

Focus the Discrepancy: Intra- and Inter-Correlation Learning for Image Anomaly Detection

Xincheng Yao¹, Ruoqi Li¹, Zefeng Qian¹, Yan Luo¹, Chongyang Zhang^{1,2*}

¹School of Electronic Information and Electrical Engineering, Shanghai Jiao Tong University

²MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

{i-Dover, nilponi, zefeng-qian, luoyan_bb, sunny_zhang}@sjtu.edu.cn¹

Abstract

Humans recognize anomalies through two aspects: larger patch-wise representation discrepancies and weaker patch-to-normal-patch correlations. However, the previous AD methods didn't sufficiently combine the two complementary aspects to design AD models. To this end, we find that Transformer can ideally satisfy the two aspects as its great power in the unified modeling of patch-wise representations and patch-to-patch correlations. In this paper, we propose a novel AD framework: *FOCUS-the-Discrepancy (FOD)*, which can simultaneously spot the patch-wise, intra- and inter-discrepancies of anomalies. The major characteristic of our method is that we renovate the self-attention maps in transformers to *Intra-Inter-Correlation (I2Correlation)*. The *I2Correlation* contains a two-branch structure to first explicitly establish intra- and inter-image correlations, and then fuses the features of two-branch to spotlight the abnormal patterns. To learn the intra- and inter-correlations adaptively, we propose the *RBF-kernel-based target-correlations* as learning targets for self-supervised learning. Besides, we introduce an entropy constraint strategy to solve the mode collapse issue in optimization and further amplify the normal-abnormal distinguishability. Extensive experiments on three unsupervised real-world AD benchmarks show the superior performance of our approach. Code will be available at <https://github.com/xcyao00/FOD>.

1. Introduction

The goal of anomaly detection (AD) is to distinguish an instance containing anomalous patterns from those normal samples and further localize those anomalous regions. Anomalies are defined as opposite to normal samples and are usually rare, which means that we need to tackle AD tasks under the unsupervised setting with only normal sam-

*Corresponding Author.

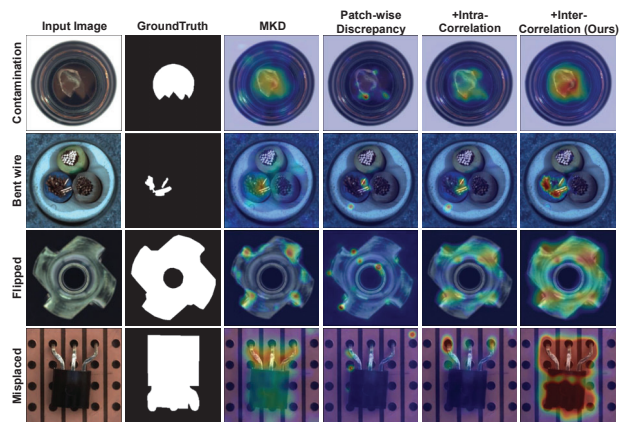


Figure 1. Anomaly detection examples on MVTEC-AD [3]. Multiresolution Knowledge Distillation (MKD) [41] adopts the conventional patch-wise representation discrepancies. Row 1 shows the hard global anomalies (*i.e.*, they are not significantly different from normal visuals). Rows 3 and 4 show the logical anomalies (*i.e.*, they may be easily recognized as normal if only from the patch-wise discrepancy).

ples accessible. The core idea of most unsupervised AD methods is to compare with normal samples to distinguish anomalies [46, 35, 60, 11, 28, 58]. Even for humans, we also recognize anomalies in this way, specifically through three discrepancies, *i.e.*, 1. patch patterns that differentiate from the normal visuals; 2. image regions that destroy textures or structures; 3. novel appearances that deviate from our accumulated knowledge of normality. Namely, anomalous patches usually have three characteristics: their patch-wise representations are different from the normal visuals; they are different from most patches within one image; they deviate from our accumulated knowledge of normality. These views intrinsically reveal that humans' recognition of anomalies depends on two aspects: patch-wise representations (1) and intra- and inter-image correlations (2, 3).

Previous methods mainly follow the former aspect to learn distinguishable representations or reconstructions,

such as reconstruct-based methods [6, 57, 1] and knowledge distillation AD models [4, 41]. The goal of these methods is to generate reconstructed samples or feature representations, and larger patch-wise representation discrepancies can appear in the abnormal patches. However, only the patch-wise representation discrepancies are insufficient for detecting more complex anomalies (e.g., rows 3 and 4 in Figure 1), since the patch-wise errors can't provide comprehensive descriptions of the spatial context. Other mainstream AD methods, such as embedding-based [11, 35] and one-class-classification-based (OCC) [39, 60] methods, are much similar to the latter aspect. These methods achieve anomaly detection by measuring the distances between the features of test samples and normal features. Compared with the non-learnable feature distances, the explicit intra- and inter-image correlations in our method are more effective to detect diverse anomalies (see Table 1, 2, 3). Moreover, patch-wise representation discrepancies and intra- and inter-correlation discrepancies are complementary, and can be combined to develop more powerful AD models.

Recently, with the self-attention mechanism and long-range modeling ability, transformers [50] have significantly renovated many computer vision tasks [14, 24, 7, 67, 56] and recently popular language-vision multimodal tasks [30, 22]. Transformers have shown great power in the unified modeling of patch-wise representations and patch-to-patch correlations. Transformers are quite suitable for AD tasks as their modeling ability can satisfy the two aspects of anomaly recognition quite well. Some works [26, 18, 8, 62] also attempt to employ transformers to construct AD models. However, these methods only use transformers to extract vision features, which didn't sufficiently adapt transformers' long-range correlation modeling capability to AD tasks. Different from these works, we explicitly exploit transformers' self-attention maps to establish the intra- and inter-image correlations. The correlation distribution of each patch can provide more informative descriptions of the spatial context, which can reveal more intricate and semantic anomaly patterns.

In this paper, motivated by humans' anomaly recognition process, we propose a novel AD framework: FOCUS-the-Discrepancy (FOD), which can exploit transformers' unified modeling ability to simultaneously spot the patch-wise, intra- and inter-discrepancies. Our key designs are composed of three recognition branches: the patch-wise discrepancy branch is to reconstruct the input patch features for distinguishing simple anomalies; the intra-correlation branch is to explicitly model patch-to-patch correlations in one image for distinguishing hard global anomalies (e.g., row 1 in Figure 1); the inter-correlation branch is to explicitly learn inter-image correlations with known normal patterns from the whole normal training set. To implement the intra- and inter-correlation branches, we adapt

Transformer and renovate the self-attention mechanism to the I2Correlation, which contains a two-branch structure to first separately model the intra- and inter-correlation distribution of each image patch, and then fuse the features of two-branch to spotlight the abnormal patterns. To learn the intra- and inter-correlations adaptively, we propose the RBF-kernel-based target-correlations as learning targets for self-supervised learning, the RBF kernel is used to present the neighborhood continuity of each image patch. Besides, an entropy constraint strategy is applied in the two branches, which can solve the mode collapse issue in optimization and further amplify the normal-abnormal distinguishability.

In summary, we make the following main contributions:

1. We propose a novel AD framework: FOD, which can effectively detect anomalies by simultaneously spotting the patch-wise, intra- and inter-discrepancies.
2. We renovate the self-attention mechanism to the I2Correlation, which can explicitly establish intra- and inter-correlations in a self-supervised way with the target-correlations. An entropy constraint strategy is proposed to further amplify the normal-abnormal distinguishability.
3. Our method can achieve SOTA results on three real-world AD datasets, this shows our method can more effectively determine anomalies from complementary views.

2. Related Work

Anomaly Detection. In this paper, we divide the mainstream AD methods into five categories: reconstruction-based, embedding-based, OCC-based methods, knowledge distillation and normalizing flow AD models. The reconstruction-based methods are the most popular AD methods and also widely studied, where the assumption is that models trained by normal samples only can reconstruct normal regions but fail in abnormal regions. Many previous works attempt to train AutoEncoders [6, 28, 57, 8, 17, 64], Variational AutoEncoders [23] and GANs [42, 1, 29, 40] to reconstruct the input images. Overfitting to the input images is the most serious issue of these methods, which means that the anomalies are also well reconstructed [62].

Recently, some embedding-based methods [2, 33, 9, 11, 35] show better AD performance by using ImageNet pre-trained networks as feature extractors. In [33], the authors fit a multivariate Gaussian to model the image-level features for further Mahalanobis distance measurement. PaDiM [11] extends the above method to localize pixel-level anomalies. PatchCore [35] extends on this line by utilizing locally aggregated, mid-level features and introducing greedy coreset subsampling to form nominal feature banks. However, these methods directly utilize pre-trained networks without any adaptation to the target dataset. Some works [31, 10, 32] attempt to adapt pre-trained features to the target data distribution. There are also some other methods for using pre-trained networks by freezing them and only

optimizing a sub-network, *e.g.*, via knowledge distillation [4, 41, 12, 38], or normalizing flows [36, 15, 37, 63].

OCC is another classical AD paradigm, the earliest works are mainly to extend the OCC models such as OC-SVM [44] or SVDD [47, 39] for anomaly detection. Recently, in [60], a patch-based SVDD that contains multiple cores rather than a single core in [39] is proposed to enable anomaly localization. In [25], a Fully Convolutional Data Description combined with receptive field upsampling is proposed to generate anomaly maps. In [48], the authors further extend the PatchSVDD [60] model by the proposed multi-scale patch-based representation learning method.

Transformer-based Anomaly Detection Methods. Recently, transformers [50] have shown great power in modeling long-range dependencies. For image anomaly detection, some works [26, 18, 62, 8, 59] also attempt to exploit transformers to design AD models. However, most of these methods only use transformers as backbones to extract vision features, and don't effectively adapt the long-range modeling capacity of transformers to AD tasks. Unlike the previous usage of transformers, we explicitly exploit the self-attention maps of transformers to establish intra- and inter-correlations. Our work shares some similarities with a recent work [20]. But we point out that our work has some significant differences from [20]: 1) Different Insight: [20] aims to unify the pointwise representation and pairwise association for time series AD, whereas our work aims to more sufficiently detect image anomalies through three complementary recognition views. 2) Novel Method: we propose entropy constraint, inter-correlation branch, and I2Correlation by effectively combining intra- and inter-correlation branches. 3) Different Task: [20] focuses on time-series AD, whereas we focus on image AD.

3. Approach

3.1. Model Overview

Figure 2 overviews our proposed approach. The model consists of three branches: patch-wise discrepancy branch, intra-correlation branch, and inter-correlation branch. The input 2D image is first sent into a pre-trained backbone to extract multi-scale feature maps. At each feature level, we construct a subsequent transformer network to explicitly model the intra- and inter-image correlations for spotting the intra- and inter-discrepancies. Each network is implemented by stacking the I2Correlation blocks and feed-forward layers alternately. Suppose each network contains K layers with length- N input features $X \in \mathbb{R}^{N \times d}$, the output of k th layer $X_k \in \mathbb{R}^{N \times d_m}$ is calculated as follows:

$$Z_k = \text{LN}(\text{I2Correlation}(X_{k-1}, X_f) + X_{k-1}) \quad (1)$$

$$X_k = \text{LN}(\text{FeedForward}(Z_k) + Z_k) \quad (2)$$

where X_f is the reference features used by the inter-correlation branch, LN means LayerNorm, and $Z_k \in \mathbb{R}^{N \times d_m}$ is the k th layer's hidden features. The final output features are calculated by linear projection: $\hat{X} = Z_K W_o$, where $W_o \in \mathbb{R}^{d_m \times d}$ is the output projection matrix.

3.2. Patch-Wise Reconstruction Discrepancy

Adopted from previous reconstruction-based AD methods, we employ feature reconstruction as our patch-wise recognition view for simplicity. In our approach, we construct a transformer network to reconstruct the input features. With the long-range dependency modeling ability, the features reconstructed by Transformer can have larger effective receptive fields [13], which are more conducive to detect hard global anomalies and logical anomalies. Moreover, with the further introduced inter-correlation branch (sec.3.4), the features can even perceive normal regions of images from the whole normal training set. Therefore, the reconstructed features generated by our model have better global perception and more discriminative semantic representation capability, which are more suitable for anomaly detection by patch-wise representation discrepancies.

Learning Objective. We can utilize classical reconstruction losses as the learning objective. We combine the ℓ_2 distance and cosine distance to measure the feature distances between the reconstructed features \hat{X} and the input features X . The loss function is defined as follows:

$$\mathcal{L}_l = \|\hat{X} - X\|_2 + \left(1 - \frac{\hat{X} \cdot X}{\|\hat{X}\|_2 \|X\|_2}\right) \quad (3)$$

3.3. Intra-Correlation Learning

The intra-correlation learning branch aims to learn informative patch-to-patch correlations from the input patch sequence adaptively. As shown in the right part of Figure 2, we explicitly take advantage of the self-attention maps of transformers as intra-correlation matrices. Formally, the intra-correlation matrix of the k th layer is calculated by:

$$\text{Intra Correlation} : S_k^g = \text{Softmax}\left(Q_k^g (K_k^g)^T / \sqrt{d_m}\right) \\ [Q_k^g, K_k^g, V_k^g] = X_{k-1} [W_k^{Qg}, W_k^{Kg}, W_k^{Vg}] \quad (4)$$

where $Q_k^g, K_k^g, V_k^g \in \mathbb{R}^{N \times d_m}$ represent the query, key, and value of the k th layer. $W_k^{Qg}, W_k^{Kg}, W_k^{Vg} \in \mathbb{R}^{d_m \times d_m}$ represent the learnable projection matrices for Q_k^g, K_k^g, V_k^g , respectively. $S_k^g \in \mathbb{R}^{N \times N}$ denotes the learned intra-image patch-to-patch correlations. Since $\text{Softmax}(\cdot)$ can convert the values in the similarity map into range $[0, 1]$ along the horizontal axis, each row of S_k^g can represent a discrete correlation distribution for each corresponding image patch. But different from the vanilla transformers, we further introduce target correlation matrices as the learning objective to explicitly optimize the intra-image correlations.

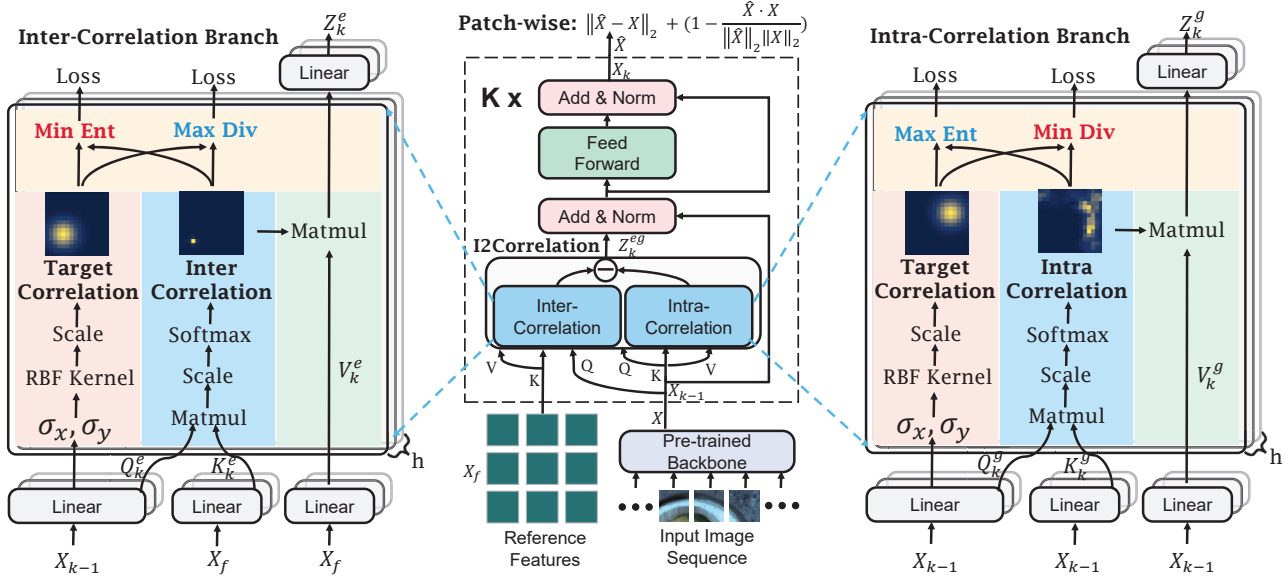


Figure 2. Model overview. The model is composed of three parts: patch-wise discrepancy branch, intra-correlation branch, and inter-correlation branch. The input image patch sequence will be transformed into the hidden features by a pre-trained backbone, and then sent into the intra-correlation branch for learning intra-image patch-to-patch correlations. The input of the inter-correlation branch is composed of hidden features (as Query) and reference features (as Key and Value). The final output features are used to reconstruct the input features.

Target Correlation. Inspired by the contrastive learning method BarlowTwins [66], we can construct a target correlation matrix $T \in \mathbb{R}^{N \times N}$ as the learning target. The objective function of BarlowTwins measures the correlation matrix between the embeddings of two identical networks fed with distorted versions of a batch of samples, and tries to make this matrix close to the identity. In our work, the role of the target correlation is to introduce a prior correlation of patches as a pretext learning target, where each patch can be highly correlated to itself and also correlated to its neighborhood patches and the correlation decreases with the increase of distance. This allows us to optimize the intra- and inter-correlations in a self-supervised way. To this end, we use the radial basis function (RBF) to construct the target correlation matrix. We further adopt two learnable kernel variances σ_x and σ_y for horizontal and vertical axes to make the target correlation of each patch can adapt to the specific pattern of itself. The target correlation matrix of the k th layer is defined as:

$$T_k^g = \frac{1}{2\pi\sigma_x\sigma_y} \exp\left(-\frac{\|x_{ij} - x_{i'j'}\|_2^2}{2(\sigma_x^2 + \sigma_y^2)}\right) \quad (5)$$

$i, i' \in \{1, \dots, H\}; j, j' \in \{1, \dots, W\}$

where $\|x_{ij} - x_{i'j'}\|_2^2$ means the Square Euclidean distance between point x_{ij} and $x_{i'j'}$, i, i' and j, j' means vertical and horizontal coordinates, respectively.

Next, we need to measure the distance between target- and intra-correlation distributions. This can usually be

achieved by calculating the KL divergence. We can obtain a KL divergence value from each level of the network. Thus, we average all KL divergence values to combine the patch-to-patch correlations from multi-layer features into a more informative measure as follows:

$$\text{Div}(T^g, S^g) = \frac{1}{K} \sum_{k=1}^K \left(KL(T_k^g \| S_k^g) + KL(S_k^g \| T_k^g) \right) \quad (6)$$

where $KL(\cdot \| \cdot) \in \mathbb{R}^N$ and its each element means the KL divergence between two discrete distributions corresponding to each row of T_k^g and S_k^g .

Due to the rarity of anomalies and the dominance of normal patterns, the normal patches should build strong correlations with most patches in the whole image, while the weights of abnormal correlation distributions are harder to distribute to most patches and are more likely to concentrate on the adjacent image patches due to the neighborhood continuity. Since normal and abnormal patches have different correlation distributions, this is a distinguishable criterion for anomaly detection. Note that the intra-correlation branch explicitly exploits the spatial dependencies of each image patch, which are more informative than the patch-wise representations for anomaly detection.

Entropy constraint. Since the normal image patterns are usually diverse, the learned correlation distributions of the normal patches may also easily concentrate on the adjacent patches, which can cause the distinguishability between normal and abnormal to be downscaled. To address

this, we further introduce an entropy constraint item for making normal patches establish strong associations with most normal patches in the whole image as much as possible. The entropy constraint item is defined as:

$$\text{Ent}(\mathcal{S}^g) = \frac{1}{K} \sum_{k=1}^K \left(\sum_{i=1}^N \sum_{j=1}^N (-S_{k(i,j)}^g \log(S_{k(i,j)}^g)) \right) \quad (7)$$

We will maximize the entropy constraint item. The loss function for the intra-correlation branch is defined as:

$$\mathcal{L}_g = \lambda_1 \text{Div}(\mathcal{T}^g, \mathcal{S}^g) - \lambda_2 \text{Ent}(\mathcal{S}^g) \quad (8)$$

where λ_1 and λ_2 are used to trade off the loss items. The optimization of intra-correlation \mathcal{S}^g is actually an alternating process with the guidance of target correlation \mathcal{T}^g (see App. A.1), ultimately resulting that each normal patch can establish strong correlations with most normal patches.

3.4. Inter-Correlation Learning

Through the intra-correlation branch, we can establish patch-to-patch correlations within a single image. However, an image usually doesn't contain all possible normal patterns, which may cause it difficult to distinguish some ambiguous abnormal patches (see rows 2, 3, 4 in Figure 1) only through the intra-correlations. To address this, we should effectively take advantage of the known normal patterns from the normal training set, which are more likely to contain more informative normal patterns. Specifically, we further propose an inter-correlation learning branch to explicitly model pairwise correlations with the whole normal training set. In this branch, the features of each patch establish a discrete inter-correlation distribution with the reference features extracted from all normal samples (see Figure 2). The inter-correlation matrix of the k th layer is similar to the corresponding intra-correlation, and is defined as:

$$\text{Inter Correlation} : S_k^e = \text{Softmax} \left(Q_k^e (K_k^e)^T / \sqrt{d_m} \right) \\ Q_k^e, K_k^e, V_k^e = X_{k-1} W_k^{Q_e}, X_f W_k^{K_e}, X_f W_k^{V_e} \quad (9)$$

where $X_f \in \mathbb{R}^{N_e \times d_e}$ represents the external reference features, N_e is the length of the reference features and d_e is the feature dimension. $W_k^{Q_e} \in \mathbb{R}^{d_m \times d_m}$ and $W_k^{K_e}, W_k^{V_e} \in \mathbb{R}^{d_e \times d_m}$ are learnable matrices for Q_k^e, K_k^e, V_k^e . $S_k^e \in \mathbb{R}^{N \times N_e}$ denotes the learned inter-image correlations.

Loss. The loss function for the inter-correlation branch has the opposite optimization direction to \mathcal{L}_g . Because the external reference features contain more comprehensive normal patterns, normal patches can establish stronger correlations with the closest reference normal patterns, instead of establishing strong correlations with most reference patterns. By contrast, it shall be harder for anomalous patches to establish strong correlations with any of the reference

patterns. So the inter-correlation distributions of anomalies are more dispersed, while the normal inter-correlation distributions are more likely to be concentrated. To this end, we maximize the KL divergence and minimize the entropy item in the training process. The practical optimization strategy is also opposite to the intra-correlation branch (see App. A.1 for details). The loss function is defined as:

$$\mathcal{L}_e = -\lambda_1 \text{Div}(\mathcal{T}^e, \mathcal{S}^e) + \lambda_2 \text{Ent}(\mathcal{S}^e) \quad (10)$$

External Reference Features. External reference features are used for providing accumulated knowledge of normality for the inter-correlation learning branch. Thus, these features should represent all possible normal patterns of all normal samples from the whole normal training set. To this end, we can employ many methods to generate the reference features, such as sampling key features by coresets subsampling algorithm [35], generating prototype features by memory module [28], or learning codebook features through vector quantization [65] or sparse coding techniques [55]. However, because the RBF-kernel in \mathcal{T}^e is position-sensitive, our reference features are better to preserve the positional information. The ablation results in App.B.2 show that the methods that can't preserve the position information perform worse. From comprehensive ablation studies, we find that using patch-wise averaged features as the external reference features is a simple but effective way. We think that features extracted by deep neural networks are highly redundant [16, 52], and different normal patterns generally correspond to larger activation values at different channels in the feature vector [61]. So feature averaging will not eliminate some rare normal patterns [11], these patterns may be preserved at specific channels. And averaging can greatly reduce the feature redundancy, making the obtained reference features more representative. Formally, for position (i, j) , we first extract the set of patch features at (i, j) , $X_{ij} = \{x_{ij}^k\}, k \in [1, N]$ from the N normal training images. Then, the reference features at position (i, j) is computed as $x_{ij}^f = \frac{1}{N} \sum_{k=1}^N x_{ij}^k$. The final external reference features are composed of averaged features at all locations and then flattened into 1D: $X_f = \text{Flatten}(\{x_{ij}^f\})$.

3.5. I2Correlation

We further combine the intra-correlation and inter-correlation branches to form the I2Correlation block. The output features of the intra- and inter-correlation branches are defined as: $Z_k^g = S_k^g V_k^g$ and $Z_k^e = S_k^e V_k^e$, respectively. Then, we use the residual feature $Z_k^{ge} = Z_k^g - Z_k^e$ as the output of the I2Correlation block. The feature Z_k^e is generated from the external reference features, which can contain rich normal patterns. Thus, by subtracting Z_k^e from the feature Z_k^g , it is conducive to spotlight the abnormal patterns in the Z_k^{ge} . This is beneficial for anomaly detection.

The total loss function consists of the three branch loss functions, and is combined as follows:

$$\mathcal{L}_{total} = \mathcal{L}_l + \mathcal{L}_g + \mathcal{L}_e \quad (11)$$

3.6. Anomaly Scoring

We utilize reconstruction errors as the baseline anomaly criterion and incorporate the normalized correlation distribution distances into the reconstruction criterion. The final anomaly score of the i th patch in the input patch sequence is shown as follows:

$$s_i = \left[\|X_{i,:} - \hat{X}_{i,:}\|_2 + \left(1 - \frac{X_{i,:} \cdot \hat{X}_{i,:}}{\|X_{i,:}\|_2 \|\hat{X}_{i,:}\|_2}\right) \right] \odot \left(1 - \text{Softmax}(-\text{Div}(\mathcal{T}_{i,:}^e, \mathcal{S}_{i,:}^e))\right) \quad (12)$$

where \odot is the element-wise multiplication.

4. Experiments

4.1. Experimental Setup

Datasets. We extensively evaluate our approach on two widely used industrial AD datasets: the MVTEC3D [3] and BTAD [26], and one recent challenging dataset: the MVTEC3D-3D [5]. MVTEC3D is established as a standard benchmark for evaluating unsupervised image anomaly detection methods. This dataset contains 5354 high-resolution images from 15 real-world categories. 5 classes consist of textures and the other 10 classes contain objects. A total of 73 different anomaly types are presented. BTAD is another popular benchmark for unsupervised image anomaly detection, which contains 2540 RGB images of three industrial products. All classes in this dataset belong to textures. MVTEC3D-3D is recently proposed, which contains 4147 2D RGB images paired with high-resolution 3D point cloud scans from 10 real-world categories. Even though this dataset is mainly designed for 3D anomaly detection, most anomalies can also be detected only through RGB images without 3D point clouds. Since we focus on image anomaly detection, we only use RGB images of the MVTEC3D-3D dataset. We refer to this subset as MVTEC3D-RGB.

Evaluation Metrics. The standard metric in anomaly detection, AUROC, is used to evaluate the performance of AD methods. Image-level AUROC is used for anomaly detection and a pixel-based AUROC for evaluating anomaly localization [3, 4, 11, 64].

Implementation Details. We use EfficientNet-b6 [45] to extract two levels of feature maps with $\{8\times, 16\times\}$ down-sampling ratios, the pre-trained networks are from the timm library [53]. Then, we construct a subsequent transformer network (see Figure 2) at each feature level to reconstruct patch features and learn patch-to-patch correlations. The parameters of the feature extractor are frozen in the training

Table 1. Comparison of our method with the SOTA methods on the MVTEC3D dataset. Red and blue indicate the first and the second best result, respectively. According to the anomaly recognition process, we divide these methods into the patch-wise representation discrepancy, patch-to-patch feature distance, and others.

Discrepancy Type	Method	Venue	Image-level AUROC	Pixel-level AUROC
Patch-wise Representation Discrepancy	STAD [4]	CVPR 2020	0.877	0.939
	PaDiM [11]	ICPR 2020	0.955	0.975
	DFR [57]	Neurocomputing 2021	/	0.950
	FCDD [25]	ICLR 2021	/	0.920
	MKD [41]	CVPR 2021	0.877	0.907
	Hou <i>et al.</i> [17]	ICCV 2021	0.895	/
	Metaformer [54]	ICCV 2021	0.958	/
	DRAEM [64]	ICCV 2021	0.980	0.973
	RDAD [12]	CVPR 2022	0.985	0.978
	SSPCAB [34]	CVPR 2022	0.989	0.972
	DSR [65]	ECCV 2022	0.982	/
	NSA [43]	ECCV 2022	0.972	0.963
	UniAD [62]	NIPS 2022	0.966	0.966
	UTRAD [8]	Neural Networks 2022	0.960	0.967
Patch-to-patch Feature Distance	PatchSVDD [60]	ACCV 2020	0.921	0.957
	DifferNet [36]	WACV 2020	0.949	/
	CFLOW [15]	WACV 2022	0.983	0.986
	CS-FLOW [37]	WACV 2022	0.987	/
	Tsai <i>et al.</i> [15]	WACV 2022	0.981	0.981
	PatchCore [35]	CVPR 2022	0.991	0.980
Others	CutPaste [21]	CVPR 2021	0.952	0.960
	Wang <i>et al.</i> [51]	CVPR 2021	/	0.91
	SPD [68]	ECCV 2022	0.946	0.946
Patch-wise&Intra&Inter	FOD (Ours)	-	0.992	0.983

Table 2. Detailed image-level AUROCs on the MVTEC3D dataset.

Category	Image-level Anomaly Detection					
	DRAEM [64]	PatchSVDD [60]	MKD [41]	PatchCore [35]	CFLOW [15]	FOD (Ours)
Carpet	0.978	0.963	1.000	1.000	0.987	1.000
Grid	1.000	0.892	0.975	0.992	0.996	1.000
Leather	1.000	0.953	0.956	1.000	1.000	1.000
Tile	0.998	0.969	0.999	1.000	0.999	1.000
Wood	0.991	0.989	0.989	0.985	0.991	0.991
Bottle	0.993	0.976	0.989	1.000	1.000	1.000
Cable	0.929	0.899	0.972	0.992	0.976	0.995
Capsule	0.984	0.763	0.979	0.984	0.977	1.000
Hazelnut	1.000	0.912	0.997	1.000	1.000	1.000
Metal nut	0.989	0.941	0.972	1.000	0.993	1.000
Pill	0.981	0.791	0.971	0.954	0.968	0.984
Screw	0.939	0.825	0.870	0.953	0.919	0.967
Toothbrush	1.000	0.992	0.886	0.906	0.997	0.944
Transistor	0.914	0.874	0.956	0.995	0.952	1.000
Zipper	1.000	0.982	0.981	0.989	0.985	0.997
Mean	0.980	0.915	0.966	0.983	0.983	0.992

process, only the parameters of the subsequent transformer networks are learnable. All the subsequent transformer networks in our model contain 3 layers. We set the hidden dimension d_m as $\{256, 512\}$ and the number of heads as 8. The hyperparameters λ_1 and λ_2 are set as 0.5 and 0.5 to trade off two parts of the \mathcal{L}_g and \mathcal{L}_e loss functions (see App. B.2 for hyperparameter sensitivity analysis). We use the Adam [19] optimizer with an initial learning rate of $1e^{-4}$. The total training epochs are set as 100 and the batch size is 1 by default. All the training and test images are resized and cropped to 256×256 resolution from the original resolution.

4.2. Main Results

SOTA Methods. We extensively compare our method with those published SOTA methods in the past three years.

The comparison results on MVTecAD are shown in Table 1. Then, we select five reproducible methods to report the detailed results on MVTecAD, these methods are representative and SOTA AD methods in the mainstream categories of image anomaly detection as we discuss in Related Work, including: DRAEM [64], PatchSVDD[60], PatchCore [35], MKD[41], CFLOW[15]. For a fair comparison, we reproduce all these methods with the same backbone as in our model. Thus, despite using the unmodified code from the official repositories, we are not able to exactly reproduce the original results, but our numbers are very close. The detailed image-level AUROC results are shown in Table 2, and detailed pixel-level AUROC results are in the App. Table 6. Most of the methods in Table 1 don't report results on the BTAD and MVTec3D-RGB datasets, so we reproduce the five representative methods on the two datasets for comparison. The image-level and pixel-level AUROC results on BTAD and MVTec3D-RGB are shown in Table 3. Additional detailed results are in the App. Table 7, 8.

Anomaly Detection. On MVTecAD, we set the SOTA performance on the mean detection AUROC, which is slightly higher than the best competitor, PatchCore [35]. Note that the results in Table 2 show that our method can achieve much better results than PatchCore when using the same backbone. What's more, in addition to the pill, screw and toothbrush classes, our method achieves more than 99% AUROC in all other classes, while other methods only achieve more than 99% AUROC in most nine classes. This shows that our method is more stable and effective in real-world applications. In the classes (*e.g.* metal nut, screw, and transistor) with more global and logical anomalies, our method can achieve significantly better results than those methods depending on patch-wise discrepancy (*e.g.* DRAEM [64] and MKD [41]), and can also achieve better results compared with other methods. This verifies that the three discrepancies are complementary factors and our model can simultaneously spot these discrepancies to recognize harder global and logical anomalies. On BTAD, our FOD can achieve 96.0% mean detection AUROC, which can outperform the best competitor CFLOW [15] by a margin of 1.2%. On MVTec3D-RGB, we can outperform all previous SOTA methods by a margin of 3.3%. Note that this dataset is much more challenging than the MVTecAD dataset when comparing the best results (99.2% for MVTecAD vs. 88.4% AUROC for MVTec3D-RGB). This demonstrates the robustness of our method.

Anomaly Localization. Our method can achieve comparable results with the best competitors on all three datasets. Our method is slightly lower than CFLOW [15] and PatchCore [35] on the three datasets, but we observe that our model that considers the intra- and inter-correlations outperforms the vanilla reconstruction AD models (*e.g.* DRAEM [64] and MKD [41]) on all three

Table 3. Comparison of our method with the SOTA methods for image-level anomaly detection and pixel-level anomaly localization performance on the BTAD and MVTec3D-RGB datasets.

Method	DRAEM [64]	PatchSVDD [60]	MKD [41]	PatchCore [35]	CFLOW [15]	FOD (Ours)
BTAD Dataset						
Image-level AUROC	0.922	0.924	0.935	0.934	0.948	0.960
Pixel-level AUROC	0.942	0.964	0.965	0.976	0.978	0.975
MVTec3D-RGB Dataset						
Image-level AUROC	0.757	0.743	0.688	0.839	0.851	0.884
Pixel-level AUROC	0.974	0.852	0.970	0.977	0.974	0.976

datasets, which verifies the effectiveness of correlation modeling. Compared with patch-wise reconstruction methods, our method is more effective and robust to detect hard global and logical anomalies (see Figure 3).

Qualitative Results. Samples in Figure 1 show qualitative impressions of accurate anomaly localization from our method. It can be found that our approach can achieve the best anomaly score maps by combining the three discrepancies. Additional visualizations are in the App. Figure 6.

4.3. Ablation Study and Model Analysis

Ablation Study. To explain how our model works effectively, we further investigate the effect of the three key designs in our model: recognition views, entropy constraint and reference features. The quantitative results are shown in Table 4, more results can be found in App. Table 9. The entropy constraint is quite effective and necessary in the intra-correlation branch. Specifically, it brings a remarkable 7.7% averaged absolute AUROC promotion, which demonstrates that the entropy constraint strategy is really conducive to increase the distinguishability between abnormal and normal. Only utilizing pure reconstruction criterion or pure KL divergence can't get the most superior detection results, the combined criterion can outperform each single criterion consistently by a margin of 1.0% and 9.8%. Thus the reconstruction errors and the intra- and inter-correlations can collaborate to improve detection performance. For the external reference features, we compare the simple mean features with more elaborate coreset features [35] (see App. Table 15 for more comparison methods). The results show that the simple mean features are more effective than the coreset features (0.923 vs. 0.836), which means that mean features are effective enough to represent all possible normal patterns and preserve the positional information. Finally, our proposed FOD surpasses the pure reconstruction Transformer by 4.0% absolute improvement. These verify that our proposed explicit correlation learning approach is effective.

Effect for Different Anomalies. To illustrate the effectiveness of our model intuitively, we explore the anomaly localization quality under different types of anomalies (see Figure 3). *E.g.*, simple local anomalies (Column 1) are obviously different from normal visuals; hard global anomalies (Columns 2,3) have less obvious visual appearances,

Table 4. Ablation results in recognition views, entropy constraint, external reference features and anomaly scoring. *Patch-wise*, *Intra* and *Inter* mean patch-wise discrepancy, intra- and inter-correlation, respectively. *w/o* and *w/* mean without and with entropy constraint. *Mean* and *Coreset* refer to mean and coreset features [35] as the external reference features (see more details in App. B.3). *Rec*, *Div* and *Rec&Div* mean the pure reconstruction criterion, pure KL divergence (E.q.6) and the combined criterion (E.q.12).

Recognition Views	Entropy Constraint	Reference Features	Anomaly Scoring	MVTecAD	BTAD	MVTec3D-RGB	Avg
Patch-wise	/	/	Rec	0.972	0.954	0.790	0.905
Intra	w/o	/	Div	0.700	0.811	0.708	0.740
	w/	/	Div	0.911	0.822	0.717	0.817
	w/	/	Rec&Div	0.974	0.952	0.818	0.915
Inter	w/	Mean	Rec&Div	0.980	0.958	0.832	0.923
	w/	Coreset	Rec&Div	0.925	0.884	0.700	0.836
Intra+Inter	w/	Mean	Div	0.896	0.922	0.814	0.877
Patch-wise+Intra+Inter (Ours)	w/	Mean	Rec&Div	0.992	0.960	0.884	0.945

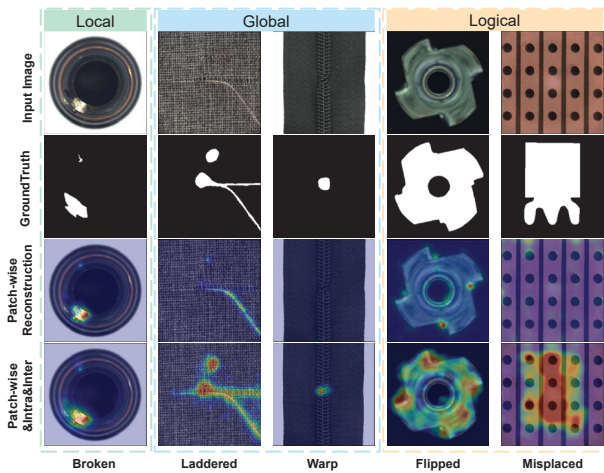


Figure 3. Visualization of different anomalies. The patch-wise reconstruction is only suitable for simple local anomalies, while for hard global and logical anomalies, it cannot detect or only detect partial anomalies. In contrast, our FOD can detect these anomalies more effectively.

which usually need to be effectively compared with normal regions to detect; logical anomalies (Columns 4,5) may be locally normal and can be detected correctly through the overall semantic understanding. We can find that our FOD is more distinguishable in general. For simple local anomalies, patch-wise reconstruction methods can also achieve good results. While for hard global anomalies, these methods cannot detect or only detect partial anomalies. However, for logical anomalies, these methods cannot detect them at all. In contrast, our method is conducive to detect diverse anomalies because of its complementary recognition views. This verifies that our method can make more precise detection and reduce the false-negative rate compared with the pure patch-wise reconstruction methods.

Correlations Visualization. To explain what our model has learned intuitively, we visualize some learned correlations in Figure 4 with the MVTEC-AD dataset. For the intra-




Figure 4. Visualization of intra- and inter-correlations. Each red point in the figure can represent an image patch, we plot its discrete correlation distribution with all the other image patches.

correlation branch, the normal patches learn to build strong correlations with most patches in the whole image, while the correlation distributions of anomalies usually concentrate in the adjacent image patches. It can be found that the normal correlations are much spread and the abnormal correlations are more concentrated in the adjacent regions. For the inter-correlation branch, most normal correlations only concentrate in one point and the abnormal correlations are more spread, which exactly means that each normal patch can build strong correlations with one special normal pattern and abnormal patches are harder to establish correlations with normal patterns.

5. Conclusion

Humans recognize anomalies through two aspects: patch-wise representation discrepancies and weak patch-to-patch correlations. In this paper, we propose a novel AD framework: FOCUS-the-Discrepancy, to simultaneously spot the patch-wise, intra- and inter-discrepancies. The patch-wise discrepancies and intra- and inter-correlations are complementary factors, and can be combined to detect more complex and diverse anomalies. The major characteristic of our method is that we renovate the self-attention maps in transformers to I2Correlation to explicitly establish intra- and inter-image correlations for AD modeling. The combination of explicit correlation learning and transformer architecture can match the core idea of anomaly detection quite well.

Acknowledgements

This work was supported in part by the National Natural Science Fund of China (61971281), the Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), and the Science and Technology Commission of Shanghai Municipality (22DZ2229005).

References

- [1] Samet Akcay, Amir Atapour-Abarghouei, and Toby P. Breckon. Ganomaly: Semi-supervised anomaly detection via adversarial training. *In ACCV*, pages 622–637, 2018. [2](#)
- [2] Liron Bergman, Niv Cohen, and Yedid Hoshen. Deep nearest neighbor anomaly detection. *arXiv preprint arXiv:2002.10445*, 2020. [2](#)
- [3] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad - a comprehensive real-world dataset for unsupervised anomaly detection. *In CVPR*, 2019. [1](#), [6](#)
- [4] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. *In CVPR*, 2020. [2](#), [3](#), [6](#)
- [5] Paul Bergmann, Xin Jin, David Sattlegger, and Carsten Steger. The mvtec 3d-ad dataset for unsupervised 3d anomaly detection and localization. *arXiv preprint arXiv:2112.09045*, 2021. [6](#)
- [6] Paul Bergmann, Sindy Löwe, Michael Fauser, and David Sattlegger. Improving unsupervised defect segmentation by applying structural similarity to autoencoders. *In ICCVTA*, 2019. [2](#)
- [7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. *In ECCV*, 2020. [2](#)
- [8] Liyang Chen, Zhiyuan You, Nian Zhang, Juntong Xi, and Xinyi Le. Utrad: Anomaly detection and localization with u-transformer. *In Neural Networks*, 2022. [2](#), [3](#), [6](#)
- [9] Niv Cohen and Yedid Hoshen. Sub-image anomaly detection with deep pyramid correspondences. *arXiv preprint arXiv:2005.02357v3*, 2020. [2](#)
- [10] Lucas Deecke, Lukas Ruff, Robert A. Vandermeulen, and Hakan Bilen. Transfer-based semantic anomaly detection. *In ICML*, 2021. [2](#)
- [11] Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier. Padim: a patch distribution modeling framework for anomaly detection and localization. *In The 1st International Workshop on Industrial Machine Learning*, pages 475–489, 2021. [1](#), [2](#), [5](#), [6](#)
- [12] Hanqiu Deng and Xingyu Li. Anomaly detection via reverse distillation from one-class embedding. *In CVPR*, 2022. [3](#), [6](#)
- [13] Xiaohan Ding, Xiangyu Zhang, Yizhuang Zhou, Jungong Han, Guiguang Ding, and Jian Sun. Scaling up your kernels to 31x31: Revisiting large kernel design in cnns. *In CVPR*, 2022. [3](#)
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *In ICLR*, 2020. [2](#)
- [15] Denis Gudovskiy, Shun Ishizaka, and Kazuki Kozuka. Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows. *In WACV*, 2022. [3](#), [6](#), [7](#), [12](#), [13](#), [14](#)
- [16] Kai Han, Yunhe Wang, Qi Tian, Jianyuan Guo, Chunjing Xu, and Chang Xu. Ghostnet: More features from cheap operations. *In CVPR*, 2020. [5](#)
- [17] Jinlei Hou, Yingying Zhang, Qiaoyong Zhong, Di Xie, Shiliang Pu, and Hong Zhou. Divide-and-assemble: Learning block-wise memory for unsupervised anomaly detection. *In ICCV*, 2021. [2](#), [6](#)
- [18] Keng Chaia Jonathan Pirnaya. Inpainting transformer for anomaly detection. *arXiv preprint arXiv:2104.13897*, 2021. [2](#), [3](#)
- [19] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *In ICLR*, 2015. [6](#)
- [20] Boyi Li, Kilian Q. Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Anomaly transformer: Time series anomaly detection with association discrepancy. *In ICLR*, 2022. [3](#), [12](#)
- [21] Chun-Liang Li, Kihyuk Sohn, Jinsung Yoon, and Tomas Pfister. Cutpaste: Self-supervised learning for anomaly detection and localization. *In CVPR*, 2021. [6](#)
- [22] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and Jianfeng Gao. Grounded language-image pre-training. *In CVPR*, 2022. [2](#)
- [23] Wenqian Liu, Runze Li, Meng Zheng, Srikrishna Karanam, Ziyang Wu, Bir Bhanu, Richard J. Radke, and Octavia Camps. Towards visually explaining variational autoencoders. *In CVPR*, 2020. [2](#)
- [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 ICCV*, 2021. [2](#)
- [25] Philipp Liznerski, Lukas Ruff, Robert A. Vandermeulen, Billy Joe Franks, Marius Kloft, and Klaus-Robert Müller. Explainable deep one-class classification. *In ICLR*, 2021. [3](#), [6](#)
- [26] Pankaj Mishra, Riccardo Verk, Daniele Fornasier, Claudio Picciarelli, and Gian Luca Foresti. Vt-adl: A vision transformer network for image anomaly detection and localization. *arXiv preprint arXiv:2104.10036*, 2021. [2](#), [3](#), [6](#)
- [27] Radford M. Neal. Pattern recognition and machine learning. *Technometrics*, 2007. [12](#), [13](#)
- [28] Hyunjong Park, Jongyoun Noh, and Bumsub Ham. Learning memory-guided normality for anomaly detection. *In CVPR*, 2020. [1](#), [2](#), [5](#), [14](#), [15](#), [16](#)
- [29] Stanislav Pidhorsky, Ranya Almohsen, Domald A. Adjeroh, and Gianfranco Doretto. Generative probabilities novelty detection with adversarial autoencoders. *In NeurIPS*, 2018. [2](#)
- [30] Alec Radford, JongWook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *In ICML*, 2021. [2](#)
- [31] Tal Reiss, Niv Cohen, Liron Bergman, and Yedid Hoshen. Panda: Adapting pretrained features for anomaly detection and segmentation. *In CVPR*, 2021. [2](#)

- [32] Tal Reiss and Yedid Hoshen. Mean-shifted contrastive loss for anomaly detection. *arXiv preprint arXiv:2106.03844*, 2021. **2**
- [33] Oliver Rippel, Patrick Mertens, and Dorit Merhof. Modeling the distribution of normal data in pre-trained deep features for anomaly detection. *arXiv preprint arXiv: 2005.14140*, 2020. **2**
- [34] Nicolae-Catalin Ristea, Neelu Madan, Radu Tudor Ionescu, Kamal Nasrollahi, Fahad Shahbaz Khan, Thomas B. Moeslund, and Mubarak Shah. Self-supervised predictive convolutional attentive block for anomaly detection. *In CVPR*, 2022. **6**
- [35] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Scholkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. *In CVPR*, 2022. **1, 2, 5, 6, 7, 8, 12, 13, 14, 15, 16**
- [36] Marco Rudolph, Bastian Wandt, and Bodo Rosenhahn. Same same but different: Semi-supervised defect detection with normalizing flows. *In WACV*, 2021. **3, 6**
- [37] Marco Rudolph, Tom Wehrbein, Bodo Rosenhahn, and Bastian Wandt. Fully convolutional cross-scale-flows for image based defect detection. *In WACV*, 2022. **3, 6**
- [38] Marco Rudolph, Tom Wehrbein, Bodo Rosenhahn, and Bastian Wandt. Asymmetric student-teacher networks for industrial anomaly detection. *In WACV*, 2023. **3**
- [39] Lukas Ruff, Nico Görnitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Robert Vandermeulen, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. *In ICML*, pages 4390–4399, 2018. **2, 3**
- [40] Mohammad Sabokrou, Mohammad Khaloori, Mahmood Fathy, and Ehsan Adeli. Adversarially learned one-class classifier for novelty detection. *In CVPR*, 2018. **2**
- [41] Mohammadreza Salehi, Niousha Sadjadi, Soroos Hossein Rohban, and Hamid R.Rabiee. Multiresolution knowledge distillation for anomaly detection. *In CVPR*, 2021. **1, 2, 3, 6, 7, 12, 13, 14**
- [42] Thomas Schlegl, Philipp Seebock, Sebastian M. Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. *In ICIPMI*, pages 146–157, 2017. **2**
- [43] Hannah M. Schluter, Jeremy Tan I, Benjamin Hou, and Bernhard Kainz. Natural synthetic anomalies for self-supervised anomaly detection and localization. *In ECCV*, 2022. **6**
- [44] Bernhard Schölkopf, John C Platt, John Shawe-Taylor, Alex J Smola, and Robert C Williamson. Estimating the support of a high-dimensional distribution. *In Neural Computation 13, 7 (2001)*, pages 1443–1471, 2001. **3**
- [45] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *In ICML*, 2019. **6**
- [46] Xian Tao, Xinyi Gong, Xin Zhang, Shaohua Yan, and Chandranath Adak. Deep learning for unsupervised anomaly localization in industrial images: A survey. *In IEEE Transactions on Instrumentation and Measurement*, 2022. **1**
- [47] David Martinus Johannes Tax and Robert P W Duin. Support vector data description. *In Machine Learning 54, 1 (2004)*, pages 45–66, 2004. **3**
- [48] Chin-Chia Tsai, Tsung-Hsuan Wu, and Shang-Hong Lai. Multi-scale patch-based representation learning for image anomaly detection and segmentation. *In WACV*, 2022. **3**
- [49] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. *In NIPS*, 2018. **15, 16**
- [50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *In NIPS*, 2017. **2, 3**
- [51] Shenzhi Wang, Liwei Wu, Lei Cui, and Yujun Shen. Glancing at the patch: Anomaly localization with global and local feature comparison. *In CVPR*, 2021. **6**
- [52] Xudong Wang and Stella X. Yu. Tied block convolution: Leaner and better cnns with shared thinner filters. *In AAAI*, 2021. **5**
- [53] Ross Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019. **6**
- [54] Jhih-Ciang Wu, Ding-Jie Chen, Chiou-Shann Fuh, and Tyng-Luh Liu. Learning unsupervised metaformer for anomaly detection. *In ICCV*, 2021. **6**
- [55] Jhih-Ciang Wu, He-Yen Hsieh, Ding-Jie Chen, Chiou-Shann Fuh, and Tyng-Luh Liu. Self-supervised sparse representation for video anomaly detection. *In ECCV*, 2022. **5, 14**
- [56] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M. Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *In NIPS*, 2021. **2**
- [57] Jie Yang, Yong Shi, and Zhiqian Qi. Dfr: Deep feature reconstruction for unsupervised anomaly segmentation. *arXiv preprint arXiv: 2012.07122*, 2020. **2, 6**
- [58] Xincheng Yao, Ruoqi Li, Jing Zhang, Jun Sun, and Chongyang Zhang. Explicit boundary guided semi-push-pull contrastive learning for supervised anomaly detection. *In CVPR*, 2023. **1**
- [59] Xincheng Yao, Chongyang Zhang, Ruoqi Li, Jun Sun, and Zhenyu Liu. One-for-all: Proposal masked cross-class anomaly detection. *In AAAI*, 2023. **3**
- [60] Jihun Yi and Sungroh Yoon. Patch svdd: Patch-level svdd for anomaly detection and segmentation. *In ACCV*, 2020. **1, 2, 3, 6, 7, 12, 13, 14**
- [61] Jason Yosinski, Jeff Clune, Anh Nguyen, Thomas Fuchs, and Hod Lipson. Understanding neural networks through deep visualization. *arXiv preprint arXiv:1506.06579*, 2015. **5**
- [62] Zhiyuan You, Lei Cui, Yujun Shen, Kai Yang, Xin Lu, Yu Zheng, and Xinyi Le. A unified model for multi-class anomaly detection. *In NIPS*, 2022. **2, 3, 6**
- [63] Jiawei Yu, Ye Zheng, Xiang Wang, Wei Li, Yushuang Wu, Rui Zhao, and Liwei Wu. Fastflow: Unsupervised anomaly detection and localization via 2d normalizing flows. *arXiv preprint arXiv:2111.07677*, 2021. **3**
- [64] Vitjan Zavrtanik, Matej Kristan, and Danijel Skocaj. Draem: A discriminatively trained reconstruction embedding for surface anomaly detection. *In ICCV*, 2021. **2, 6, 7, 12, 13, 14**
- [65] Vitjan Zavrtanik, Matej Kristan, and Danijel Skocaj. Dsr – a dual subspace re-projection network for surface anomaly detection. *In ECCV*, 2022. **5, 6, 14, 16**

- [66] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. *In ICML*, 2021. [4](#)
- [67] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip H.S. Torr, and Li Zhang. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. *In CVPR*, 2021. [2](#)
- [68] Yang Zou, Jongheon Jeong, Latha Pemula, Dongqing Zhang, and Onkar Dabeer. Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. *In ECCV*, 2022. [6](#)