

Boosting Fine-Grained Visual Anomaly Detection with Coarse-Knowledge-Aware Adversarial Learning

Qingqing Fang¹, Qinliang Su^{1,2*}, Wenxi Lv¹, Wenchao Xu³, Jianxing Yu^{4,5}

¹ School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou, China

² Guangdong Key Laboratory of Big Data Analysis and Processing, Guangzhou, China

³ Department of Computing, The Hong Kong Polytechnic University, Hong Kong SAR

⁴ School of Artificial Intelligence, Sun Yat-sen University, Guangdong, China

⁵ Pazhou Lab, Guangzhou, 510330, China

{fangqq3, lvwx8}@mail2.sysu.edu.cn, {suqliang, yujx26}@mail.sysu.edu.cn, wenchao.xu@polyu.edu.hk

Abstract

Many unsupervised visual anomaly detection methods train an auto-encoder to reconstruct normal samples and then leverage the reconstruction error map to detect and localize the anomalies. However, due to the powerful modeling and generalization ability of neural networks, some anomalies can also be well reconstructed, resulting in unsatisfactory detection and localization accuracy. In this paper, a small coarsely-labeled anomaly dataset is first collected. Then, a coarse-knowledge-aware adversarial learning method is developed to align the distribution of reconstructed features with that of normal features. The alignment can effectively suppress the auto-encoder's reconstruction ability on anomalies and thus improve the detection accuracy. Considering that anomalies often only occupy very small areas in anomalous images, a patch-level adversarial learning strategy is further developed. Although no patch-level anomalous information is available, we rigorously prove that by simply viewing any patch features from anomalous images as anomalies, the proposed knowledge-aware method can also align the distribution of reconstructed patch features with the normal ones. Experimental results on four medical datasets and two industrial datasets demonstrate the effectiveness of our method in improving the detection and localization performance.

anomalies, in this paper, we focus our attention on the fine-grained anomaly detection. In addition to detecting anomalies, some applications also require locating the anomalies, like automatically marking the defective areas (Roth et al. 2022a; Liu et al. 2023b). Thus, it is of great value to develop methods that can simultaneously detect and locate anomalies accurately.

Considering the extreme diverseness of anomalies, existing methods (Ruff et al. 2018; Akcay, Atapour-Abarghouei, and Breckon 2019; Tack et al. 2020; Roth et al. 2022b; Deng and Li 2022; Reiss and Hoshen 2023) mostly only assume the access to normal datasets. By using it to train a model that characterizes the normal patterns, anomalies can then be detected by assessing their conformity to it. The reconstruction-based method (Akcay, Atapour-Abarghouei, and Breckon 2019; Vasilev et al. 2020; Zavrtnik, Kristan, and Skočaj 2021b; Guo et al. 2023), as one of the most widely-used fine-grained anomaly detection methods, is a concrete example of this idea, which first trains an auto-encoder on normal images and then uses the reconstruction error map to detect and localize anomalies. However, due to the strong modeling and generalization ability of neural networks, it is found that auto-encoders sometimes can also reconstruct the anomalous areas well due to the leakage of visual information from other parts of images, causing the reconstruction error map unreliable for anomaly detection and localization (You et al. 2022).

Code — <https://github.com/Faustinaqq/CKAAD>

Extended version — <https://arxiv.org/abs/2412.12850>

Introduction

Visual anomaly detection (AD) aims to detect images that show significant deviation from normal ones. In practice, the deviation could happen at image-level or at the more fine-grained pixel-level, where anomalies only appear at certain local areas, while the remaining large region looks to be normal (Yang et al. 2022). Fine-grained anomalies happen a lot in real-world applications, such as the X-ray/CT images with nodules or other lesions in medical diagnosis (Fernando et al. 2021), images of products with defects like undesired scratches or small holes in industrial manufacturing *etc* (Liu et al. 2024). Because of the ubiquitousness of fine-grained anomalies and their intrinsic distinction from image-level

The reason why the reconstruction-based methods cannot lead to competitive performance can be largely attributed to the only usage of normal samples when training the auto-encoder, which makes the decision boundary between normal and anomalous samples blurred (Xia et al. 2022). To boost the detection and localization accuracy, we propose to introduce a small dataset comprised of anomalous images. We argue that it is feasible to collect a small number of anomalies in many real-world applications (Ruff et al. 2019; Pang et al. 2021a; Ding, Pang, and Shen 2022). As exemplified in medical diagnosis and industrial manufacturing, there often exist some X-Ray images that have been confirmed to be problematic by veteran doctors, or images that are taken for products with defects. But compared with datasets in classification tasks, the anomaly datasets possess two unique characteristics: 1) *Coarseness*: anomalies only

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

occupy a small area of the image, but only the entire image is known to be anomalous; 2) *Incompleteness*: the dataset only covers a tiny subset of all possible anomaly types due to the diverseness of anomalies. Thus, the dataset can be viewed as a kind of coarse anomaly knowledge, and our goal is to make use of it for better anomaly detection and localization. We note that some recent works have also proposed to collect an additional anomaly dataset and then use it to help detect anomalies like the Deep SAD (Ruff et al. 2019), PANDA-OE (Reiss et al. 2021), DevNet (Pang et al. 2021a), DRA (Ding, Pang, and Shen 2022), AA-BiGAN (Tian, Su, and Yin 2022), PReNet (Pang et al. 2023), *etc.* However, these methods simply view an entire image in the anomaly dataset as anomalous, and rarely explicitly take the coarseness characteristics of anomalies into account. Thus, when they are applied to fine-grained anomaly detection problems, they may not make the best use of the knowledge hidden in the anomaly dataset, resulting in a performance that is not good enough. Moreover, all of these methods only consider how to improve anomaly detection accuracy, but rarely bear in mind the anomaly localization task that is also important in many scenarios.

In this work, we proposed to leverage the coarse anomaly knowledge to boost the anomaly detection and localization performance of feature-reconstruction-based methods. To this end, a Coarse-Knowledge-Aware adversarial learning Anomaly Detection (CKAAD) strategy is developed to suppress the auto-encoder’s reconstruction ability on anomalies by seeking to align the distribution of reconstructed features with that of normal features, where the features are extracted from a pre-trained ResNet in advance. Specifically, we first utilize the coarse anomalous knowledge to develop an energy-based discriminator, which is trained to assign low energy for normal features, but high energy for reconstructed features and anomalous features. The discriminator can effectively incentivize the auto-encoder to output normal features, while explicitly avoiding outputting anomalous features, thereby achieving the alignment with the distribution of normal features. Considering that anomalies often only occur in small areas in the anomalous images, a patch-level adversarial learning strategy is further developed. Because there are no patches that are known to be anomalous surely, we show that we can simply view all patches from an anomalous image as anomalies and then apply a similar adversarial learning strategy at image-level to guide the auto-encoder’s reconstruction. We rigorously prove that the proposed patch-level adversarial learning strategy can also lead to the alignment between the reconstructed features and normal features, even though the patch-based anomalous knowledge is not completely correct. Experimental results on real-world anomaly datasets including four medical datasets and two industrial datasets demonstrate that by leveraging a small number of coarsely labeled anomalies, our proposed coarse-knowledge-aware method can boost the detection and localization performance remarkably.

Related Work

Unsupervised methods Existing unsupervised methods only use normal samples. Among them, reconstruction-

based methods commonly use auto-encoder (Zavrtanik, Kristan, and Skočaj 2021b; Zhang, Wang, and Chen 2022), variational auto-encoder (Vasilev et al. 2020) or generative adversarial network (Schlegl et al. 2017; Akcay, Atapour-Abarghouei, and Breckon 2019) to reconstruct normal images, assuming that unseen anomalous images can not be reconstructed well, so the difference between the input and reconstructed images can be used to detect anomalies. Because these image-level methods hardly localize anomalous areas accurately, the feature reconstruction methods (Salehi et al. 2021; Deng and Li 2022; Tien et al. 2023; Guo et al. 2023) are further proposed. But whether reconstructing in the image or feature level, these reconstruction-based methods face the same problem that the anomalies are also restored well, decreasing the detection performance. Besides the reconstruction-based methods, one-class methods (Ruff et al. 2018; Yi and Yoon 2020; Reiss et al. 2021; Gui et al. 2022) map normal images or patches to a compact space, then determine anomalies based on the distance. Embedding-based methods (Tack et al. 2020; Defard et al. 2021; Cho, Seol, and Lee 2021; Reiss and Hoshen 2023) discriminate anomalies according to the learned embedding features. Other unsupervised methods (Reiss et al. 2021; Zavrtanik, Kristan, and Skočaj 2021a; Li et al. 2021; Liu et al. 2023b) utilize synthesized pseudo anomalies for directly predicting anomaly labels or masks, can result in poor detection performance on real anomalies.

Weakly-supervised methods In recent years, some weakly-supervised methods have been proposed to enhance anomaly detection performance by utilizing observed partial anomalies. Based on pre-trained models, PANDA-OE (Reiss et al. 2021) directly employs classification models, but directly using a discriminator to find a boundary can result in misclassifications of unseen anomalies. Instead of training a classifier, Deep SAD (Ruff et al. 2019) maps normal samples to a compact space, pushing anomalous samples away. DPLAN (Pang et al. 2021b) guides the model to find anomaly samples and assigns higher anomaly scores through reinforcement learning. GAN-based methods (Kimura et al. 2020; Tian, Su, and Yin 2022; Tian, Su, and Yu 2023; Lv et al. 2024) avoid generating anomalous samples and get the anomaly score based on the norm of latent feature or the reconstruction error. DevNet (Pang et al. 2021a) and DRA (Ding, Pang, and Shen 2022) use deviation loss or score heads to separate anomalous patches from normal ones. PReNet (Pang et al. 2023) detects anomalies based on pairwise anomaly scores learned by three kinds of instance pairs. However, these methods directly enlarging the distinction of anomaly scores between normal and anomalous samples, are prone to over-fit on these collected samples. Besides, most of them primarily use these incomplete anomalies for detection, rarely leveraging such weakly supervised information to achieve accurate localization performance.

Preliminaries on Feature Reconstruction-Based Anomaly Detection

Due to the diverseness of anomalies, current anomaly detection methods are generally established on the assumption of

access to a normal dataset $\mathcal{X}^+ = \{x_1^+, x_2^+, \dots, x_n^+\}$. By seeking to characterize the normal patterns in normal samples as accurately as possible, anomalies can be identified by checking whether they conform to the normal patterns. Regarding methods following this approach, the most widely used one is the reconstruction-based method, which basically trains an auto-encoder (AE) on normal samples and then uses the reconstruction error to detect anomalies.

For image data, to yield better performance, it is often suggested to reconstruct feature maps $\{F^\ell\}_{\ell=1}^L$ that are extracted from pre-trained ResNet (He et al. 2016), instead of reconstructing original images, with F^ℓ and L denoting the feature map from the ℓ -th block of ResNet and the total number of blocks. Since the feature maps from different blocks are of different sizes, when implementing the encoder, we first down-sample maps to the same size as the smallest one by a sequence of 3×3 convolutional operators and then concatenate them along the channel dimension. The concatenated maps are then further fed into a ResNet block to produce the final latent representation z , as shown in Figure 1. The whole procedure can be described as $z = \mathcal{E}_\theta(\{F^s\}_{s \in \mathcal{S}})$, where $\mathcal{S} \subset \{1, 2, \dots, L\}$ represents a subset of L ResNet blocks that are chosen to participate in reconstruction. Given the latent code z , residual deconvolution operations are then employed to up-sample it stage by stage to recover the feature maps of different sizes \tilde{F}^s for $s \in \mathcal{S}$, that is, $\{\tilde{F}^s\}_{s \in \mathcal{S}} = \mathcal{D}_\theta(z)$. By viewing the original feature maps $\{F^s\}_{s \in \mathcal{S}}$ as input and the recovered maps $\{\tilde{F}^s\}_{s \in \mathcal{S}}$ as output, the encoder $\mathcal{E}_\theta(\cdot)$ and decoder $\mathcal{D}_\theta(\cdot)$ can be jointly written as

$$\{\tilde{F}^s\}_{s \in \mathcal{S}} = G_\theta(\{F^s\}_{s \in \mathcal{S}}), \quad (1)$$

where $G_\theta(\cdot) = \mathcal{D}_\theta(\mathcal{E}_\theta(\cdot))$.

The auto-encoder $G_\theta(\cdot)$ is trained to reduce the difference between F^s and \tilde{F}^s for $s \in \mathcal{S}$ on all training normal samples by minimizing the loss

$$L_{rec}^+ = \frac{1}{|\mathcal{X}^+|} \sum_{x \in \mathcal{X}^+} \sum_{s \in \mathcal{S}} \left(1 - \cos(\text{flatten}(\tilde{F}^s), \text{flatten}(F^s))\right) \quad (2)$$

where $\text{flatten}(\cdot)$ and $\cos(\cdot, \cdot)$ mean flattening the tensor into one-dimensional vector and the cosine similarity. Since the loss is trained only on normal samples, the auto-encoder tends to reconstruct the input normal images well, but poorly for anomalous ones. That is why the error maps between F^s and \tilde{F}^s can be employed to detect and localize the anomalies. However, due to the generalization ability of neural networks, it is found that some anomalous images can also be well reconstructed, making the error map not always reliable in detecting and localizing the anomalies.

Boosting Reconstruction-Based Anomaly Detection via Coarse-Knowledge-Aware Adversarial Learning

To improve the detection and localization accuracy, in addition to the normal dataset \mathcal{X}^+ , we further assume the availability of a small anomaly dataset

$$\mathcal{X}^- = \{x_1^-, x_2^-, \dots, x_m^-\}. \quad (3)$$

As elaborated in the Introduction, due to the diverseness of anomalies, the collected anomalies in \mathcal{X}^- may only cover a very tiny number of all possible anomaly types. Moreover, due to the annotating cost, we assume we only know the images in \mathcal{X}^- are anomalous, but have no knowledge of the exact location of the anomalies. This paper aims to leverage this coarse anomalous knowledge to better detect and locate anomalies.

One of the most direct ways to leverage the dataset \mathcal{X}^- is to minimize the following loss

$$L_{rec} = L_{rec}^+ - \lambda L_{rec}^-, \quad (4)$$

where L_{rec}^- denotes the reconstruction error computed over samples in \mathcal{X}^- ; and $\lambda \in \mathbb{R}^+$ is a positive weighting factor. However, the loss L_{rec} only drives the auto-encoder to output feature maps \tilde{F}^s that look different from the input F^s for anomalous samples, but never restricting how the difference should look like. Because anomalies often only occupy a small area in anomalous images in many real-world applications, simply enlarging the reconstruction error of anomalous samples has the risk of turning normal areas into anomalies, deteriorating the detection and localization performance. To enable the auto-encoder to detect and localize the anomalies accurately, we argue that we should encourage it to possess the following two properties:

- P1) The auto-encoder should manage to transform an anomalous image to its normal counterpart, *e.g.*, transforming an anomalous image with a small hole into one that has the hole removed.
- P2) The reconstruction quality of normal samples should not be compromised too much by the introduced anomalies.

If the two appealing properties hold, the reconstruction error map between input and output features could be used to detect and localize the anomalies more reliably.

Aligning the Distributions of Reconstructed Features with Normal Features

To achieve the appealing properties above, we propose to view the auto-encoder as a generator $G_\theta(\cdot)$. For the simplicity of presentation, we only present how to process one feature map $F^s \in \{F^s\}_{s \in \mathcal{S}}$, and denote the chosen map F^s as F for conciseness, while the remaining feature maps can be processed similarly. Since we now have two datasets available, \mathcal{X}^+ and \mathcal{X}^- , we can use them to construct a distribution regarding input feature maps as

$$\mathbb{P}(F) = \alpha \mathbb{P}^+(F) + (1 - \alpha) \mathbb{P}^-(F), \quad (5)$$

where $\mathbb{P}^+(F)$ and $\mathbb{P}^-(F)$ represent feature map distributions of normal samples $x^+ \in \mathcal{X}^+$ and anomalous samples $x^- \in \mathcal{X}^-$, respectively; $\alpha \in [0, 1]$ is a parameter used to control the mixture ratio of the two distributions. For any input feature F drawn from the distribution $\mathbb{P}(F)$, the generator (auto-encoder) will output a feature map $G_\theta(F)$. The distribution of the generated feature map $G_\theta(F)$ can be represented as

$$\mathbb{P}_g(F) = \alpha \mathbb{P}_g^+(F) + (1 - \alpha) \mathbb{P}_g^-(F), \quad (6)$$

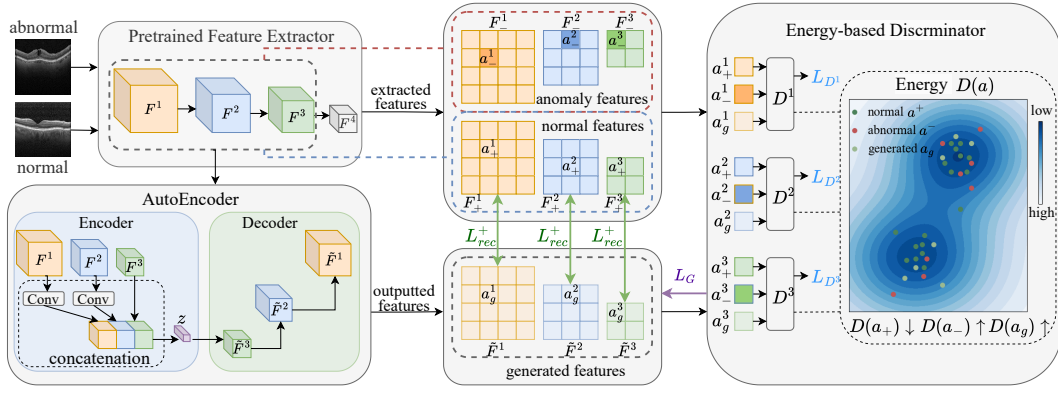


Figure 1: Framework of our proposed CKAAD. With the normal and anomalous feature maps extracted from pre-trained ResNet as well as the generated feature maps output from the auto-encoder, the energy-based discriminator D^ℓ is trained to distinguish between normal, anomaly, and generated patch features output from the ℓ -th layer of ResNet. The auto-encoder is then incentivized to always output normal features regardless of the input types.

where $\mathbb{P}_g^+(F)$ and $\mathbb{P}_g^-(F)$ denote the distributions of generated (reconstructed) feature map $G_\theta(F)$ when the input F is drawn from $\mathbb{P}^+(F)$ and $\mathbb{P}^-(F)$, respectively. If the distribution of $\mathbb{P}_g(F)$ is forced to align with the distribution of normal features $\mathbb{P}^+(F)$, it amounts to driving $\mathbb{P}_g^-(F)$ towards $\mathbb{P}^+(F)$, that is, having the reconstructed features of anomalous samples to look similar to normal ones. In this way, the auto-encoder transforms an anomalous sample into a normal one effectively, thereby partially achieving the appealing property P1.

Image-Level Alignment To align the generative distribution $\mathbb{P}_g(F)$ with normal feature distribution $\mathbb{P}^+(F)$, a generative adversarial network (GAN) (Goodfellow et al. 2014) of the form below can be used to regularize the generator

$$\min_G \max_D \mathbb{E}_{F \sim \mathbb{P}^+(F)} \log[D(F)] + \mathbb{E}_{F \sim \mathbb{P}_g(F)} \log[1 - D(F)] \quad (7)$$

After convergence, the distribution of $G_\theta(F)$ (i.e., $\mathbb{P}_g(F)$) will be equal to $\mathbb{P}^+(F)$. Although the vanilla GAN can align the generative distribution to the normal feature distribution, the valuable collected anomalies are only used to construct the input distribution $\mathbb{P}(F)$. Since the discriminator has not been informed of any anomaly information, the discriminator could mistakenly recognize anomalies as normal samples, compromising the model's ability of transforming anomaly images into normal ones. To address this issue, we propose to use the anomaly dataset \mathcal{X}^- to enhance the discriminator's ability to distinguish between normal and anomalous samples. To this end, we propose an energy-based discriminator $D(\cdot) : \mathcal{F} \rightarrow [0, +\infty)$, which is designed to assign low energy to normal features and high energy to generated and anomalous features. The discriminator is trained by minimizing the following loss

$$L_D(D, G_\theta) = \mathbb{E}_{F \sim \mathbb{P}^+(F)} D(F) + \gamma \mathbb{E}_{F \sim \mathbb{P}_g(F)} [0, a - D(F)]^+ + (1 - \gamma) \mathbb{E}_{F \sim \mathbb{P}^-(F)} [0, a - D(F)]^+, \quad (8)$$

where $[u, v]^+ = \max(u, v)$, $a > 0$ is a threshold, and $\gamma \in (0, 1]$ is the balance coefficient. With the trained dis-

criminator, the generator is encouraged to output features to get lower energy by minimizing

$$L_G(D, G_\theta) = \gamma \mathbb{E}_{F \sim \mathbb{P}_g(F)} D(F) + (1 - \gamma) \mathbb{E}_{F \sim \mathbb{P}^-(F)} D(F) - \mathbb{E}_{F \sim \mathbb{P}^+(F)} D(F). \quad (9)$$

When (8) and (9) are used to update the discriminator and generator, it can be proved that the output feature distribution $\mathbb{P}_g(F)$ can align with the normal feature distribution $\mathbb{P}^+(F)$. We give the following theorem.

Theorem 1. Let $\mathbb{P}^+(F)$, $\mathbb{P}^-(F)$ and $\mathbb{P}_g(F)$ be the distributions of normal, anomalous, and generated feature maps. Assuming the $\mathbb{P}^+(F)$ and $\mathbb{P}^-(F)$ are two disjoint distributions, and $\gamma \in (0, 1]$. When Discriminator $D(\cdot) : \mathcal{F} \rightarrow [0, +\infty)$ and generator $G(\cdot)$ are updated according to (8) and (9), there exists a Nash equilibrium of the system (D^*, G^*) such that $\mathbb{P}_g(F) = \mathbb{P}^+(F)$.

Proof. Please see the proof in the extended version.

The discriminator trained with loss (8) has the ability to identify anomalous features from normal ones. If an anomalous image is generated, it will be recognized by the discriminator, which will then incentivize the auto-encoder to transform it into normal one. On the other hand, when input to the encoder is a normal sample, output from the generator will be similar to the input, thus the adversarial learning will almost not affect this sample's reconstruction. In this way, we achieve the two appealing properties P1 and P2.

Patch-Level Alignment The proposed knowledge-aware discriminator in (8) distinguishes between normal and anomalous features at image-level. In practice, anomalies often only occupy small areas of an image. Hence, to achieve more accurate detection and localization, it is better to perform the alignment at the patch level. If we directly divide an image into numerous patches and then pass them through the pre-trained ResNet and the auto-encoder, it will be very expensive. Fortunately, the feature at position (h, w) (i.e.,

$F(h, w)$) corresponds to a patch in the origin image, thus it can be viewed as the feature of a patch. Similarly, the generated feature $\tilde{F}(h, w)$ at position (h, w) can be viewed as the generated patch feature. As shown in Figure 1, a_+^ℓ and a_-^ℓ represent a patch feature of a normal and abnormal image, respectively.

Denote the distribution of features of normal patches as $\mathbb{P}^+(p)$. Obviously, features of all patches from normal images follow the distribution $\mathbb{P}^+(p)$. For the generated feature $\tilde{F}(h, w)$ at position (h, w) , we denote its distribution as $\mathbb{P}_g^{(h,w)}(p)$, and further denote the generator that reconstructs the (h, w) -th patch feature as $G_\theta^{(h,w)}(\cdot)$. To align the distribution of generated patch features with that of normal patch features, unlike the alignment at the image-level, no path-level anomaly annotations are available. For an anomalous image, it contains both normal and anomalous patches. Thus, for images from \mathcal{X}^- , we only know that its patch features follow a mixture distribution

$$\mathbb{P}_m(p) = \beta\mathbb{P}^+(p) + (1 - \beta)\mathbb{P}^-(p), \quad (10)$$

where $\beta \in (0, 1)$ is the unknown mixed ratio of $\mathbb{P}^+(p)$ and $\mathbb{P}^-(p)$, with $\mathbb{P}^-(p)$ denoting the distribution of anomalous patches. Although the available anomalous knowledge is noisy, the following theorem shows that we can still make use of it to enhance the auto-encoder’s ability to recover normal patches from anomalous ones.

Theorem 2. *Suppose the $\mathbb{P}^+(p)$ and $\mathbb{P}^-(p)$ are two disjoint distributions and $\gamma \in (0, 1]$. When Discriminator $D^{(h,w)}(\cdot) : \mathcal{P} \rightarrow [0, +\infty)$ and generator $G^{(h,w)}(\cdot)$ are updated by minimizing the two losses*

$$\begin{aligned} L_D(D^{(h,w)}, G_\theta^{(h,w)}) &= \gamma \mathbb{E}_{p \sim \mathbb{P}_g^{(h,w)}(p)} [0, a - D^{(h,w)}(p)]^+ \\ &\quad + (1 - \gamma) \mathbb{E}_{p \sim \mathbb{P}_m(p)} [0, a - D^{(h,w)}(p)]^+ \\ &\quad + \mathbb{E}_{p \sim \mathbb{P}^+(p)} D^{(h,w)}(p), \end{aligned} \quad (11)$$

$$\begin{aligned} L_G(D^{(h,w)}, G_\theta^{(h,w)}) &= \gamma \mathbb{E}_{p \sim \mathbb{P}_g^{(h,w)}(p)} D^{(h,w)}(p) \\ &\quad + (1 - \gamma) \mathbb{E}_{p \sim \mathbb{P}_m(p)} D^{(h,w)}(p) \\ &\quad - \mathbb{E}_{p \sim \mathbb{P}^+(p)} D^{(h,w)}(p), \end{aligned} \quad (12)$$

there exists a Nash equilibrium of the system $(D^{*(h,w)}, G^{*(h,w)})$ such that $\mathbb{P}_g^{(h,w)}(p) = \mathbb{P}^+(p)$.

Proof. Please see the proof in the extended version.

According to (11), the discriminator is designed to assign low energy to normal patch features, while assigning high energy to patch features that are generated or from anomalous images. Although the patch discriminator seems contradictory in the sense of encouraging to output low and high energy to normal patch features simultaneously, due to the term $\mathbb{E}_{p \sim \mathbb{P}^+(p)} D(p)$ in (11) has greater weight, its influence will override that of anomalous images $\mathbb{E}_{p \sim \mathbb{P}_m(p)} [0, a - D^{(h,w)}(p)]^+$, thus leading to the overall low energy for normal patches. With the discriminator capable of outputting low and high energy for normal, anomalous

and generated patches, by driving the generator to generate patches with low energy, we can align the distribution of generated patches with normal ones’. Thus, by comparing the input patch features with the reconstructed ones, anomalous areas can be localized accurately. In experiments, a common patch feature discriminator $D(p)$ is used for all patch feature generators $G^{(h,w)}(p)$.

Training and Testing

Training For patch feature from layer $s \in \mathcal{S}$, the discriminator $D^s(p)$ can be used to distinguish patch features of all positions as (11). Combining discriminators of all layers that participate in the reconstruction, we can get a total discriminator loss and corresponding generator loss $L_D = \sum_{s \in \mathcal{S}} \frac{1}{H^s W^s} \sum_{h=1}^{H^s} \sum_{w=1}^{W^s} L_D(D^s, G_\theta^{(h,w)})$ and $L_G = \sum_{s \in \mathcal{S}} \frac{1}{H^s W^s} \sum_{h=1}^{H^s} \sum_{w=1}^{W^s} L_G(D^s, G_\theta^{(h,w)})$, where H^s and W^s are the height and width of the feature map F^s . As the structure shown in Figure 1, we use the knowledge-aware adversarial learning to regularize the output features of auto-encoder, and train the model by minimizing the following two losses alternatively

$$\min_D L_D, \quad \min_\theta L_{rec}^+ + \lambda L_G. \quad (13)$$

Testing For any test sample x , we first obtain its pre-trained feature maps F^s , and reconstructed feature maps \tilde{F}^s . For pixel-level anomaly detection, we compute the value

$$S_{map}^s(h, w) = 1 - \cos(F^s(h, w), \tilde{F}^s(h, w)) \quad (14)$$

of position (h, w) in anomaly maps for $s \in \mathcal{S}$, and upsample them to the same size as the input image by bi-linear interpolation. Finally, we sum the upsampled error maps of all layers together and smooth the summed map with a Gaussian kernel, yielding the final anomaly score map S_{map} . For image-level anomaly detection, we obtain the average result of the top-k values in the anomaly maps to produce the final anomaly score S_{img} .

Experiments

Experimental Setups

Datasets We mainly conduct experiments on real-world anomaly detection medical datasets and industrial datasets. **Medical datasets:** *i) ISIC2018* (Tschandl, Rosendahl, and Kittler 2018; Codella et al. 2019): The ISIC2018 challenge dataset (task 3) contains 7 categories and we classify NV (nevus) as normal, the rest 6 categories as abnormal. *ii) Chest X-rays* (Kermany et al. 2018) contains normal and abnormal Chest X-rays scans. *iii) Br35H* (Hamada 2020; Zhou et al. 2024): Brain Tumor Detection dataset contains non-tumorous and various tumorous images. *iv) OCT* (Kermany et al. 2018): Retinal optical coherence tomography (OCT) contains normal OCT scans and three types of scans with diseases. **Industrial dataset:** *i) MVTEC AD* (Bergmann et al. 2019) is a widely known industrial dataset comprising 15 classes with 5 textures and 10 objects. *ii) Visa* (Zou et al. 2022) is an industrial dataset containing 12 classes.

	r_l	Methods	ISIC2018			Chest X-ray			Br35H			OCT		
			AUC	F1	ACC	AUC	F1	ACC	AUC	F1	ACC	AUC	F1	ACC
Unsupervised	0%	GANomaly	72.54	60.95	56.99	76.14	81.18	74.36	84.43	78.89	75.83	90.52	91.09	86.16
		DifferNet	76.52	64.58	64.25	86.41	84.62	79.17	91.66	84.98	85.33	94.27	92.85	89.05
		Fastflow	80.81	69.95	70.98	90.25	86.02	80.77	93.85	87.67	86.83	94.08	92.60	88.43
		AE-FLOW	87.79	80.56	84.97	92.00	88.92	85.58	N/A	N/A	N/A	98.15	96.36	94.42
		ReContrast	90.15	81.12	86.01	84.20	83.45	77.24	96.77	93.95	93.67	99.60	98.53	97.80
		Recon	87.51	79.50	82.90	85.59	83.70	78.04	97.37	96.23	96.16	99.33	97.84	96.80
Weakly Supervised	1%	PReNet	73.84	67.02	69.26	86.66	84.83	80.45	81.86	76.66	75.33	60.26	85.88	75.83
		Deep SAD	89.52	77.13	80.83	86.82	83.40	79.33	86.83	81.41	80.67	93.04	90.93	86.50
		AA-BiGAN	87.83	78.96	83.25	90.63	87.52	84.13	88.87	84.04	83.67	96.02	94.41	91.70
		DRA	90.10	78.51	82.04	92.35	88.36	85.59	90.41	84.18	84.33	91.52	91.34	86.83
		DevNet	88.98	77.80	80.83	92.53	89.41	85.57	93.94	88.29	88.33	99.38	98.75	98.13
		CKAAD (Ours)	91.02	82.04	85.84	93.22	88.77	86.22	98.89	97.88	97.83	99.88	99.33	99.00
	2%	PReNet	74.28	66.87	68.83	89.69	87.47	84.29	82.33	75.00	76.33	65.25	85.98	75.90
		Deep SAD	89.86	79.18	83.94	90.16	86.56	83.97	85.91	79.48	76.33	95.26	93.18	90.00
		AA-BiGAN	88.20	78.10	83.25	92.13	89.86	87.02	84.46	79.87	79.17	97.15	95.44	93.20
		DRA	90.65	79.70	82.80	92.59	89.43	86.69	95.00	89.92	89.50	94.98	92.94	89.50
		DevNet	89.36	78.80	83.25	94.27	90.48	87.66	94.89	88.05	88.33	99.31	98.56	97.83
		CKAAD (Ours)	91.68	82.86	86.87	95.23	91.17	88.94	99.07	98.03	98.00	99.91	99.42	99.13
	5%	PReNet	74.29	67.36	71.24	91.01	87.82	83.81	85.12	80.48	78.33	67.67	86.27	76.47
		Deep SAD	89.81	78.29	82.38	92.87	88.49	85.58	91.54	83.50	83.33	95.74	93.88	91.03
		AA-BiGAN	88.91	78.95	82.30	92.80	89.95	86.86	95.96	91.87	91.83	97.69	96.04	94.10
		DRA	91.05	79.80	83.07	93.61	89.59	86.86	95.48	90.03	89.67	91.91	91.39	86.57
		DevNet	89.93	78.63	82.64	95.77	91.23	88.94	99.52	97.80	97.83	99.60	98.93	98.40
		CKAAD (Ours)	92.66	84.00	87.65	96.57	93.87	92.47	99.79	98.20	98.17	99.93	99.55	99.33

Table 1: Performance (Average AUC(%), best F1(%)) and ACC(%) on medical datasets. With $k = 1$, we increase the ratio of labeled anomalies r_l in the training set from 0% to 5%.

Implementation All images are resized to 256. α, γ are set to 0.5. λ is set to 0.02. Adam optimizer is utilized with $\beta = (0.5, 0.999)$. The learning rate for the auto-encoder is set to 1e-03 for medical datasets, 5e-03 for industrial datasets, and 1e-04 for the discriminator. Resnet18 is chosen as the backbone and $\mathcal{S} = \{2, 3\}$ for medical datasets. WideResnet50 and $\mathcal{S} = \{1, 2, 3\}$ are set for industrial datasets because the anomalies are more subtle.

Metrics We report the AUC (Area Under receiver operating characteristic Curve), best F1, Accuracy (ACC) for medical datasets, and image-level AUC, pixel-level AUC and PRO (Per-Region Overlap) for industrial datasets.

Experimental Scenarios For medical datasets, we evaluate the performance with consideration of two key parameters for the incomplete anomalous information: 1) the number of types of collected anomalies k ; 2) the ratio of collected anomalies r_l in the training data. Specifically, we first fix $k = 1$ but increase r_l gradually. In this case, anomalies from any anomaly category can be used as the only labeled anomaly class during training, and performance averaged over all anomaly types is reported. Then, we fix r_l but increase k gradually. The k anomaly types are randomly selected from all possible anomaly categories, so the performance is averaged over 5 random experiments. For industrial datasets, there are no anomalies for training. So elastic transformation (see details in the extended version) is used to distort normal images to produce corresponding anomalies

and the average performance of all classes is reported.

Detection Results on Medical Datasets

We compared our proposed methods with several unsupervised methods GANomaly(Akay, Atapour-Abarghouei, and Breckon 2019), DifferNet(Rudolph, Wandt, and Rosenhahn 2021), Fastflow(Yu et al. 2021), AE-FLOW(Zhao, Ding, and Zhang 2023), ReContrast(Guo et al. 2023), and weakly-supervised methods Deep SAD(Ruff et al. 2019), DevNet(Pang et al. 2021a), AA-BiGAN(Tian, Su, and Yin 2022), DRA(Ding, Pang, and Shen 2022), PReNet (Pang et al. 2023). All methods are run under the same settings.

Performance under different ratios of labeled anomalies As shown in Table 1, with $k = 1$, the performance of unsupervised and weakly supervised methods is reported. Compared to the reconstruction baseline (Recon), our proposed method can enhance the reconstruction model’s performance by utilizing labeled anomalies. Using only 1% anomalies, our method improves the AUC by 3.51% for ISIC2018, 7.63% for Chest X-ray, 1.52% for Br35H. While adding r_l to 2% and 5%, the performance continuously increases. On the other hand, compared to other weakly-supervised methods, our method exhibits the best overall performance of three metrics on all medical datasets.

Performance under different types of labeled anomalies As Figure 2 shows, we evaluate the performance with different numbers of labeled anomaly types on ISIC2018 which




Figure 2: Histograms of performance (Average AUC and F1 of 5 random experiments) with different number (k) of labeled anomaly classes on ISIC2018 under $r_l = 5\%$.

Dataset	MVTec			Visa		
	I-AUC	P-AUC	PRO	I-AUC	P-AUC	PRO
PaDiM	95.4	97.5	92.1	89.1	98.1	85.9
DRAEM	98.0	97.3	93.0	88.7	93.5	73.1
RD4AD	98.5	97.8	93.9	96.0	90.1	70.9
PatchCore	99.1	98.1	93.4	95.1	98.8	91.2
DesTseg	98.6	97.9	92.6	92.3	98.4	92.3
RD++	99.4	98.3	95.0	95.9	98.7	93.4
DMAD	99.5	98.2	90.6	95.5	98.6	91.3
D3AD	97.2	97.4	93.3	96.0	97.9	94.1
CKAAD (Ours)	99.5	98.4	95.2	96.7	98.8	94.0

Table 2: Performance on industrial datasets. Average Image-level / Pixel-level AUC (I-AUC / P-AUC) and PRO of all classes are reported.

has 6 kinds of anomalies. With only 5% anomalies, the type k is increased from 1 to 3, and the histogram exhibits the average AUC and best F1 of 5 weakly-supervised methods. With the increasing number of observed anomaly types, weakly-supervised methods can continuously improve detection performance. Among them, our method gets the best overall performance under all k settings.

Detection Results on Industrial Datasets

We also evaluate fine-grained anomaly detection on industrial datasets and report the performance in Table 2. Compared to recent state-of-the-art approaches PaDim (Defard et al. 2021), DRAEM (Zavrtanik, Kristan, and Skočaj 2021a), RD4AD (Deng and Li 2022), PatchCore (Roth et al. 2022b), DesTseg (Zhang et al. 2023), RD++ (Tien et al. 2023), DMAD (Liu et al. 2023a) and D3AD (Tebbe and Tayyub 2024), our method still gets competitive performance in both image-level and pixel-level detection, demonstrating our effectiveness in fine-grained anomaly detection.

Ablation Study

There are several strategies of utilizing anomalies, including the Recon (2), ReconSub (4), pure GAN (7), coarse-knowledge-aware image-level adversarial learning (CKAImg)(8)(9) and patch-level adversarial learning (CKAPatch)(11)(12). To compare these strategies, we conduct experiments with $r_l = 5\%$, $k = 1$ for medical datasets and distorted anomalies for industrial datasets. As shown in Table 3, compared to Recon, ReconSub can decrease performance a lot on MVTec, indicating it is unsuitable for fine-grained

Strategy	ISIC2018		Chest X-Ray		MVTec		
	AUC	F1	AUC	F1	I-AUC	P-AUC	PRO
Recon	87.5	79.5	85.6	83.7	99.1	98.2	94.4
ReconSub	89.2	79.0	92.3	87.3	91.2	95.0	86.2
GAN	91.0	82.0	95.2	91.1	99.2	98.2	94.9
CKAImg	91.4	82.3	96.0	92.0	99.3	98.3	95.1
CKAPatch	92.6	84.0	96.6	93.9	99.5	98.4	95.2

Table 3: Ablation of different strategies.




Figure 3: Visualization of detection results. Images, anomaly maps of Recon, ReconSub, and our method are shown in the first column and last three columns respectively. The second column in (b) circles anomalous areas by red lines.

anomaly localization. Compared to pure GAN, CKAImg improves the detection performance more, that is because the energy discriminator is knowledge-aware of anomalies, leading the reconstructed feature to be better aligned with normal ones. Furthermore, employing CKAPatch can make the discriminator be better aware of anomalous areas, thus guiding the auto-encoder to output normal patch features more surely, further boosting the detection performance.

Figure 3 exhibits image error maps of different strategies. The reconstruction quality can be reflected by the error maps in which high values represent the area not reconstructed well. The visualization result shows that only reconstructing normal samples can also generalize to anomalies, thus anomalous areas can not be detected in the error maps. ReconSub can negatively influence the reconstruction quality of normal samples while not accurately localizing the anomalous areas. The proposed patch-level knowledge-aware method can reconstruct normal images well while having the ability to localize the anomalous areas.

Conclusion

In this paper, we propose to leverage a small coarse-labeled anomaly dataset to boost the anomaly detection performance of the feature reconstruction model. To achieve our goal, we first introduce the energy-based discriminators which can be aware of anomalous information to guide the feature alignment at image-level, then transfer the alignment to patch-level by only using the coarse image labels to better localize the anomalous area. Experimental results on medical and industrial datasets prove the effectiveness of our method.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No. 62276280, 62276279), Guangzhou Science and Technology Planning Project (No. 2024A04J9967), Guangdong Basic and Applied Basic Research Foundation (2024B1515020032).

References

- Akçay, S.; Atapour-Abarghouei, A.; and Breckon, T. P. 2019. Ganomaly: Semi-supervised anomaly detection via adversarial training. In *14th Asian Conference on Computer Vision*, 622–637.
- Bergmann, P.; Fauser, M.; Sattlegger, D.; and Steger, C. 2019. MVTEC AD-A comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9592–9600.
- Cho, H.; Seol, J.; and Lee, S.-g. 2021. Masked Contrastive Learning for Anomaly Detection. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, 1434–1441.
- Codella, N.; Rotemberg, V.; Tschandl, P.; Celebi, M. E.; Dusza, S.; Gutman, D.; Helba, B.; Kalloo, A.; Liopyris, K.; Marchetti, M.; et al. 2019. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). arXiv:1902.03368.
- Defard, T.; Setkov, A.; Loesch, A.; and Audigier, R. 2021. Padim: a patch distribution modeling framework for anomaly detection and localization. In *International Conference on Pattern Recognition*, 475–489.
- Deng, H.; and Li, X. 2022. Anomaly detection via reverse distillation from one-class embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9737–9746.
- Ding, C.; Pang, G.; and Shen, C. 2022. Catching both gray and black swans: Open-set supervised anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7388–7398.
- Fernando, T.; Gammulle, H.; Denman, S.; Sridharan, S.; and Fookes, C. 2021. Deep learning for medical anomaly detection—a survey. *ACM Computing Surveys (CSUR)*, 54(7): 1–37.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Gui, X.; Wu, D.; Chang, Y.; and Fan, S. 2022. Constrained Adaptive Projection with Pretrained Features for Anomaly Detection. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, 2059–2065.
- Guo, J.; Lu, S.; Jia, L.; Zhang, W.; and Li, H. 2023. ReContrast: Domain-Specific Anomaly Detection via Contrastive Reconstruction. In *Advances in Neural Information Processing Systems*, 10721–10740.
- Hamada. 2020. Br35h: Brain tumor detection 2020. <https://www.kaggle.com/datasets/ahmedhamada0/brain-tumor-detection>, 2020.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Kermany, D. S.; Goldbaum, M.; Cai, W.; Valentim, C. C.; Liang, H.; Baxter, S. L.; McKeown, A.; Yang, G.; Wu, X.; Yan, F.; et al. 2018. Identifying medical diagnoses and treatable diseases by image-based deep learning. *cell*, 172(5): 1122–1131.
- Kimura, D.; Chaudhury, S.; Narita, M.; Munawar, A.; and Tachibana, R. 2020. Adversarial discriminative attention for robust anomaly detection. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 2172–2181.
- Li, C.-L.; Sohn, K.; Yoon, J.; and Pfister, T. 2021. Cutpaste: Self-supervised learning for anomaly detection and localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9664–9674.
- Liu, J.; Xie, G.; Wang, J.; Li, S.; Wang, C.; Zheng, F.; and Jin, Y. 2024. Deep industrial image anomaly detection: A survey. *Machine Intelligence Research*, 21(1): 104–135.
- Liu, W.; Chang, H.; Ma, B.; Shan, S.; and Chen, X. 2023a. Diversity-measurable anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12147–12156.
- Liu, Z.; Zhou, Y.; Xu, Y.; and Wang, Z. 2023b. SimpNet: A simple network for image anomaly detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20402–20411.
- Lv, W.; Su, Q.; Wan, H.; Xu, H.; and Xu, W. 2024. Contamination-Resilient Anomaly Detection via Adversarial Learning on Partially-Observed Normal and Anomalous Data. In *Forty-first International Conference on Machine Learning*.
- Pang, G.; Ding, C.; Shen, C.; and Hengel, A. v. d. 2021a. Explainable deep few-shot anomaly detection with deviation networks. arXiv:2108.00462.
- Pang, G.; Shen, C.; Jin, H.; and van den Hengel, A. 2023. Deep weakly-supervised anomaly detection. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1795–1807.
- Pang, G.; van den Hengel, A.; Shen, C.; and Cao, L. 2021b. Toward deep supervised anomaly detection: Reinforcement learning from partially labeled anomaly data. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, 1298–1308.
- Reiss, T.; Cohen, N.; Bergman, L.; and Hoshen, Y. 2021. Panda: Adapting pretrained features for anomaly detection and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2806–2814.
- Reiss, T.; and Hoshen, Y. 2023. Mean-shifted contrastive loss for anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2155–2162.

- Roth, K.; Pemula, L.; Zepeda, J.; Schölkopf, B.; Brox, T.; and Gehler, P. 2022a. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14318–14328.
- Roth, K.; Pemula, L.; Zepeda, J.; Schölkopf, B.; Brox, T.; and Gehler, P. 2022b. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14318–14328.
- Rudolph, M.; Wandt, B.; and Rosenhahn, B. 2021. Same same but different: Semi-supervised defect detection with normalizing flows. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 1907–1916.
- Ruff, L.; Vandermeulen, R.; Goernitz, N.; Deecke, L.; Siddiqui, S. A.; Binder, A.; Müller, E.; and Kloft, M. 2018. Deep one-class classification. In *International conference on machine learning*, 4393–4402.
- Ruff, L.; Vandermeulen, R. A.; Görnitz, N.; Binder, A.; Müller, E.; Müller, K.-R.; and Kloft, M. 2019. Deep Semi-Supervised Anomaly Detection. In *International Conference on Learning Representations*.
- Salehi, M.; Sadjadi, N.; Baselizadeh, S.; Rohban, M. H.; and Rabiee, H. R. 2021. Multiresolution knowledge distillation for anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14902–14912.
- Schlegl, T.; Seeböck, P.; Waldstein, S. M.; Schmidt-Erfurth, U.; and Langs, G. 2017. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International conference on information processing in medical imaging*, 146–157.
- Tack, J.; Mo, S.; Jeong, J.; and Shin, J. 2020. CSI: novelty detection via contrastive learning on distributionally shifted instances. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 11839–11852.
- Tebbe, J.; and Tayyub, J. 2024. Dynamic Addition of Noise in a Diffusion Model for Anomaly Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 3940–3949.
- Tian, B.; Su, Q.; and Yin, J. 2022. Anomaly Detection by Leveraging Incomplete Anomalous Knowledge with Anomaly-Aware Bidirectional GANs. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, 2255–2261.
- Tian, B.; Su, Q.; and Yu, J. 2023. Leveraging contaminated datasets to learn clean-data distribution with purified generative adversarial networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 9989–9996.
- Tien, T. D.; Nguyen, A. T.; Tran, N. H.; Huy, T. D.; Duong, S.; Nguyen, C. D. T.; and Truong, S. Q. 2023. Revisiting reverse distillation for anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24511–24520.
- Tschandl, P.; Rosendahl, C.; and Kittler, H. 2018. The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 5(1): 1–9.
- Vasilev, A.; Golkov, V.; Meissner, M.; Lipp, I.; Sgarlata, E.; Tomassini, V.; Jones, D. K.; and Cremers, D. 2020. q-Space novelty detection with variational autoencoders. In *Computational Diffusion MRI: MICCAI Workshop, Shenzhen, China, October 2019*, 113–124.
- Xia, X.; Pan, X.; Li, N.; He, X.; Ma, L.; Zhang, X.; and Ding, N. 2022. GAN-based anomaly detection: A review. *Neurocomputing*, 493: 497–535.
- Yang, J.; Xu, R.; Qi, Z.; and Shi, Y. 2022. Visual anomaly detection for images: A systematic survey. *Procedia computer science*, 199: 471–478.
- Yi, J.; and Yoon, S. 2020. Patch svdd: Patch-level svdd for anomaly detection and segmentation. In *Proceedings of the Asian conference on computer vision*, 375–390.
- You, Z.; Cui, L.; Shen, Y.; Yang, K.; Lu, X.; Zheng, Y.; and Le, X. 2022. A unified model for multi-class anomaly detection. *Advances in Neural Information Processing Systems*, 35: 4571–4584.
- Yu, J.; Zheng, Y.; Wang, X.; Li, W.; Wu, Y.; Zhao, R.; and Wu, L. 2021. Fastflow: Unsupervised anomaly detection and localization via 2d normalizing flows. arXiv:2111.07677.
- Zavrtanik, V.; Kristan, M.; and Skočaj, D. 2021a. Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8330–8339.
- Zavrtanik, V.; Kristan, M.; and Skočaj, D. 2021b. Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition*, 112: 107706.
- Zhang, J.; Wang, S.; and Chen, S. 2022. Reconstruction Enhanced Multi-View Contrastive Learning for Anomaly Detection on Attributed Networks. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, 2376–2382.
- Zhang, X.; Li, S.; Li, X.; Huang, P.; Shan, J.; and Chen, T. 2023. DeSTSeg: Segmentation Guided Denoising Student-Teacher for Anomaly Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3914–3923.
- Zhao, Y.; Ding, Q.; and Zhang, X. 2023. AE-FLOW: Autoencoders with Normalizing Flows for Medical Images Anomaly Detection. In *The Eleventh International Conference on Learning Representations*.
- Zhou, Q.; Pang, G.; Tian, Y.; He, S.; and Chen, J. 2024. AnomalyCLIP: Object-agnostic Prompt Learning for Zero-shot Anomaly Detection. In *The Twelfth International Conference on Learning Representations*.
- Zou, Y.; Jeong, J.; Pemula, L.; Zhang, D.; and Dabeer, O. 2022. Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. In *European Conference on Computer Vision*, 392–408.