1 - \frac{\langle \mathbf{z}_i^r, \mathbf{z}_i^o \rangle}{\|\mathbf{z}_i^r\| \cdot \|\mathbf{z}_i^o\|}\right). \quad (6)$$

### 3.4. Sample-Aware Reweighting Mechanism For One-for-all Segmentation

In addition to improving segmentation performance, we have also incorporated a Sample-Aware Reweighting (SAR) mechanism to initialize more sample-specific accurate prototypes as queries to boost One-for-All performance.

**Motivation:** Features extracted from various samples, particularly those from distinct categories, often demonstrate substantial disparities in high-level characteristics. Employing a shared learnable query to memorize patterns across

all categories can be suboptimal for this variability. We argue that incorporating sample-specific information into the query can help One-for-All reconstruction.

**Design:** Before training, the query  $\mathbf{q}_0^{ori} \in \mathbb{R}^{H_K \times W_K \times C'}$  is randomly initialized, whose shape is the same as the highest patch embedding  $\mathbf{h}_K$ . We utilize channel and spatial weights in CBAM [39] to reweight query  $\mathbf{q}_0^{ori}$  based on  $\mathbf{h}_K$  with the richest semantic. The multiplied  $\mathbf{q}_0 \in \mathbb{R}^{H_k \times W_k \times C'}$  is the query of the first transformer layer. This reweighted query incorporates sample-specific information, making it distinct and better suited to the variability across samples (see in Supplementary Material for details).

## 4. Experiments

### 4.1. Metrics

As discussed in Sec. 3.1.2, we analyze Anomaly Rate (AR) to show the degree of imbalance in the datasets. Accordingly, we adopt pAP and DSC apart from AUROC for anomaly segmentation evaluation. pAP evaluates the average precision across different recall levels, prioritizing successfully detecting positive pixels and keeping robust under class imbalance [15, 47, 51]. Meanwhile, DSC, a common metric in segmentation tasks, demonstrates robustness regardless of object size [17, 42]. In our implementation, we determine the segmentation threshold by maximizing the sum of precision and recall on the Precision-Recall (PR) curve for each category.

### 4.2. Datasets

#### 4.2.1 MVTec-AD [1]

MVTec-AD is a widely-used industrial anomaly detection dataset of real-world scenarios. It comprises 15 classes, to-

Dataset	AR(%)	(Originally) One-for-One			One-for-All					
		SimpleNet [25]	RD4AD [11]	DeSTSeg [51]	UniAD [45]	OmniAL [52]	HVQ-Trans [27]	DiAD [15]	ViTAD [50]	UniAS(Ours)
MVTec	13.17	96.78	96.15	94.41	96.79(96.80*)	<b>98.30*</b>	97.25(97.30*)	96.80*	97.70*	98.22
VisA	3.00	96.82	97.04	94.28	98.36	96.60*	98.41( <b>98.70*</b> )	96.00*	98.20*	97.80

Table 1. **The performance of recent SOTA models in pixel-AUROC(%)**. Results with \* are from the original paper, and other results are obtained from our reproduction. Best results are **bold**. Works are chronically arranged.

taling 5,354 high-resolution images spanning various categories, with one object in an image and a clean background. Each class consists of normal training samples, while the test set includes both normal and anomalous samples. For each anomalous sample in the test set, pixel-level ground-truth annotations are provided for segmentation evaluation.

#### 4.2.2 VisA [54]

VisA is a recent industrial anomaly detection dataset containing more challenging scenarios. This dataset comprises a total of 10,821 images, covering 12 object categories with 78 different types of anomalies, including scratches, dents, color spots, cracks, and structural defects. Images in the dataset exhibit complex structures, objects positioned in various locations, with tiny anomalous sizes and noisy backgrounds.

### 4.3. Implementation Details

All images in MVTEC-AD and VisA are resized to  $224 \times 224$ . We adopt a pretrained and fixed EfficientNet [34] to extract features from stage-1 to stage-4. Channel dimension  $C'$  for patch embedding layers is 256. AdamW [26] with weight decay 0.0001 is adopted for optimization. Our UniAS decoder is trained from scratch for 1000 epochs on a RTX 3090 GPU with batch size 64. The learning rate is initially set to  $1 \times 10^{-4}$ , and reduced by a factor of 0.1 every 400 epochs. For more details, please refer to the Supplementary Material.

### 4.4. Limitation of AUROC

We first present our results in AUROC in Tab. 1, compared with recent works. It can be observed that recent works show a noticeable saturation trend in reported AUROC across both datasets. However, it is important to note that the AUROC scores reported in the literature may be inflated and misleading for evaluating segmentation performance due to the imbalanced distribution between normal and abnormal pixels, as indicated in the AR column. Specifically, in the VisA dataset, which has a notably low anomaly rate, HVQ-Trans shows the highest AUROC values. Despite this, a substantial segmentation performance gap exists, as shown in the segmentation visualization in Fig. 4. Consequently, we shift our focus to other metrics such as pAP and DSC in the following discussion. Detailed AUROC results are provided in the Supplementary Material.

### 4.5. Qualitative Results

We visualize the results from both datasets in Fig. 4, comparing with two recent One-for-All SOTA models, UniAD [45] and HVQ-Trans [27]. While they can successfully distinguish anomalies, they exhibit numerous false positive predictions (e.g., Screw, Candle, and Capsule). In contrast, our design demonstrates more accurate segmentation performance across various anomaly categories. Particularly, UniAS mitigates erroneously FP predictions, thereby achieving better overall segmentation precision and offering valuable assistance in anomaly monitoring and removal.

We attribute the improvement to our design of multi-level feature reconstruction, which successfully model features at multiple levels, leveraging their complementary roles to suppress false predictions. In particular, higher-level features concentrate on semantics, facilitating accurate overall anomaly localization, while lower-level features have rich textural information aiding in delineating the shape of the anomalous region. Multiplying them together maximizes the complementary benefits, ensuring a precise segmentation result. We also visualize the anomaly maps of each layer, as depicted in Fig. 5. We label four levels of EfficientNet-extracted features 1 to 4, from the lowest to the highest. As shown, while each layer’s anomaly map consistently highlights the anomalous areas, their respective roles differ, contributing to the eventual precise result.

### 4.6. Quantitative Results

We choose seven of the best SOTA models to compare. These models include a boundary-based approach with a feature memory bank incorporated, **PatchCore** [29]; an image reconstruction-based method with pseudo-anomalies, **DRAEM** [47]; an image reconstruction-based method with diffusion, **DiAD** [15]; and four feature reconstruction-based methods, namely, **DeSTSeg** [51] and **RD4AD** [11] with CNN structures, as well as **UniAD** [45] and **HVQ-Trans** [27] with Transformer structures. UniAD and HVQ-Trans are two closely related transformer-based One-for-All works, while the remaining methods are originally designed as one-for-one approaches. For a fair comparison, we adapt these methods within the One-for-All scheme with the special designs unchanged. We re-trained the models using their publicly available official code with the default implementation for measuring pAP and DSC.

As shown in Tab. 2 and Tab. 3, our UniAS outper-

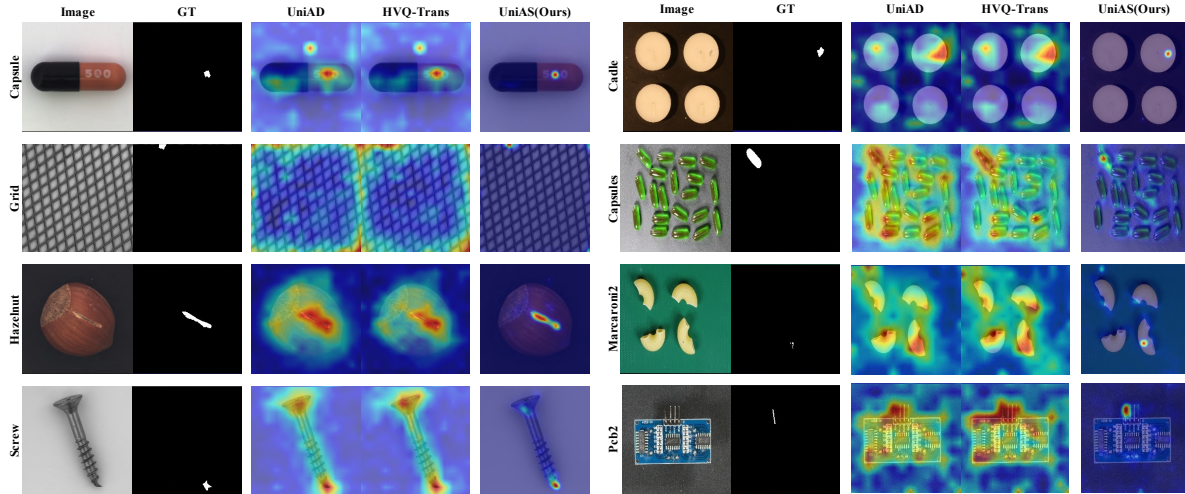


Figure 4. **Qualitative results on MVTec-AD(left) and VisA(right).** We visualize the anomaly maps and the generated masks. UniAS localizes anomaly precisely, exhibiting meaningful segmentation results.




Figure 5. (Left) Examples of input images and corresponding anomaly maps from feature level 4 to level 1. (Right) Aggregated anomaly maps and the anomaly ground truth masks. Anomaly maps play complementary roles, detecting anomalous pixels from coarse to fine.

Method	MVTec	VisA	Mean
PatchCore <a href="#">22'CVPR</a>	55.73	35.64	45.69
DRAEM <a href="#">21'ICCV</a>	49.86	34.49	42.18
DeSTSeg <a href="#">23'CVPR</a>	59.20	33.35	46.28
RD4AD <a href="#">22'CVPR</a>	51.41	38.23	44.82
UniAD <a href="#">22'NIPS</a>	44.42	34.32	39.37
HVQ-Trans <a href="#">23'NIPS</a>	48.15	33.17	40.66
DiAD <a href="#">24'AAAI</a>	52.60*	26.10*	39.35*
<b>UniAS (Ours) <a href="#">25'WACV</a></b>	<b>65.12</b>	<b>40.06</b>	<b>52.29</b>

Table 2. Quantitative Results on MVTec-AD and VisA in pAP(%). The best results are **bold**, and results with \* are from the original paper.

forms all the competitive models, setting new SOTA to 65.12/59.33 and 40.06/32.50 in the percentage of average pAP and DSC for two datasets, respectively. The performance has improved by a large margin, indicating the considerably improved segmentation ability of our UniAS. Specifically, for recent One-for-All SOTA models, namely UniAD and HVQ-Trans, feature downsampling in these

Method	MVTec	VisA	Mean
PatchCore <a href="#">22'CVPR</a>	50.74	28.15	39.44
DRAEM <a href="#">21'ICCV</a>	27.93	17.35	22.64
DeSTSeg <a href="#">23'CVPR</a>	43.61	23.48	33.55
RD4AD <a href="#">22'CVPR</a>	39.73	8.34	24.02
UniAD <a href="#">22'NIPS</a>	30.56	18.93	24.75
HVQ-Trans <a href="#">23'NIPS</a>	32.63	26.71	29.67
<b>UniAS (Ours) <a href="#">25'WACV</a></b>	<b>59.33</b>	<b>32.50</b>	<b>46.03</b>

Table 3. Quantitative Results on MVTec-AD and VisA in DSC(%). The best results are **bold**.

methods damages the ability of anomaly localization and segmentation. Without this operation, UniAS outperforms the latest SOTA model HVQ-Trans [27], by 16.97 $\uparrow$ /26.70 $\uparrow$  and 6.89 $\uparrow$ /5.79 $\uparrow$ . Moreover, the improvement compared to PatchCore [29] indicates the effectiveness of utilizing the SAR learnable query as a memory matrix compared to a computationally expensive memory bank [29].

		Components			Results	
Multi-level	SAR Q.	Hybrid Structure			Metric	
		Conv3	MG-CNN	MGG-CNN	pAP	DSC
-	-	-	-	-	40.14(24.98)	25.24(34.09)
✓	-	-	-	-	56.47( 8.65)	39.06(20.27)
✓	✓	-	-	-	59.84( 5.78)	48.35(10.98)
✓	✓	✓	-	-	61.14( 1.73)	49.27(10.06)
✓	✓	-	✓	-	63.39(1.13)	51.56( 7.77)
✓	✓	-	-	✓	<b>65.12</b>	<b>59.33</b>

Table 4. **Structural Component Study** The performance gap from the default setting is shown in red. The decoder of the baseline in the first line is a vanilla Transformer featuring self-attention layers exclusively. Cross-attention is employed to build a multi-level pipeline in the second line, with the deepest level still being formed by self-attention.

#Levels	Levels	Metric	
		pAP	DSC
1	{4}	44.44(20.68)	36.88(22.45)
	{3}	53.49(11.63)	38.22(21.11)
	{2}	56.52( 8.60)	37.47(21.86)
	{1}	40.18(24.94)	21.11(38.22)
2	{4,3}	54.20(10.92)	45.55(13.78)
	{4,1}	58.63( 6.49)	45.63(13.70)
3	{4,3,2}	61.57( 3.55)	55.03( 4.30)
4	{4,3,2,1}	<b>65.12</b>	<b>59.33</b>

Table 5. **Feature Combination Study** The performance gap from the default setting is in (red). Feature levels are labeled 1,2,3,4 from lowest to highest level accordingly. {·} means the levels included.

## 4.7. Ablation studies

### 4.7.1 Structural Component Study

In Tab. 4, we ablate each component of our structural designs on the MVTech-AD dataset:

1) **Multi-level:** Following UniAD [45], the baseline model downsamples and reconstructs multi-level features at the lowest resolutions in the highest level. Compared to our UniAS with multi-level design, the performance experiences a notable drop of 24.98 ↓ /34.09 ↓. However, upon simply incorporating the multi-level reconstruction design, the performance improves significantly to 56.47/39.06 in terms of pAP and DSC. This improvement underscores the complementary roles and necessity of multi-level features.

2) **Hybrid Structure:** We insert one layer of Conv3, Multi-granularity CNN (MG-CNN) without gate, and our MGG-CNN into Transformer layers. The results in the last three rows of Tab. 4 show improvement against the third line, proving the importance of precisely modeling local details. Moreover, MG-CNN with gated branches of different receptive fields is more effective than a simple Conv3 layer, while the activation layer (“gate”) helps to aggregate multi-scale features, forming MGG-CNN.

3) **Query:** We prove that simply incorporating a randomly initialized query for all datasets is less optimal in the One-for-All scheme. After deploying the SAR mechanism, significant performance improvements can be observed in the second line in Tab. 4.

### 4.7.2 Feature Combination Study

We confirm the essentiality and efficacy of amalgamating various feature levels, as presented in Tab. 5. For the 4 levels of feature extracted by EfficientNet, we infer the trained model with various combinations of feature levels. It is observed that incorporating more levels leads to better performance. Although we discover that feature levels 2 and 3 contain more valuable information compared to levels 1 and 4, combining all feature levels leads to the best performance, which validates the effectiveness of our multi-level pipeline.

## 5. Conclusion and Discussion

Anomaly segmentation is critical for anomaly monitoring and removal but remains under-explored. Recent excellent AUROC results have obscured this issue in imbalanced AD settings. In this study, we carefully discuss the limitations of AUROC and introduce other metrics for anomaly segmentation, namely pAP and DSC. Moreover, we propose our model, UniAS, which advances anomaly segmentation within a class-unified framework. UniAS hierarchically reconstructs normal features through a multi-level hybrid pipeline, combining transformer and CNN layers, and leverages multi-granularity gated CNN to facilitate local multi-granular detailed filtering and reconstruction. Additionally, a sample-aware reweighting mechanism is integrated to enhance the robustness of One-for-All anomaly segmentation. Together, UniAS effectively discriminates defect-corrupted features at the pixel level, achieving new SOTA anomaly segmentation performance with significant improvements.

Nevertheless, the performance of the proposed method in complex scenarios (*e.g.* several classes in VisA) still offers potential for further enhancement. Challenges persist in accurately delineating the contours of anomalous regions, particularly in scenarios involving multiple objects or background noise. Segmenting anomalies accurately in intricate scenes remains an open challenge for future studies.

## Acknowledgement

This work is supported by Natural Science Foundation of China under Grant 62271465, Suzhou Basic Research Program under Grant SYG202338, Open Fund Project of Guangdong Academy of Medical Sciences, China (No. YKY-KF202206), and Jiangsu Province Science Foundation for Youths (NO. BK20240464).



## References

- [1] 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. 1, 3, 5
- [2] Paul Bergmann, Sindy Löwe, Michael Fauser, David Sattlegger, and Carsten Steger. Improving unsupervised defect segmentation by applying structural similarity to autoencoders. *arXiv preprint arXiv:1807.02011*, 2018. 2
- [3] Joao PC Bertoldo, Dick Ameln, Ashwin Vaidya, and Samet Akçay. Aupimo: Redefining visual anomaly detection benchmarks with high speed and low tolerance. *arXiv preprint arXiv:2401.01984*, 2024. 2, 3
- [4] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang. Swin-UNET: UNet-like pure transformer for medical image segmentation. In *European conference on computer vision*, pages 205–218. Springer, 2022. 4
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 3, 4
- [6] Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L Yuille, and Yuyin Zhou. TransUNet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306*, 2021. 4
- [7] Liyang Chen, Zhiyuan You, Nian Zhang, Juntong Xi, and Xinyi Le. Utrad: Anomaly detection and localization with u-transformer. *Neural Networks*, 147:53–62, 2022. 2, 3
- [8] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. 3, 4
- [9] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34:17864–17875, 2021. 3, 4
- [10] Thomas Defard, Aleksandr Setkov, Angélique 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, 3
- [11] Hanqiu Deng and Xingyu Li. Anomaly detection via reverse distillation from one-class embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9737–9746, 2022. 2, 3, 6
- [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. 4
- [13] Tharindu Fernando, Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. Deep learning for medical anomaly detection—a survey. *ACM Computing Surveys (CSUR)*, 54(7):1–37, 2021. 1
- [14] Denis Gudovskiy, Shun Ishizaka, and Kazuki Kozuka. Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 98–107, 2022. 2, 3
- [15] Haoyang He, Jiangning Zhang, Hongxu Chen, Xuhai Chen, Zhishan Li, Xu Chen, Yabiao Wang, Chengjie Wang, and Lei Xie. Diad: A diffusion-based framework for multi-class anomaly detection. *arXiv preprint arXiv:2312.06607*, 2023. 1, 2, 3, 5, 6
- [16] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 4
- [17] Qixin Hu, Yixiong Chen, Junfei Xiao, Shuwen Sun, Jieneng Chen, Alan L Yuille, and Zongwei Zhou. Label-free liver tumor segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7422–7432, 2023. 5
- [18] Jitesh Jain, Jiachen Li, Mang Tik Chiu, Ali Hassani, Nikita Orlov, and Humphrey Shi. Oneformer: One transformer to rule universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2989–2998, 2023. 3
- [19] Danilo Jimenez Rezende and Shakir Mohamed. Variational inferezing with normalizing flows. *arXiv e-prints*, pages arXiv–1505, 2015. 2
- [20] Daehyun Kim, Sungyong Baik, and Tae Hyun Kim. Sanflow: Semantic-aware normalizing flow for anomaly detection. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 2
- [21] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9404–9413, 2019. 3
- [22] Chun-Liang Li, Kihyuk Sohn, Jinsung Yoon, and Tomas Pfister. Cutpaste: Self-supervised learning for anomaly detection and localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9664–9674, 2021. 2, 3
- [23] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 3
- [24] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 4
- [25] Zhikang Liu, Yiming Zhou, Yuansheng Xu, and Zilei Wang. Simplenet: A simple network for image anomaly detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20402–20411, 2023. 2, 3, 4, 6

- [26] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. **6**
- [27] 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. *arXiv preprint arXiv:2310.14228*, 2023. **1, 2, 3, 6, 7**
- [28] Mehdi Rafiei, Toby P Breckon, and Alexandros Iosifidis. On pixel-level performance assessment in anomaly detection. *arXiv preprint arXiv:2310.16435*, 2023. **2, 3**
- [29] 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. **1, 2, 3, 4, 6, 7**
- [30] Marco Rudolph, Tom Wehrbein, Bodo Rosenhahn, and Bastian Wandt. Fully convolutional cross-scale-flows for image-based defect detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1088–1097, 2022. **2**
- [31] Takaya Saito and Marc Rehmsmeier. The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets. *PLoS one*, 10(3):e0118432, 2015. **3**
- [32] Mohammadreza Salehi, Niousha Sadjadi, Soroosh Baselizadeh, Mohammad H Rohban, and Hamid R Rabiee. Multiresolution knowledge distillation for anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14902–14912, 2021. **5**
- [33] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International conference on information processing in medical imaging*, pages 146–157. Springer, 2017. **1**
- [34] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. **6**
- [35] Fenghe Tang, Bingkun Nian, Jianrui Ding, Quan Quan, Jie Yang, Wei Liu, and S Kevin Zhou. Mobileutr: Revisiting the relationship between light-weight cnn and transformer for efficient medical image segmentation. *arXiv preprint arXiv:2312.01740*, 2023. **4**
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. **4**
- [37] Guodong Wang, Shumin Han, Errui Ding, and Di Huang. Student-teacher feature pyramid matching for anomaly detection. *arXiv preprint arXiv:2103.04257*, 2021. **2**
- [38] Meng Wang, Tian Lin, Lianyu Wang, Aidi Lin, Ke Zou, Xinxing Xu, Yi Zhou, Yuanyuan Peng, Qingquan Meng, Yiming Qian, et al. Uncertainty-inspired open set learning for retinal anomaly identification. *Nature Communications*, 14(1):6757, 2023. **2**
- [39] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. **5**
- [40] Tiange Xiang, Yixiao Zhang, Yongyi Lu, Alan L Yuille, Chaoyi Zhang, Weidong Cai, and Zongwei Zhou. Squid: Deep feature in-painting for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23890–23901, 2023. **3**
- [41] Qingsong Yao, Zecheng He, Yuexiang Li, Yi Lin, Kai Ma, Yefeng Zheng, and S Kevin Zhou. Adversarial medical image with hierarchical feature hiding. *IEEE Transactions on Medical Imaging*, 2023. **1**
- [42] Qingsong Yao, Li Xiao, Peihang Liu, and S Kevin Zhou. Label-free segmentation of covid-19 lesions in lung ct. *IEEE transactions on medical imaging*, 40(10):2808–2819, 2021. **1, 5**
- [43] Xincheng Yao, Ruoqi Li, Zefeng Qian, Yan Luo, and Chongyang Zhang. Focus the discrepancy: Intra-and inter-correlation learning for image anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6803–6813, 2023. **1, 2**
- [44] Xincheng Yao, Chongyang Zhang, Ruoqi Li, Jun Sun, and Zhenyu Liu. One-for-all: Proposal masked cross-class anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 4792–4800, 2023. **1, 2, 3**
- [45] 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, 3, 6, 8**
- [46] Zhiyuan You, Kai Yang, Wenhan Luo, Lei Cui, Yu Zheng, and Xinyi Le. Adtr: Anomaly detection transformer with feature reconstruction. In *International Conference on Neural Information Processing*, pages 298–310. Springer, 2022. **3**
- [47] Vitjan Zavrtnik, Matej Kristan, and Danijel Skočaj. Draema: a discriminatively trained reconstruction embedding for surface anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8330–8339, 2021. **2, 5, 6**
- [48] Gongjie Zhang, Kaiwen Cui, Tzu-Yi Hung, and Shijian Lu. Defect-gan: High-fidelity defect synthesis for automated defect inspection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2524–2534, 2021. **2**
- [49] Hui Zhang, Zuxuan Wu, Zheng Wang, Zhineng Chen, and Yu-Gang Jiang. Prototypical residual networks for anomaly detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16281–16291, 2023. **2, 3**
- [50] Jiangning Zhang, Xuhai Chen, Yabiao Wang, Chengjie Wang, Yong Liu, Xiangtai Li, Ming-Hsuan Yang, and Dacheng Tao. Exploring plain vit reconstruction for multi-class unsupervised anomaly detection. *arXiv preprint arXiv:2312.07495*, 2023. **6**

- [51] Xuan Zhang, Shiyu Li, Xi Li, Ping Huang, Jiulong Shan, and Ting Chen. Destseg: Segmentation guided denoising student-teacher for anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3914–3923, 2023. 2, 3, 5, 6
- [52] Ying Zhao. Omnia: A unified cnn framework for unsupervised anomaly localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3924–3933, 2023. 1, 6
- [53] Bo Zong, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Daeki Cho, and Haifeng Chen. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In *International conference on learning representations*, 2018. 1
- [54] Yang Zou, Jongheon Jeong, Latha Pemula, Dongqing Zhang, and Onkar Dabeer. Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. In *European Conference on Computer Vision*, pages 392–408. Springer, 2022. 6