

CutPaste: Self-Supervised Learning for Anomaly Detection and Localization

Chun-Liang Li*, Kihyuk Sohn*, Jinsung Yoon, Tomas Pfister
Google Cloud AI Research

{chunliang, kihyuks, jinsungyoon, tpfister}@google.com

Abstract

We aim at constructing a high performance model for defect detection that detects unknown anomalous patterns of an image without anomalous data. To this end, we propose a two-stage framework for building anomaly detectors using normal training data only. We first learn self-supervised deep representations and then build a generative one-class classifier on learned representations. We learn representations by classifying normal data from the CutPaste, a simple data augmentation strategy that cuts an image patch and pastes at a random location of a large image. Our empirical study on MVTec anomaly detection dataset demonstrates the proposed algorithm is general to be able to detect various types of real-world defects. We bring the improvement upon previous arts by 3.1 AUCs when learning representations from scratch. By transfer learning on pretrained representations on ImageNet, we achieve a new state-of-the-art 96.6 AUC. Lastly, we extend the framework to learn and extract representations from patches to allow localizing defective areas without annotations during training.

1. Introduction

Anomaly detection aims to detect an instance containing anomalous and defective patterns that are different from those seen in normal instances. Many problems from different vision applications are anomaly detection, including manufacturing defect detection [9, 5], medical image analysis [50, 48], and video surveillance [2, 31, 53]. Unlike a typical supervised classification problem, anomaly detection faces unique challenges. First, due to the nature of the problem, it is difficult to obtain a large amount of anomalous data, either labeled or unlabeled. Second, the difference between normal and anomalous patterns are often *fine-grained* as defective areas might be small and subtle in high-resolution images.

Due to limited access to anomalous data, constructing an anomaly detector is often conducted under semi-supervised or one-class classification settings using normal data only.

Since the distribution of anomaly patterns is unknown in advance, we train models to learn patterns of normal instances and determine anomaly if the test example is not represented well by these models. For example, an autoencoder that is trained to reconstruct normal data is used to declare anomalies when the data reconstruction error is high. Generative models declare anomalies when the probability density is below a certain threshold. However, the anomaly score defined as an aggregation of pixel-wise reconstruction error or probability densities lacks to capture a high-level semantic information [42, 37].

Alternative methods using high-level learned representations have shown more effective for anomaly detection. For example, deep one-class classifier [46] demonstrates an effective end-to-end trained one-class classifiers parameterized by deep neural networks. It outperforms its shallow counterparts, such as one-class SVMs [49] and reconstruction-based approaches such as autoencoders [34]. In self-supervised representation learning, predicting geometric transformations of an image [20, 24, 4], such as rotation or translation, and contrastive learning [54, 52] have shown to be successful in distinguishing normal data from outliers. However, most existing works focus on detecting semantic outliers (e.g., visual objects from different classes) from object-centric natural images. In Section 4.1, we show these methods do not generalize well in detecting fine-grained anomalous patterns as in defect detection.

In this work, we tackle a one-class defect detection problem, a special case of image anomaly detection, where various forms of unknown anomalous patterns present locally in the high-resolution images. We follow the two-stage framework [52], where we first learn self-supervised representations by solving a proxy task, then build a generative one-class classifier on learned representations to distinguish data with anomalous patterns from normal ones. Our innovation is at designing a novel proxy task for self-supervised learning of representations. Specifically, we formulate a proxy classification task between normal training data and the ones augmented by the *CutPaste*, the proposed data augmentation strategy that cuts an image patch and pastes at a random location of an image. *CutPaste* augmentation is mo-

*Equal contributions.

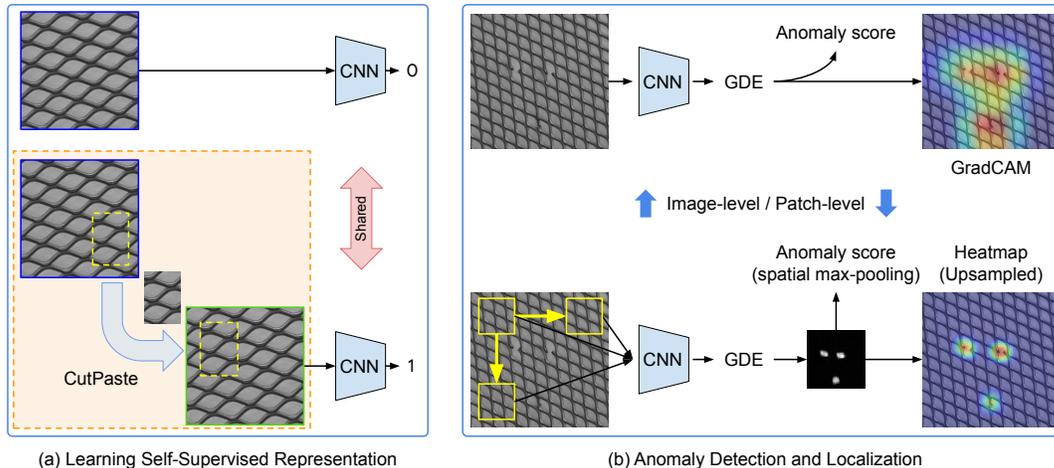


Figure 1: An overview of our method for anomaly detection and localization. (a) A deep network (CNN) is trained to distinguish images from normal (blue) and augmented (green) data distributions by *CutPaste* (orange dotted box), which cuts a small rectangular region (yellow dotted box) from normal data and pastes it at random location. Representations are trained either from the whole image or local patches. (b, top) An image-level representation makes a holistic decision for anomaly detection and is used to localize defect via GradCAM [51]. (b, bottom) A patch-level representation extracts dense features from local patches to produce anomaly score map, which is then max-pooled for detection or upsampled for localization [32].

tivated to produce a spatial *irregularity* to serve as a coarse approximation of real defects, which we have no access at training. Rectangular patches of different sizes, aspect ratios, and rotation angles are pasted to generate diverse augmentations. Although *CutPaste* augmented samples (Figure 2(e)) are easily distinguishable from real defects and thus might be a crude approximation of a real anomaly distribution, we show that representations learned by detecting irregularity introduced by *CutPaste* augmentations generalize well on detecting real defects.

We evaluate our methods on MVTEC anomaly detection dataset [5], a real-world industrial visual inspection benchmark. By learning deep representations from scratch, we achieve **95.2 AUC** on image-level anomaly detection, which outperforms existing works [25, 61] by at least 3.1 AUC. Moreover, we report state-of-the-art **96.6 image-level AUC** by transfer learning from an ImageNet pretrained model. Moreover, we explain how learned representations could be used to localize the defective areas in high-resolution images. Without using any anomaly data, a simple patch model extension can achieve **96.0 pixel-level localization AUC**, which improves upon previous state-of-the-art [61] (95.7 AUC). We conduct an extensive study using different types of augmentation and proxy tasks to show the effectiveness of *CutPaste* augmentations for self-supervised representation learning on unknown defect detection.

2. A Framework for Anomaly Detection

In this section, we present our anomaly detection framework for high-resolution image with defects in local regions. Following [54], we adopt a two-stage framework for

building an anomaly detector, where in the first stage we learn deep representations from normal data and then construct an one-class classifier using learned representations. Subsequently, in Section 2.1, we present a novel method for learning self-supervised representations by predicting *CutPaste* augmentation, and extend to learning and extracting representations from local patches in Section 2.4.

2.1. Self-Supervised Learning with *CutPaste*

Defining good pretext tasks is essential for self-supervised representation learning. While popular methods including rotation prediction [19] and contrastive learning [60, 12] have been studied in the context of semantic one-class classification [20, 24, 4, 54, 52], our study in Section 4.1 shows that naively applying existing methods, such as rotation prediction or contrastive learning, is sub-optimal for detecting local defects as we will show in Section 4.1.

We conjecture that geometric transformations [20, 24, 4], such as rotations and translations, are effective in learning representation of semantic concepts (e.g., objectness), but less of regularity (e.g., continuity, repetition). As shown in Figure 2(b), anomalous patterns of defect detection typically include irregularities such as cracks (bottle, wood) or twists (toothbrush, grid). Our aim is to design an augmentation strategy creating local irregular patterns. Then we train the model to identify these local irregularity with the hope that it can generalize to unseen real defects at test time.

A popular augmentation method that could create a local irregularity in image is *Cutout* [18] (Figure 2(c)), which wipes out a randomly selected small rectangular area of an image. *Cutout* is found to be a useful data augmenta-

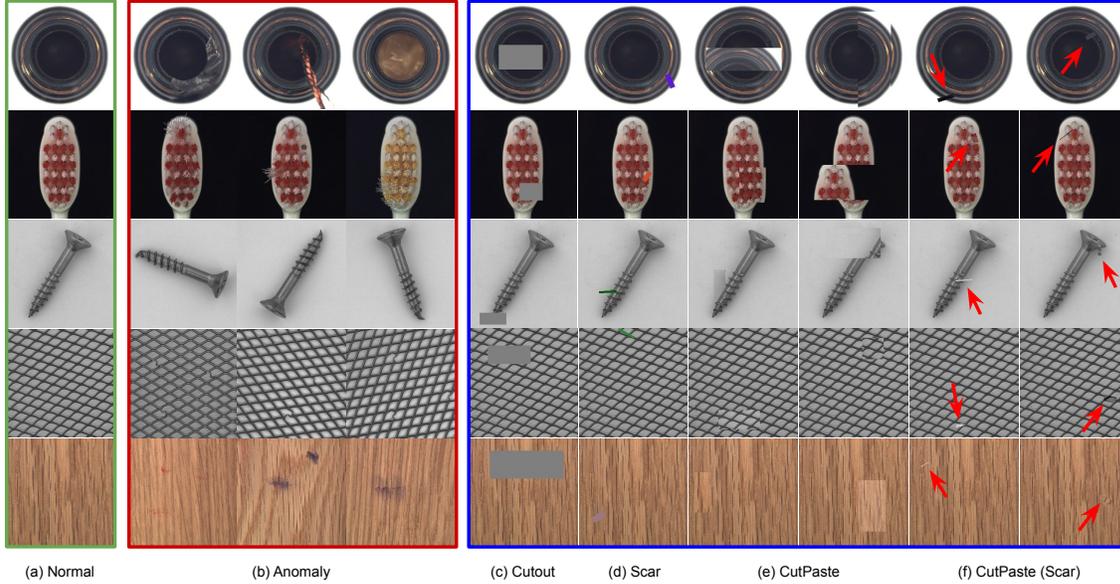


Figure 2: Visualization of (a, green) normal, (b, red) anomaly, and (c–h, blue) augmented normal samples from bottle, toothbrush, screw, grid, and wood classes of MVTEC anomaly detection dataset [5]. Augmented normal samples are generated by baseline augmentations including (c) Cutout and (d) Scar, and our proposed (e) CutPaste and (f) CutPaste (Scar). We use red arrows in (f) to highlight the pasted patch of scar shape, a thin rectangle with rotation.

tion that enforces *invariance*, leading to improved accuracy on multi-class classification tasks. In contrast, we start by *discriminating* Cutout images from the normal ones. At first glance, the task seems easy to solve by well-crafted low level image filters. Surprisingly, as we will show in Section 4, without the hindsight of knowing this, a deep convolution network does not learn these shortcuts. Using Cutout in the algorithm design for defect detection can also be found in [32, 57]. We can make the task harder by randomly choosing colors and the scale as shown in Figure 2(d) to avoid naive shortcut solutions.

To further prevent learning naive decision rules for discriminating augmented images and encouraging the model to learn to detect irregularity, we propose the *CutPaste* augmentation as follows:

1. Cut a small rectangular area of variable sizes and aspect ratios from a normal training image.
2. Optionally, we rotate or jitter pixel values in the patch.
3. Paste a patch back to an image at a random location.

We show the CutPaste augmentation process in the orange dotted box of Figure 1 and more examples in Figure 2(e). Following the idea of rotation prediction [19], we define the training objective of the proposed self-supervised representation learning as follows:

$$\mathcal{L}_{CP} = \mathbb{E}_{x \in \mathcal{X}} \{ \text{CE}(g(x), 0) + \text{CE}(g(\text{CP}(x)), 1) \} \quad (1)$$

where \mathcal{X} is the set of normal data, $\text{CP}(\cdot)$ is a CutPaste augmentation and g is a binary classifier parameterized by deep

networks. $\text{CE}(\cdot, \cdot)$ refers to a cross-entropy loss. In practice, data augmentations, such as translation or color jitter, are applied before feeding x into g or CP.

2.2. CutPaste Variants

CutPaste-Scar. A special case of Cutout called “scar” using a long-thin rectangular box of random color, as in Figure 2(d), is proposed in [16] for defect detection. Similarly, in addition to original CutPaste using a large rectangular patch, we propose a *CutPaste-Scar* using a scar-like (long-thin) rectangular box filled with an image patch (Figure 2(f)).

Multi-Class Classification. While CutPaste (large patch) and CutPaste-Scar share a similarity, the shapes of an image patch of two augmentations are very different. Empirically, they have their own advantages on different types of defects. To leverage the strength of both scales in the training, we formulate a finer-grained 3-way classification task among normal, CutPaste and CutPaste-Scar by treating CutPaste variants as two separate classes. Detailed study will be presented in Section 5.2.

Similarity between CutPaste and real defects. The success of CutPaste may be understood from outlier exposure [23], where we generate the pseudo anomalies (CutPaste) during the training. Apart from using natural images as in [23], CutPaste creates examples preserving more local structures of the normal examples (i.e., the pasted patch is

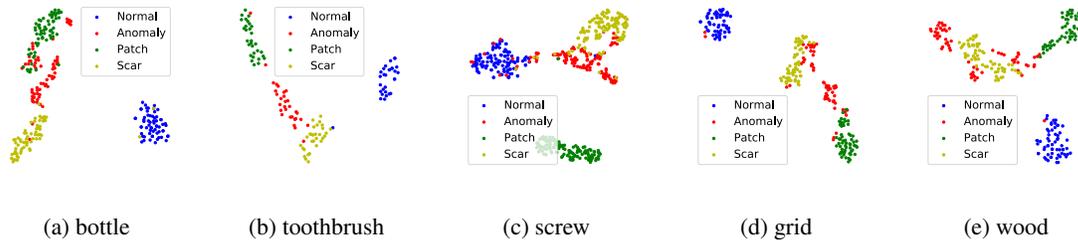


Figure 3: t-SNE visualization of representations of models trained with 3-way CutPaste prediction task. We plot embeddings of normal (blue), anomaly (red), and augmented normal by CutPaste (“Patch”, green) and CutPaste-scar (“Scar”, yellow).

from the same domain), which is more challenging for the model to learn to find this *irregularity*.

On the other hand, CutPaste does look similar to some real defects. A natural question is if the success of CutPaste is from a good mimic of real defects. In Figure 3, we show the t-SNE plots of the representations from the trained model. Clearly, the CutPaste examples are almost not overlapped with real defect examples (anomaly), but the learned representation is able to distinguish between normal example, different CutPaste augmented samples and real defects. It suggests (1) CutPaste is still not a perfect simulation of real defects and (2) learning on it to find irregularity generalizes well on unseen anomalies.

2.3. Computing Anomaly Score

There exist various ways to compute anomaly scores via one-class classifiers. In this work, we build generative classifiers like kernel density estimator [52] or Gaussian density estimator [43], on representations f . Below, we explain how to compute anomaly scores and the trade-offs.

Although nonparametric KDE is free from distribution assumptions, it requires many examples for accurate estimation [58] and could be computationally expensive. With limited normal training examples for defect detection, we consider a simple parametric Gaussian density estimator (GDE) whose log-density is computed as follows:

$$\log p_{\text{gde}}(x) \propto \left\{ -\frac{1}{2} (f(x) - \mu)^\top \Sigma^{-1} (f(x) - \mu) \right\} \quad (2)$$

where μ and Σ are learned from normal training data.¹

2.4. Localization with Patch Representation

While we present a method for learning a holistic representation of an image, learning a representation of an image patch would be preferred if we want to localize defective regions [38, 6, 61] in addition to image-level detection. By learning and extracting representations from an image

¹We note that a mixture of Gaussian, which is a middle ground between KDE and GDE, can also be used for more expressive density modeling. We do not observe significant performance gain empirically.

patch, we can build an anomaly detector that is able to compute the score of an image patch, which then can be used to localize the defective area.

CutPaste prediction is readily applicable to learn a patch representation – all we need to do at training is to crop a patch before applying CutPaste augmentation. Similar to Equation (1), the training objective can be written as:

$$\mathbb{E}_{x \in \mathcal{X}} \left\{ \mathbb{CE}(g(c(x)), 0) + \mathbb{CE}(g(\text{CP}(c(x))), 1) \right\} \quad (3)$$

where $c(x)$ crops a patch at random location of x . At test time, we extract embeddings from all patches with a given stride. For each patch, we evaluate its anomaly score and use a Gaussian smoothing to propagate the score to every pixel [32]. In Section 4.2, we visualize a heatmap using patch-level detector for defect localization, along with that of an image-level detector using visual explanation techniques such as GradCAM [51].

3. Related Work

Anomaly detection under one-class classification setting, where we assume only the normal data is given during the training, has been widely studied [49, 56, 46, 63, 66, 13, 42, 36, 27]. Recent success of self-supervised learning in computer vision [39, 19, 8, 60, 40, 12, 21] has also been demonstrated effective for one-class classification and anomaly detection. One major family is by predicting geometric transformations [20, 24, 4], such as rotation, translation or flips. The other family includes variants of contrastive learning with geometric augmentations [54, 52]. However, the success has been limited to semantic anomaly detection benchmarks, such as CIFAR-10 [28] or ImageNet [17], and as we show in Section 4.1, methods relying on geometric transformations perform poorly on defect detection benchmarks.

Because of practical applications, such as industrial inspection or medical diagnosis, defect detection [9, 5] has received lots of attention. The initial steps have been taken with methods including autoencoding [9, 7, 25, 59], generative adversarial networks [48, 3], using pretrained models on ImageNet [38, 45, 6, 14, 43, 44], and self-supervised learning by solving different proxy tasks with augmentations [61, 47, 57, 15]. The proposed CutPaste prediction

task is not only shown to have strong performance on defect detection, but also amenable to combine with existing methods, such as transfer learning from pretrained models for better performance or patch-based models for more accurate localization, which we demonstrate in Section 4.

3.1. Relation to Other Augmentations

Although Cutout [18] and RandomErasing [65] are similar to CutPaste, they create irregularities by a small rectangular region filled with either zero or uniformly sampled pixel values instead of a structural image patch as CutPaste. Moreover, unlike typical use of augmentations for learning invariant representations, we learn a representation that is *discriminative* to these augmentations.

Scar augmentation [16] (Figure 2(d)) is a special case of Cutout, which uses a long-thin rectangle with random colors. While it demonstrates strong performance, we show that CutPaste with the same scale (Figure 2(f)), which fills a long-thin rectangle by a patch from the same image, improves upon representations trained by predicting Cutout.

CutMix [62], which extracts a rectangular image patch from an image and pastes at random location of another image, is related to CutPaste in terms of pasting operations. One main difference is CutMix leverages existing image labels with MixUp [64] in the objective while CutPaste prediction is a self-supervised learning without the need of image labels. The other difference is CutMix studies standard supervised tasks, while we aim for one-class classification.

[11] presents a denoising autoencoder with patch-swap augmentation as noise process. [26] proposes to learn representations by predicting local augmentations using GAN. Our method is simpler (e.g., no need to train decoder or GAN) while highly performant, thus more practical.

4. Experiments

We conduct most experiments on MVTec Anomaly Detection dataset [5] that contains 10 object and 5 texture categories for anomaly detection. The dataset is composed of normal images for training and both normal and anomaly images with various types of defect for testing. It also provides pixel-level annotations for defective test images. The dataset is relatively small scale in number of images, where the number of training images varies from 60 to 391, posing a unique challenge for learning deep representations.

We follow one-class classification protocol, also known as semi-supervised anomaly detection [10],² where we train a one-class classifier for each category on its respective normal training examples. Following [52], we learn the representations by augmentation prediction from scratch with

²While previous works [5, 6] have used *unsupervised* to describe their settings, it could be misleading as training data is curated to include normal data only.

ResNet-18 [22] plus an MLP projection head on top of average pooling layer followed by the last linear layer. We construct a Gaussian density estimation (GDE) as Equation (2) for anomaly detector based on the top pooled features.

We train a model on 256×256 image. We note that the same training strategy, such as the selection of hyperparameters or data augmentations, is applied to all categories. Detailed settings of training can be found in Appendix A.

4.1. Main Results

We report the anomaly detection performance in Table 1. We run experiments 5 times with different random seeds and report the mean AUC and standard error for each category. We also report the average of mean and standard errors for texture, object, and all categories.

We test representations trained with different proxy tasks of self-supervised learning, including baselines such as rotation [20], Cutout or scar predictions, the proposed CutPaste, CutPaste-Scar predictions, and using both with 3-way classification. We also compare with previous works, including deep one-class classifier (DOCC) [45], uninformed student [5], and patch SVDD [61]. We note that some of these methods use ImageNet pretrained model for transfer learning, either by fine-tuning (DOCC) or distillation (uninformed student). The results are in Table 1.

Rotation prediction is demonstrated to be powerful in semantic anomaly detection [52]. However, it results in unsatisfactory 73.1 AUC in defect detection compared with the Scar prediction (85.0), a Cutout variant. Some failure of rotation prediction is due to the unaligned objects, such as `screw` shown in Figure 2. For aligned objects, although it performs well on `toothbrush`, it is sub-optimal on `capsule`. Detailed ablation study of Cutout variants can be found in Section 5.

CutPaste and CutPaste-Scar, which improve Cutout and Scar prediction by avoiding potential naive solutions, outperform other augmentation predictions with 90.9 and 93.5 AUCs, respectively. With a finer-grained 3-way classification to leverage different scale of CutPaste, we achieve the best 95.2 AUC, which surpasses existing works on learning from scratch, such as P-SVDD [61] (92.1 AUC). The proposed data-driven approach via CutPaste is also better than existing works leveraging pretrained networks, including DOCC [45] (87.9 AUC) with pretrained VGG16 and Uninformed Student [6] (92.5 AUC) with pretrained ResNet18. Last, we further improve the AUC to **96.1** by ensembling anomaly scores from 5 CutPaste (3-way) models.

4.2. Defect Localization

We conduct anomaly localization experiments using our representations trained with 3-way classification task. One challenge to accurate localization of defect is that it is difficult to use a heatmap-style approach for localization as our

Table 1: Anomaly detection performance on MVTec AD dataset [5]. We report AUCs of representations trained to classify CutPaste, CutPaste (scar), both (3-way), and baseline augmentations such as rotation, Cutout, or scar. For comparison, we report those of deep one-class classifier [45], uninformed student [6] and patch-SVDD [61]. We report mean and standard error tested with 5 random seeds. Lastly, we report the AUC using ensemble of 5 CutPaste (3-way) models. The best performing model and those within standard error are bold-faced.

Category		DOCC [45]	U-Student [6]	P-SVDD [61]	Rotation	Cutout	Scar	CutPaste	CutPaste (scar)	CutPaste (3-way)	Ensemble
texture	carpet	90.6	95.3	92.9	29.7±1.4	35.3±2.3	92.7±0.4	67.9±1.8	94.6±0.6	93.1±1.1	93.9
	grid	52.4	98.7	94.6	60.5±7.0	57.5±3.0	74.4±2.5	99.9 ±0.1	95.5±0.3	99.9 ±0.1	100.0
	leather	78.3	93.4	90.9	55.2±1.4	67.7±1.5	99.9 ±0.1	99.7±0.1	100.0 ±0.0	100.0 ±0.0	100.0
	tile	96.5	95.8	97.8	70.1±1.9	71.8±4.0	96.7 ±0.9	95.9 ±1.0	89.4±2.8	93.4±1.0	94.6
	wood	91.6	95.5	96.5	95.8±1.1	92.0±0.8	98.9 ±0.2	94.9±0.5	98.7 ±0.3	98.6 ±0.5	99.1
	average	81.9	95.7	94.5	62.3±2.6	64.9±2.3	92.5±0.8	91.7±0.7	95.7±0.8	97.0 ±0.5	97.5
object	bottle	99.6	96.7	98.6	95.0±0.7	88.7±0.8	98.5±0.2	99.2±0.2	98.0±0.5	98.3±0.5	98.2
	cable	90.9	82.3	90.3	85.3±0.8	80.2±1.4	78.3±1.7	87.1±0.8	78.8±2.9	80.6±0.5	81.2
	capsule	91.0	92.8	76.7	71.8±1.8	69.5±1.1	82.9±0.7	87.9±0.7	95.3 ±0.8	96.2 ±0.5	98.2
	hazelnut	95.0	91.4	92.0	83.6±0.8	69.7±1.3	98.9 ±0.2	91.3±0.6	96.7±0.4	97.3±0.3	98.3
	metal nut	85.2	94.0	94.0	72.7±0.5	84.6±0.7	86.9±1.5	96.8±0.5	97.9±0.2	99.3 ±0.2	99.9
	pill	80.4	86.7	86.1	79.2±1.4	78.7±0.7	82.2±1.4	93.4 ±0.9	85.8±1.3	92.4 ±1.3	94.9
	screw	86.9	87.4	81.3	35.8±2.9	17.6±4.2	11.3±2.2	54.4±1.7	83.7±0.7	86.3 ±1.0	88.7
	toothbrush	96.4	98.6	100.0	99.1±0.2	98.1±0.6	94.8±1.0	99.2±0.2	96.7±0.4	98.3±0.9	99.4
	transistor	90.8	83.6	91.5	88.9±0.4	82.5±1.2	92.0±0.7	96.4 ±0.7	91.1±0.6	95.5 ±0.5	96.1
	zipper	92.4	95.8	97.9	74.3±1.6	75.7±1.0	86.8±0.9	99.4 ±0.1	99.5 ±0.1	99.4 ±0.2	99.9
average	90.9	90.9	90.8	78.6±1.1	74.5±1.3	81.3±1.1	90.5±0.6	92.4±0.8	94.3 ±0.6	95.5	
average	87.9	92.5	92.1	73.1±1.6	71.3±1.6	85.0±1.0	90.9±0.7	93.5±0.8	95.2 ±0.6	96.1	

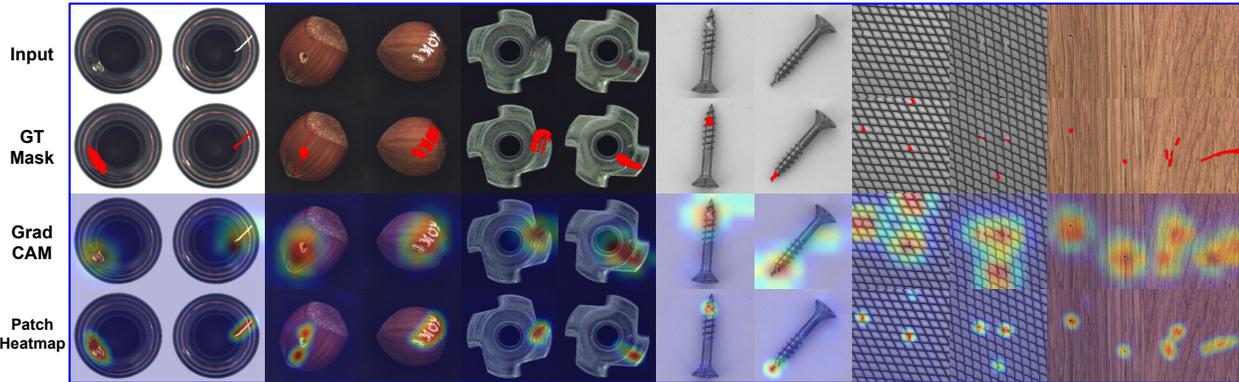


Figure 4: Defect localization on bottle, hazelnut, metal nut, screw, wood and grid classes of MVTec datasets. From top to bottom, input images, those with ground-truth localization mask in red, GradCAM results using image-level detector, and heatmaps using patch-level detector. We provide more examples in Appendix B.

model learns a holistic representation of an image. Instead, we use visual explanation techniques, GradCAM [51], to highlight the area affecting the decision of anomaly detector. We show qualitative results in the second row of Figure 4, which are visually pleasing. We further evaluate the pixel-wise localization AUC, achieving 88.3.

Instead, we learn a representation of an image patch using CutPaste prediction, as in Section 2.4. We train models of 64×64 patches from 256×256 image. At test time, we densely extract anomaly scores with a stride of 4 and propagate the anomaly scores via receptive field upsampling with Gaussian smoothing [32]. We report a localization AUC in Table 2. Our patch-based model achieves 96.0

AUC. Specifically, our model shows strong performance on texture categories over previous state-of-the-art (96.3 AUC compared to 93.7). We also outperforms the DistAug contrastive learning [52], which only results in 90.4 localization AUC. Finally, we visualize representative samples for localization in Figure 4, showing accurate localization even when defects are tiny. More comprehensive results on defect localization are given in Appendix B.

4.3. Transfer Learning with Pretrained Models

In Section 4.1, we have shown the proposed data-driven approach is better than leveraging pretrained networks, such as DOCC [45] and Uninformed Student [6]. It is consistent

Table 2: Pixel-wise localization AUC on MVTEc dataset. The best and models within standard error are bold-faced.

Category		FCDD [32]	P-SVDD [61]	CutPaste (3-way)
texture	carpet	96	92.6	98.3 \pm 0.0
	grid	91	96.2	97.5 \pm 0.1
	leather	98	97.4	99.5 \pm 0.0
	tile	91	91.4	90.5 \pm 0.2
	wood	88	90.8	95.5 \pm 0.1
	average	93	93.7	96.3 \pm 0.1
object	bottle	97	98.1	97.6 \pm 0.1
	cable	90	96.8	90.0 \pm 0.2
	capsule	93	95.8	97.4 \pm 0.1
	hazelnut	95	97.5	97.3 \pm 0.1
	metal nut	94	98.0	93.1 \pm 0.4
	pill	81	95.1	95.7 \pm 0.1
	screw	86	95.7	96.7 \pm 0.1
	toothbrush	94	98.1	98.1 \pm 0.0
	transistor	88	97.0	93.0 \pm 0.2
	zipper	92	95.1	99.3 \pm 0.0
	average	91	96.7	95.8 \pm 0.1
average	92	95.7	96.0 \pm 0.1	

Table 3: Detection performance on MVTEc dataset using representations of EfficientNet (B4) [55] pretrained on ImageNet [17] and fine-tuned by the CutPaste (3-way). The number is bold when it is better than its pretrained or fine-tuned counterpart under the same feature (pool v.s. level-7).

Category		Pool		Level-7	
		Pretrain	Finetune	Pretrain	Finetune
texture	carpet	98.3	100.0 \pm 0.0	97.6	100.0 \pm 0.0
	grid	96.4	98.8 \pm 0.1	98.2	99.1 \pm 0.0
	leather	100.0	100.0 \pm 0.0	100.0	100.0 \pm 0.0
	tile	99.9	98.9 \pm 0.2	99.9	99.8 \pm 0.2
	wood	99.7	99.8 \pm 0.0	99.6	99.8 \pm 0.0
	average	98.9	99.5 \pm 0.0	99	99.7 \pm 0.0
object	bottle	99.8	100.0 \pm 0.0	100.0	100.0 \pm 0.0
	cable	91.2	93.9 \pm 0.1	96.5	96.2 \pm 0.3
	capsule	93	94.3 \pm 0.3	94.7	95.4 \pm 0.1
	hazelnut	96.6	99.7 \pm 0.0	100.0	99.9 \pm 0.0
	metalnut	94.3	98.7 \pm 0.1	97.7	98.6 \pm 0.0
	pill	81.9	91.3 \pm 0.2	91	93.3 \pm 0.2
	screw	86.3	86 \pm 0.1	92.0	86.6 \pm 0.2
	toothbrush	89.3	92.8 \pm 0.2	90.3	90.7 \pm 0.1
	transistor	94.6	95.6 \pm 0.2	97.5	97.5 \pm 0.2
	zipper	95.6	99.9 \pm 0.0	97	99.9 \pm 0.1
	average	92.3	95.2 \pm 0.1	95.7	95.8 \pm 0.1
average	94.5	96.6 \pm 0.1	96.8	97.1 \pm 0.0	

with the prior study on semantic anomaly detection [52]. On the other hand, pretrained EfficientNet [55] is found useful for defect detection [43]. As shown in Table 3, without fine-tuning, the representation from the pretrained EfficientNet (B4) results in 94.5 AUC, which is competitive with the proposed CutPaste prediction (95.2 from Table 1).

Here we demonstrate that the proposed self-supervised learning via CutPaste is versatile, which can also be used to improve the pretrained networks to better adapt to the data. We use pretrained EfficientNet (B4) as a backbone, and follow the standard fine-tuning steps to train with the same

Table 4: Detection AUCs of representations trained to predict Cutout, with mean pixel value, with random color, Confetti noise [32], or the proposed CutPaste.

Category	Cutout	Cutout (Mean)	Cutout (Color)	Confetti	CutPaste
texture	64.9 \pm 2.3	65.5 \pm 1.8	70.5 \pm 2.2	80.1 \pm 2.3	91.7 \pm 0.7
object	74.5 \pm 1.3	78.1 \pm 1.1	78.9 \pm 1.0	76.7 \pm 0.3	90.5 \pm 0.6
all	71.3 \pm 1.7	73.9 \pm 1.3	76.1 \pm 1.4	77.8 \pm 1.5	90.9 \pm 0.7

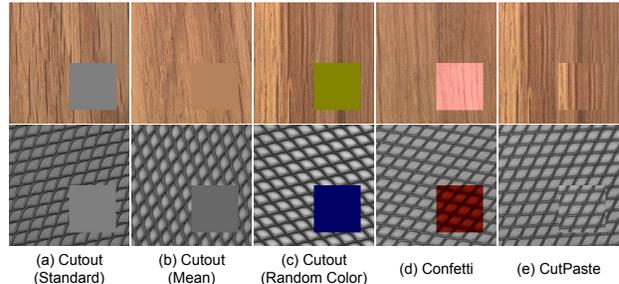


Figure 5: Visual comparison between the proposed CutPaste and Cutout variants, including filling with grey color, mean pixel values, random colors and Confetti noise [32].

CutPaste prediction (3-way) task. Detailed settings can be found in Appendix A. We show the results in Table 3. After fine-tuning via CutPaste, we achieve the **new state-of-the-art 96.6 AUC**. Furthermore, CutPaste prediction is a general and useful strategy to adapt to the data for most of the situations. For example, CutPaste improves by a large margin on class *pill* (81.9 \rightarrow 91.3). For many nearly perfect situations, such as *bottle*, CutPaste is still able to improve by a small margin. Last, as suggested by [30, 43], we investigate the performance of various deep features. We find that level-7 feature shows the best performance, and we further improve the level-7 feature of EfficientNet from 96.8 (pretrained) to **97.1 \pm 0.0** with CutPaste.

5. Ablation Study

We conduct various additional studies to provide deeper insights of the proposed CutPaste. We first compare CutPaste with different Cutout variants in addition to the standard ones reported in Section 4.1. Second, we showcase the representation learned via predicting CutPaste generalizes well to more crafted unseen defects. Last, we compare with the semantic anomaly detection.

5.1. From Cutout to CutPaste

We evaluate the performance of representations trained to predict variants of Cutout augmentations whose areas are filled by grey color (standard), mean pixel values, random color, or image patch from different location, i.e., CutPaste. We also test Confetti noise [32] that jitters a color of a local patch. We show samples from considered augmentations in Figure 5 and report the detection AUCs in Table 4. While

Table 5: Detection AUCs of representations trained with binary classification between normal and the union of CutPaste and CutPaste-scar examples and 3-way classification among normal, CutPaste and CutPaste-scar examples.

Category	CutPaste	CutPaste (scar)	Binary (Union)	3-Way
texture	91.7±0.7	95.7±0.8	97.3±0.3	97.0±0.5
object	90.5±0.6	92.4±0.8	92.8±0.5	94.3±0.6
all	90.9±0.7	93.5±0.8	94.3±0.5	95.2±0.6

achieving 71.3 AUC that already is significantly better than random guessing, predicting a standard Cutout augmentation is still a simple task and the network may have learned a naive solution from easy proxy task, as discussed in Section 2. By gradually increasing the difficulty of proxy task to avoid known trivial solutions with random color to the patch, or with structures similar to local patterns of the normal data (Confetti noise, CutPaste), the network learns to find irregularity and generalizes better to detect real defects.

5.2. Binary v.s. Finer-Grained Classification

In Table 1, although CutPaste-scar shows better performance on average than CutPaste, there is no clear winner that works the best for all. As there are diverse types of defect in practice, we leverage the strength of both augmentations for representation learning. In Section 2.2, we train a model by solving a 3-way classification task between normal, CutPaste and CutPaste-scar. Alternatively, we train to solve a binary classification task by discriminating normal examples and the union of two augmentations.

The results, along with those of representations trained with CutPaste and CutPaste-scar, are in Table 5. It is clear that using both augmentations improve the performance. Between binary with union of augmentations and 3-way, we observe better detection performance with representations trained by 3-way classification task. A plausible hypothesis on the superiority of 3-way formulation in our case is that it is more natural to model CutPaste and CutPaste-scar augmentations separately than together as there exists a systematic difference between them in the size, shape, and rotation angle of patches.

5.3. CutPaste on Synthetic Anomaly Detection

We further study the generalization of our models to unseen anomalies. Specifically, we test on synthetic anomaly datasets created by patching diverse shape masks to normal data, such as digits [29], square, ellipse, or heart [35], filled with random color or natural images. Samples of synthetic anomalies are shown in Figure 6 and detection results are in Table 6. We first note that these datasets are not trivial – a model trained by predicting Cutout augmentations achieves only 81.5. Our proposed CutPaste (3-way) model performs well on synthetic dataset, achieving 98.3 AUC on average. We highlight that some shapes (e.g., ellipse, heart) or color



Figure 6: Synthetic defects on pill class. From left to right, we use MNIST [29], square, ellipse, heart [35] with random color, and those filled with natural image patches.

Table 6: Detection AUCs on synthetic data. Various shapes, such as digit, square, ellipse, or heart, are patched to normal images with random color or natural images ([†]).

Dataset	MNIST	Square	Ellipse	Heart	Square [†]	Ellipse [†]	Heart [†]	Avg
Cutout	52.3	90.6	89.3	87.5	86.4	84.0	80.7	81.5
CutPaste	96.1	98.4	98.2	97.9	99.3	99.2	99.0	98.3

statistics inside the patch (e.g., constant color, natural images) are not seen at training, but we can still generalize to these unseen cases.

5.4. Application to Semantic Outlier Detection

We also conduct the semantic anomaly detection experiment on CIFAR-10 [28] following the protocol in [20, 52], where a single class is treated as normal and remaining 9 classes are anomalies. We make a comparison of Cutout, CutPaste and rotation prediction [52]. Cutout results in 60.2 AUC, and CutPaste achieves 69.4 AUC, which significantly improves upon Cutout (60.2). However, these are still far behind that of rotation prediction (91.3 AUC) on CIFAR-10 semantic anomaly detection. On the other hand, in Section 4.1, we have discussed the reversed situation that rotation prediction is much worse than 3-way CutPaste prediction. The results suggest the difference between semantic anomaly detection and defect detection, which needs different algorithm and augmentation designs.

6. Conclusion

We propose a data-driven approach for defect detection and localization. The key to our success is self-supervised learning of representations with CutPaste, a simple yet effective augmentation that encourages the model to find local irregularities. We show superior image-level anomaly detection performance on the real-world dataset. Furthermore, by learning and extracting patch-level representations, we demonstrate state-of-the-art pixel-wise anomaly localization performance. We envision the CutPaste augmentation could be a cornerstone for building a powerful model for semi-supervised and unsupervised defect detection.

Acknowledgment. We thank Yang Feng for sharing the implementation of uninformed student and Sercan Arik for the proofread of our manuscript.

References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *{USENIX}*, 2016. 12
- [2] Amit Adam, Ehud Rivlin, Ilan Shimshoni, and Daviv Reinitz. Robust real-time unusual event detection using multiple fixed-location monitors. *IEEE transactions on pattern analysis and machine intelligence*, 2008. 1
- [3] Samet Akcay, Amir Atapour-Abarghouei, and Toby P Breckon. Ganomaly: Semi-supervised anomaly detection via adversarial training. In *ACCV*, 2018. 4
- [4] Liron Bergman and Yedid Hoshen. Classification-based anomaly detection for general data. In *ICLR*, 2020. 1, 2, 4
- [5] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad—a comprehensive real-world dataset for unsupervised anomaly detection. In *CVPR*, 2019. 1, 2, 3, 4, 5, 6
- [6] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In *CVPR*, 2020. 4, 5, 6
- [7] Paul Bergmann, Sindy Löwe, Michael Fauser, David Sattlegger, and Carsten Steger. Improving unsupervised defect segmentation by applying structural similarity to autoencoders. *arXiv preprint arXiv:1807.02011*, 2018. 4
- [8] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018. 4
- [9] Diego Carrera, Giacomo Boracchi, Alessandro Foi, and Brendt Wohlberg. Detecting anomalous structures by convolutional sparse models. In *IJCNN*, 2015. 1, 4
- [10] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM computing surveys (CSUR)*, 2009. 5
- [11] Liang Chen, Paul Bentley, Kensaku Mori, Kazunari Misawa, Michitaka Fujiwara, and Daniel Rueckert. Self-supervised learning for medical image analysis using image context restoration. *Medical image analysis*, 2019. 5
- [12] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. 2, 4
- [13] Hyunsun Choi, Eric Jang, and Alexander A Alemi. Waic, but why? generative ensembles for robust anomaly detection. *arXiv preprint arXiv:1810.01392*, 2018. 4
- [14] Niv Cohen and Yedid Hoshen. Sub-image anomaly detection with deep pyramid correspondences. *arXiv preprint arXiv:2005.02357*, 2020. 4
- [15] Anne-Sophie Collin and Christophe De Vleeschouwer. Improved anomaly detection by training an autoencoder with skip connections on images corrupted with stain-shaped noise. *arXiv preprint arXiv:2008.12977*, 2020. 4
- [16] daisukelab. Spotting defects! — deep metric learning solution for mvtec anomaly detection dataset. <https://medium.com/analytics-vidhya/spotting-defects-deep-metric-learning-solution-for-mvtec-anomaly-detection-dataset-c77691beb1eb>. Accessed: 2020-11-12. 3, 5
- [17] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 4, 7
- [18] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 2, 5
- [19] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Un-supervised representation learning by predicting image rotations. In *ICLR*, 2018. 2, 3, 4
- [20] Izhak Golan and Ran El-Yaniv. Deep anomaly detection using geometric transformations. In *NIPS*, 2018. 1, 2, 4, 5, 8
- [21] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 4
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 5
- [23] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:1812.04606*, 2018. 3
- [24] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. In *NIPS*, 2019. 1, 2, 4
- [25] Chaoqing Huang, Jinkun Cao, Fei Ye, Maosen Li, Ya Zhang, and Cewu Lu. Inverse-transform autoencoder for anomaly detection. *arXiv preprint arXiv:1911.10676*, 2019. 2, 4
- [26] Simon Jenni, Hailin Jin, and Paolo Favaro. Steering self-supervised feature learning beyond local pixel statistics. In *CVPR*, 2020. 5
- [27] Polina Kirichenko, Pavel Izmailov, and Andrew Gordon Wilson. Why normalizing flows fail to detect out-of-distribution data. *arXiv preprint arXiv:2006.08545*, 2020. 4
- [28] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 4, 8
- [29] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998. 8
- [30] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *NIPS*, 2018. 7
- [31] Yusha Liu, Chun-Liang Li, and Barnabás Póczos. Classifier two sample test for video anomaly detections. In *BMVC*, 2018. 1
- [32] Philipp Liznerski, Lukas Ruff, Robert A Vandermeulen, Billy Joe Franks, Marius Kloft, and Klaus-Robert Müller. Explainable deep one-class classification. *arXiv preprint arXiv:2007.01760*, 2020. 2, 3, 4, 6, 7, 13
- [33] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *ICLR*, 2017. 12

- [34] Jonathan Masci, Ueli Meier, Dan Cireşan, and Jürgen Schmidhuber. Stacked convolutional auto-encoders for hierarchical feature extraction. In *ICANN*, 2011. 1
- [35] Loic Matthey, Irina Higgins, Demis Hassabis, and Alexander Lerchner. dsprites: Disentanglement testing sprites dataset. URL <https://github.com/deepmind/dsprites-dataset/>. [Accessed on: 2018-05-08], 2017. 8
- [36] Warren R Morningstar, Cusuh Ham, Andrew G Gallagher, Balaji Lakshminarayanan, Alexander A Alemi, and Joshua V Dillon. Density of states estimation for out-of-distribution detection. *arXiv preprint arXiv:2006.09273*, 2020. 4
- [37] Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, and Balaji Lakshminarayanan. Detecting out-of-distribution inputs to deep generative models using a test for typicality. *arXiv preprint arXiv:1906.02994*, 2019. 1
- [38] Paolo Napoletano, Flavio Piccoli, and Raimondo Schettini. Anomaly detection in nanofibrous materials by cnn-based self-similarity. *Sensors*, 2018. 4
- [39] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. 4
- [40] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 4
- [41] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 2011. 12
- [42] Jie Ren, Peter J Liu, Emily Fertig, Jasper Snoek, Ryan Poplin, Mark Depristo, Joshua Dillon, and Balaji Lakshminarayanan. Likelihood ratios for out-of-distribution detection. In *NIPS*, 2019. 1, 4
- [43] 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. 4, 7
- [44] Marco Rudolph, Bastian Wandt, and Bodo Rosenhahn. Same same but different: Semi-supervised defect detection with normalizing flows. *arXiv preprint arXiv:2008.12577*, 2020. 4
- [45] Lukas Ruff, Jacob R Kauffmann, Robert A Vandermeulen, Grégoire Montavon, Wojciech Samek, Marius Kloft, Thomas G Dietterich, and Klaus-Robert Müller. A unifying review of deep and shallow anomaly detection. *arXiv preprint arXiv:2009.11732*, 2020. 4, 5, 6
- [46] Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *ICML*, 2018. 1, 4
- [47] Mohammadreza Salehi, Ainaz Eftekhari, Niousha Sadjadi, Mohammad Hossein Rohban, and Hamid R Rabiee. Puzzle-ae: Novelty detection in images through solving puzzles. *arXiv preprint arXiv:2008.12959*, 2020. 4
- [48] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *IPMI*, 2017. 1, 4
- [49] Bernhard Schölkopf, Robert C Williamson, Alex J Smola, John Shawe-Taylor, and John C Platt. Support vector method for novelty detection. In *NIPS*, 2000. 1, 4
- [50] Philipp Seeböck, Sebastian Waldstein, Sophie Klimesch, Bianca S Gerendas, René Donner, Thomas Schlegl, Ursula Schmidt-Erfurth, and Georg Langs. Identifying and categorizing anomalies in retinal imaging data. *arXiv preprint arXiv:1612.00686*, 2016. 1
- [51] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, 2017. 2, 4, 6
- [52] Kihyuk Sohn, Chun-Liang Li, Jinsung Yoon, Minh Jin, and Tomas Pfister. Learning and evaluating representations for deep one-class classification. In *ICLR*, 2021. 1, 2, 4, 5, 6, 7, 8
- [53] Waqas Sultani, Chen Chen, and Mubarak Shah. Real-world anomaly detection in surveillance videos. In *CVPR*, 2018. 1
- [54] Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin. Csi: Novelty detection via contrastive learning on distributionally shifted instances. *arXiv preprint arXiv:2007.08176*, 2020. 1, 2, 4
- [55] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019. 7
- [56] David MJ Tax and Robert PW Duin. Support vector data description. *Machine learning*, 2004. 4
- [57] Tareq Tayeh, Sulaiman Aburakhia, Ryan Myers, and Abdallah Shami. Distance-based anomaly detection for industrial surfaces using triplet networks. *arXiv preprint arXiv:2011.04121*, 2020. 3, 4
- [58] Alexandre B Tsybakov. *Introduction to nonparametric estimation*. Springer Science & Business Media, 2008. 4
- [59] Shashanka Venkataramanan, Kuan-Chuan Peng, Rajat Vikram Singh, and Abhijit Mahalanobis. Attention guided anomaly localization in images. *arXiv preprint arXiv:1911.08616*, 2019. 4
- [60] Mang Ye, Xu Zhang, Pong C Yuen, and Shih-Fu Chang. Unsupervised embedding learning via invariant and spreading instance feature. In *CVPR*, 2019. 2, 4
- [61] Jihun Yi and Sungroh Yoon. Patch svdd: Patch-level svdd for anomaly detection and segmentation. *arXiv preprint arXiv:2006.16067*, 2020. 2, 4, 5, 6, 7
- [62] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. 5
- [63] Shuangfei Zhai, Yu Cheng, Weining Lu, and Zhongfei Zhang. Deep structured energy based models for anomaly detection. *arXiv preprint arXiv:1605.07717*, 2016. 4
- [64] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. MixUp: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 5
- [65] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *AAAI*, 2020. 5, 12

- [66] Bo Zong, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Daeki Cho, and Haifeng Chen. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In *ICLR*, 2018. 4