

# What Makes a Good Data Augmentation for Few-Shot Unsupervised Image Anomaly Detection?

Lingrui Zhang\*  
SUSTech

lingrui\_zhang@foxmail.com

Shuheng Zhang\*  
SUSTech

zhangsh2001@foxmail.com

Guoyang Xie  
SUSTech and University of Surrey

guoyang.xie@surrey.ac.uk

Jiaqi Liu  
SUSTech

liujq32021@mail.sustech.edu.cn

Hua Yan  
University of Warwick

Aaron.Yan.1@warwick.ac.uk

Jinbao Wang<sup>†</sup>  
SUSTech

linking@163.com

Feng Zheng<sup>†</sup>  
SUSTech

zfeng02@gmail.com

Yaochu Jin  
Bielefeld University and University of Surrey

yaochu.jin@uni-bielefeld.de

## Abstract

*Data augmentation is a promising technique for unsupervised anomaly detection in industrial applications, where the availability of positive samples is often limited due to factors such as commercial competition and sample collection difficulties. In this paper, how to effectively select and apply data augmentation methods for unsupervised anomaly detection is studied. The impact of various data augmentation methods on different anomaly detection algorithms is systematically investigated through experiments. The experimental results show that the performance of different industrial image anomaly detection (termed as IAD) algorithms is not significantly affected by the specific data augmentation method employed and that combining multiple data augmentation methods does not necessarily yield further improvements in the accuracy of anomaly detection, although it can achieve excellent results on specific methods. These findings provide useful guidance on selecting appropriate data augmentation methods for different requirements in IAD.*

## 1. Introduction

Industrial Image Anomaly Detection (IAD) is a challenging task that aims to identify defects or abnormalities in industrial images. Unlike natural images, industrial images have high similarity among normal samples, which makes it easy to model their distribution. However, the number

and variety of anomalies are very low and unpredictable, which makes it hard to collect enough labeled data for supervised learning. Therefore, IAD requires unsupervised or weakly supervised methods that can learn from normal samples only and detect anomalies based on their deviation from the normal distribution.

There are two main types of anomaly detection, feature embedding-based methods [2–4, 7, 8, 12, 15, 18–20, 26, 27, 30, 31], and reconstruction-based methods [5, 11, 21, 32–34]. Feature embedding-based methods use pre-trained models to extract high-level features from industrial images. And then clustering or density estimation techniques are applied to measure the distance or probability of each image feature from the normal cluster or distribution. Images with large distances or low probability are considered to be anomalies. Reconstruction-based methods use generative models, such as variational autoencoders or generative adversarial networks, to learn a mapping function from industrial images to a latent space and vice versa. Then they reconstruct each image from its latent representation and compute the reconstruction error between the original and reconstructed images. Images with large reconstruction errors are considered to be anomalies.

Few-shot IAD is an emerging research topic [9, 17, 18, 23] that aims to detect defects in industrial images with only a few normal samples. This scenario is common in real-world industrial applications, where collecting a large number of normal samples may be impractical or costly. For instance, in the debugging phase of a production line, only a few normal samples can be produced before the yield reaches a satisfactory level. Moreover, some industrial domains may face commercial competition or privacy issues that prevent

<sup>1</sup>Contributed equally.

<sup>2</sup>Corresponding authors.

them from sharing large amounts of training data.

Data augmentation methods have been used in unsupervised few-shot anomaly detection [9, 18, 22, 23] to increase the diversity and robustness of normal samples. However, there is a lack of systematic and comprehensive analysis of how different data augmentation methods affect the performance of IAD algorithms. Previous works have used various data augmentation methods with different settings, but without providing in-depth and detailed explanations or comparisons. For example, MAEDAY [22] only uses rotation, while [23] aggressively uses multiple augmentations. Only RegAD [9] is one of the few works that has studied the impact of data augmentation methods on few-shot anomaly detection in more detail, using four data augmentation methods on two IAD datasets. In this paper, we conduct a thorough and extensive study of the role of data augmentation methods. We apply 6 separate data augments to 3 IAD datasets and test them on 11 IAD algorithms. For the number of training images, we choose 1, 2, 4, and 8 respectively. In total, we conduct thousands of experiments and systematically summarize a large number of meaningful conclusions. The results of our experiments show that there is no single data augmentation method that can consistently improve the performance of all unsupervised IAD algorithms. However, we subdivide the 11 IAD methods into multiple classes, and the analysis shows that data augmentation methods tend to have similar effects on similar IAD algorithms. In addition, we explore the effects of combining different data augmentations to increase the diversity and complexity of the training data, and the experimental results show that the benefits of mixed data augmentation depend heavily on the IAD method. For most methods, we do not observe significant improvements from using multiple data augmentations. However, in some cases, we find that mixed data augmentation can boost accuracy remarkably. The main contribution of this work can be summarized as follows:

- We examine 6 data augmentation methods, and 11 image anomaly detection methods on 3 benchmark datasets, resulting in a total of 6,688 instances. In addition, we present a plug-and-play and modular implementation for data augmentation of few-shot IAD evaluation, which greatly benefits the future development of few-shot IAD evaluation.
- Our research highlights that the optimal choice of data augmentation method is dependent on the specific methodology employed. Interestingly, we find that comparable IAD algorithms generally respond similarly to various data augmentation methods.
- Our findings indicate that the effectiveness of combining multiple data augmentation methods substantially

depends on the specific IAD method employed. Additionally, we identify the IAD methods that are compatible with mixed data augmentation.

## 2. Related Works

### 2.1. Unsupervised IAD

There are many kinds of anomaly detection methods, and they can be divided into two main types, embedding-based [2, 3, 18] and reconstruction-based [28, 32]. The embedding-based methods are further divided into four classes, i.e., Normalizing Flow [18], Teacher-Student [2], One-Class Classification [24] and Memory Bank [3]. Next, they will be presented in detail.

#### 2.1.1 Normalizing Flow (NF)

Normalizing Flows is a generative model that can produce easy-to-handle distributions, where sampling and density assessment can be efficient and accurate [18]. CS-Flow [19] used the feature block  $x$  to input the flow model to fit the distribution  $z$  of the product of Gaussian distribution and Dirac distribution. The network is updated by the likelihood probability method of this distribution, and the evaluation index is made. FastFlow [31] is implemented by two-dimensional normalized flow and used as a probability distribution estimator, which effectively mapped image features to easily processed base distributions and pays attention to the relationship between local and global features. However, the features of normal images in industrial manufacturing may not conform to the Gaussian distribution.

#### 2.1.2 Teacher-Student

It is a kind of transfer learning. That is to say, the performance of one model is transferred to another model. The teacher network is a complex network with very good performance and generalization ability. This network can be used to guide the student network to learn so that a simpler student model with less parameter computation can also have similar performance to the teacher network, which is also a way of model compression. STPM [26] is multiplied by three different resolutions of the anomaly map to detect the anomalies. RD4AD [6] proposed a reverse knowledge distillation scheme. It took the teacher as encoder, and the student as decoder, and added a one-class bottleneck embedding module between them. But, These methods rely too much on the teacher network, which may weaken the generalization ability.

#### 2.1.3 One-Class Classification (OCC)

Since we only need to distinguish between abnormal and normal, we can turn the problem into OCC. CutPaste [14]

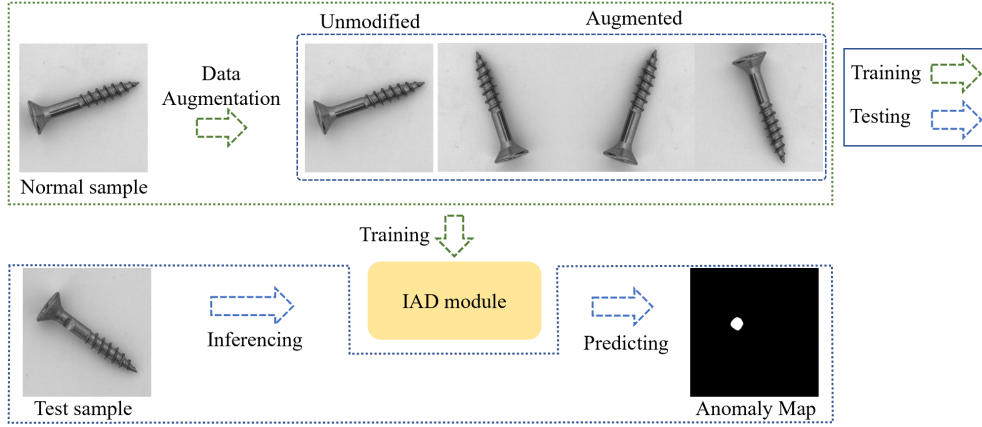


Figure 1. The overall framework of data augmentation for few-shot IAD. In training phase, the normal image and its augmented images are fed into IAD model for training. In inference phase, the test image are classified as anomalies if at least one patch is anomalous and the pixel-level anomaly segmentation is generated by the prediction result of IAD model.

Shot	Methods	Vanilla	Rotation	Flip	Scale	Translate	Color.Jitter	Perspective
1	CFA	0.934	0.933	0.93	0.905	0.902	0.937	0.927
	CSFlow	0.678	0.763	0.737	0.679	0.72	0.735	0.615
	CutPaste	0.762	0.893	0.869	0.611	0.808	0.831	0.861
	DRAEM	0.665	0.648	0.707	0.574	0.672	0.687	0.759
	FastFlow	0.5	0.868	0.74	0.797	0.782	0.787	0.754
	FAVAE	0.713	0.712	0.776	0.662	0.583	0.605	0.716
	PaDiM	0.781	0.879	0.878	0.853	0.8	0.811	0.849
	PatchCore	0.906	0.927	0.914	0.92	0.913	0.91	0.904
	RD4AD	0.844	0.913	0.918	0.879	0.908	0.887	0.886
	SPADE	-	-	-	-	-	-	-
STPM	0.876	0.89	0.915	0.912	0.888	0.92	0.908	
2	CFA	0.928	0.936	0.932	0.916	0.919	0.925	0.928
	CSFlow	0.8	0.878	0.76	0.884	0.853	0.833	0.724
	CutPaste	0.718	0.855	0.89	0.883	0.582	0.782	0.604
	DRAEM	0.619	0.752	0.772	0.724	0.532	0.749	0.779
	FastFlow	0.644	0.887	0.811	0.844	0.88	0.829	0.751
	FAVAE	0.796	0.798	0.842	0.636	0.66	0.65	0.589
	PaDiM	0.875	0.903	0.907	0.878	0.886	0.881	0.878
	PatchCore	0.911	0.913	0.91	0.909	0.916	0.895	0.92
	RD4AD	0.885	0.919	0.91	0.911	0.916	0.914	0.919
	SPADE	-	0.882	0.887	0.872	0.872	0.831	0.884
STPM	0.907	0.92	0.926	0.924	0.923	0.934	0.923	
4	CFA	0.924	0.925	0.926	0.885	0.917	0.92	0.909
	CSFlow	0.742	0.814	0.745	0.89	0.756	0.78	0.665
	CutPaste	0.778	0.896	0.891	0.775	0.81	0.753	0.608
	DRAEM	0.684	0.857	0.667	0.714	0.632	0.7	0.753
	FastFlow	0.8	0.815	0.764	0.824	0.826	0.772	0.749
	FAVAE	0.85	0.714	0.72	0.712	0.711	0.513	0.748
	PaDiM	0.904	0.914	0.923	0.911	0.918	0.907	0.912
	PatchCore	0.921	0.919	0.919	0.924	0.919	0.907	0.904
	RD4AD	0.91	0.922	0.927	0.925	0.908	0.911	0.925
	SPADE	-	0.871	0.87	0.873	0.871	0.84	0.869
STPM	0.924	0.907	0.927	0.923	0.923	0.932	0.914	
8	CFA	0.932	0.937	0.924	0.886	0.925	0.923	0.923
	CSFlow	0.892	0.884	0.895	0.924	0.932	0.943	0.877
	CutPaste	0.876	0.883	0.904	0.835	0.864	0.705	0.755
	DRAEM	0.788	0.859	0.816	0.812	0.761	0.749	0.71
	FastFlow	0.895	0.882	0.869	0.885	0.86	0.789	0.896
	FAVAE	0.701	0.819	0.813	0.644	0.606	0.688	0.756
	PaDiM	0.927	0.936	0.936	0.926	0.939	0.924	0.931
	PatchCore	0.917	0.924	0.923	0.907	0.912	0.904	0.92
	RD4AD	0.921	0.935	0.93	0.932	0.924	0.924	0.931
	SPADE	0.893	0.896	0.9	0.893	0.895	0.867	0.899
STPM	0.94	0.897	0.934	0.923	0.912	0.928	0.919	

Table 1(a). Image-level AUC-ROC on BTAD. The red one is the best augmentation for the IAD method, and the blue is the second best one.

Shot	Methods	Vanilla	Rotation	Flip	Scale	Translate	Color.Jitter	Perspective
1	CFA	0.811	0.829	0.811	0.788	0.802	0.814	0.806
	CSFlow	0.708	0.727	0.708	0.742	0.75	0.7	0.713
	CutPaste	0.65	0.701	0.702	0.68	0.679	0.652	0.703
	DRAEM	0.683	0.718	0.715	0.714	0.69	0.687	0.741
	FastFlow	0.527	0.618	0.613	0.694	0.682	0.578	0.6
	FAVAE	0.651	0.56	0.591	0.6	0.581	0.588	0.626
	PaDiM	0.684	0.697	0.683	0.669	0.683	0.681	0.674
	PatchCore	0.788	0.805	0.792	0.788	0.8	0.797	0.789
	RD4AD	0.77	0.805	0.802	0.799	0.823	0.816	0.784
	SPADE	-	-	-	-	-	-	-
STPM	0.799	0.841	0.814	0.823	0.843	0.831	0.84	
2	CFA	0.839	0.86	0.853	0.795	0.825	0.833	0.843
	CSFlow	0.745	0.781	0.773	0.783	0.768	0.754	0.778
	CutPaste	0.697	0.709	0.748	0.659	0.659	0.611	0.726
	DRAEM	0.78	0.765	0.774	0.751	0.771	0.773	0.784
	FastFlow	0.555	0.743	0.674	0.753	0.731	0.628	0.617
	FAVAE	0.667	0.648	0.659	0.595	0.62	0.642	0.631
	PaDiM	0.708	0.734	0.731	0.694	0.716	0.71	0.714
	PatchCore	0.795	0.831	0.805	0.796	0.806	0.793	0.802
	RD4AD	0.798	0.832	0.816	0.835	0.817	0.835	0.82
	SPADE	-	0.737	0.731	0.71	0.734	0.733	0.744
STPM	0.84	0.839	0.85	0.848	0.851	0.849	0.838	
4	CFA	0.879	0.891	0.861	0.818	0.86	0.895	0.883
	CSFlow	0.785	0.841	0.815	0.825	0.81	0.806	0.738
	CutPaste	0.728	0.771	0.714	0.519	0.674	0.621	0.671
	DRAEM	0.82	0.798	0.802	0.819	0.824	0.823	0.826
	FastFlow	0.693	0.741	0.745	0.79	0.782	0.723	0.734
	FAVAE	0.655	0.626	0.669	0.605	0.639	0.67	0.623
	PaDiM	0.722	0.734	0.723	0.689	0.72	0.727	0.716
	PatchCore	0.844	0.872	0.826	0.844	0.848	0.852	0.847
	RD4AD	0.838	0.897	0.871	0.866	0.876	0.872	0.861
	SPADE	-	0.764	0.739	0.749	0.758	0.757	0.759
STPM	0.864	0.869	0.865	0.876	0.875	0.868	0.885	
8	CFA	0.923	0.913	0.888	0.875	0.92	0.919	0.924
	CSFlow	0.856	0.9	0.863	0.9	0.893	0.892	0.861
	CutPaste	0.705	0.808	0.857	0.473	0.694	0.607	0.673
	DRAEM	0.872	0.892	0.905	0.892	0.9	0.904	0.919
	FastFlow	0.828	0.809	0.793	0.843	0.826	0.809	0.81
	FAVAE	0.701	0.651	0.67	0.619	0.611	0.68	0.668
	PaDiM	0.776	0.81	0.739	0.734	0.761	0.779	0.771
	PatchCore	0.891	0.916	0.876	0.889	0.898	0.883	0.891
	RD4AD	0.903	0.933	0.911	0.917	0.921	0.912	0.916
	SPADE	0.781	0.794	0.77	0.783	0.791	0.788	0.787
STPM	0.865	0.896	0.87	0.91	0.914	0.912	0.904	

Table 1(b). Image-level AUC-ROC on MVTec AD. The red one is the best augmentation for the IAD method, and the blue is the second best one.

Shot	Methods	Vanilla	Rotation	Flip	Scale	Translate	Color,Jitter	Perspective
1	CFA	0.721	0.76	0.779	0.75	0.729	0.752	0.744
	CSFlow	0.494	0.454	0.491	0.515	0.528	0.473	0.566
	CutPaste	0.651	0.719	0.732	0.601	0.627	0.677	0.606
	DRAEM	0.522	0.505	0.523	0.525	0.522	0.512	0.507
	FastFlow	0.5	0.507	0.533	0.559	0.566	0.56	0.488
	FAVAE	0.512	0.479	0.508	0.535	0.502	0.57	0.502
	PaDiM	0.518	0.677	0.676	0.598	0.615	0.604	0.616
	PatchCore	0.664	0.581	0.682	0.67	0.648	0.635	0.67
	RD4AD	0.791	0.77	0.834	0.722	0.749	0.716	0.746
	SPADE	-	-	-	-	-	-	-
STPM	0.561	0.589	0.585	0.706	0.642	0.773	0.662	
2	CFA	0.729	0.755	0.793	0.744	0.745	0.759	0.748
	CSFlow	0.513	0.469	0.515	0.524	0.548	0.514	0.547
	CutPaste	0.58	0.51	0.757	0.657	0.696	0.651	0.66
	DRAEM	0.517	0.528	0.53	0.503	0.499	0.499	0.555
	FastFlow	0.613	0.52	0.529	0.58	0.545	0.638	0.543
	FAVAE	0.51	0.528	0.51	0.512	0.505	0.569	0.518
	PaDiM	0.63	0.706	0.709	0.63	0.658	0.665	0.646
	PatchCore	0.632	0.644	0.649	0.671	0.679	0.67	0.624
	RD4AD	0.574	0.768	0.845	0.761	0.765	0.752	0.764
	SPADE	-	0.687	0.695	0.698	0.696	0.717	0.678
STPM	0.62	0.658	0.645	0.749	0.743	0.756	0.715	
4	CFA	0.804	0.783	0.804	0.717	0.779	0.803	0.767
	CSFlow	0.504	0.43	0.508	0.552	0.522	0.496	0.559
	CutPaste	0.609	0.768	0.738	0.641	0.449	0.639	0.795
	DRAEM	0.537	0.588	0.554	0.553	0.546	0.522	0.62
	FastFlow	0.504	0.584	0.763	0.499	0.658	0.621	0.553
	FAVAE	0.536	0.547	0.557	0.565	0.525	0.559	0.521
	PaDiM	0.718	0.755	0.751	0.708	0.73	0.744	0.713
	PatchCore	0.676	0.649	0.694	0.698	0.677	0.668	0.68
	RD4AD	0.815	0.8	0.828	0.78	0.808	0.805	0.794
	SPADE	-	0.686	0.689	0.695	0.696	0.706	0.652
STPM	0.691	0.697	0.647	0.751	0.748	0.765	0.745	
8	CFA	0.821	0.803	0.826	0.756	0.78	0.822	0.832
	CSFlow	0.678	0.695	0.721	0.72	0.758	0.749	0.675
	CutPaste	0.747	0.787	0.734	0.545	0.594	0.552	0.492
	DRAEM	0.637	0.715	0.698	0.559	0.635	0.664	0.726
	FastFlow	0.544	0.587	0.75	0.512	0.579	0.665	0.565
	FAVAE	0.566	0.556	0.547	0.499	0.527	0.557	0.541
	PaDiM	0.766	0.779	0.782	0.735	0.759	0.779	0.752
	PatchCore	0.711	0.666	0.693	0.662	0.729	0.69	0.71
	RD4AD	0.825	0.806	0.824	0.798	0.819	0.797	0.809
	SPADE	0.705	0.72	0.726	0.706	0.723	0.72	0.713
STPM	0.732	0.619	0.733	0.769	0.747	0.718	0.727	

Table 1(c). Image-level AUC-ROC on MTD. The red one is the best augmentation for the IAD method, and the blue is the second best one.

proposed a high-performance image defect anomaly detection model that can detect anomaly patterns without relying on abnormal data. The whole framework can be described in two-stage. Firstly, the representation of normal images is learned by a self-supervised learning method. Then, a single classifier is constructed based on the learned image representation. It mainly constructs negative samples by cutting images and then pasting them to other locations.

### 2.1.4 Memory Bank

This is a simple and effective method. In the training phase, the feature extraction of the normal image is stored in the memory bank. In the test phase, the distance between the feature of the test sample and the nearest neighbor feature is calculated to achieve anomaly detection. CFA [13] proposed a new method to obtain discriminative features through metric learning. PaDiM [4] proposed an anomaly detection and localization framework based on distributed

modeling in the one-class learning environment. In order to improve the effectiveness of features, SPADE [3] extracted multi-resolution features in pyramid architecture. PatchCore [17] extended the work of a series of unsupervised anomaly detection algorithms such as SPADE, and PaDiM. It mainly solved the problem of the slow testing speed of SPADE and did some exploration in the feature extraction part.

### 2.1.5 Reconstruction-Based Method

This method is mainly to reconstruct the abnormal image into a normal image and compared it with the original abnormal image to find the abnormal part. The method proposed by DRAEM [32] provided an arbitrary number of abnormal samples and a pixel-perfect abnormal segmentation map, which could be used to train the proposed method without real abnormal samples. This method learned the joint representation of abnormal images and their non-abnormal reconstruction and learned the decision bounds of normal and abnormal samples. This method could directly locate anomalies without additional complex post-processing of network output and can be trained using simple and general anomaly simulations. Unlike DRAEM, FAVAE [5] did not rely on external data sets to generate anomalies to further enhance the generalization of the model.

## 2.2. Few-Shot IAD

Few-shot anomaly detection (FSAD) emerged not long ago. The number of samples is small, and the amount of data is not much. For how to use the data to get a better effect, RegAD [9] proposed a comparison-based solution, which is very different from the popular methods based on reconstruction or single classification. And [18] also used 16 images to train its model. Its experimental results on the challenging and newly proposed MVTEC AD datasets and Magnetic Tile Defects datasets showed that their method outperforms the existing methods. However, no one has systematically explored the field of Few-Shot IAD.

## 2.3. Data Augmentation

Data augmentation plays a great role in few-shot training. It can help the model to learn more features from data to get a better effect. According to [29], image augmentation methods based on image erasure usually delete one or more sub-regions in the image to replace the pixel values of these sub-regions with constant or random values. And image mixed data augmentation is mainly done by mixing multiple images or sub-regions of the image into one. But, these data augmentation methods cannot be used for anomaly detection of industrial products.

### 3. Benchmark Setup

Figure 1 illustrates our data augmentation process for training the anomaly detection model. We apply different data augmentation to one image to generate multiple augmented images. These images are used to train the model to learn the normal appearance and features of the data. Then, we can feed a new image to the trained model and obtain two outputs: a confidence score that indicates how likely the image is anomalous, and a mask that highlights the anomalous regions in the image.

#### 3.1. Dataset

We use a number of datasets to investigate the effects of different methods and models on few-shot, including MVtec AD [1], MTD [10] and BTAD [16].

**MVtec AD** [1] contains five textures and ten objects in different fields. It has 4096 normal images and 1258 abnormal images with resolutions varying from  $700 \times 700$  to  $1024 \times 1024$ .

**MTD** [10] has 952 normal images and 392 abnormal images. It contains 6 types of anomalies and simulates authenticity.

**BTAD** [16] contains 2250 normal images and 580 abnormal images. It is a real-world dataset used for anomaly detection, and it has three industrial products with surface flaws.

#### 3.2. Baseline Methods

We conduct experiments on a total of eleven methods, which are divided into two categories: feature-embedding-based and reconstruction-based methods. As mentioned earlier, feature-embedding-based methods are further divided into four classes. CSFlow [19] and FastFlow [31] are methods based on normalizing flow, and CFA [13], SPADE [3], PaDiM [4], and PatchCore [17] all use the features of good samples saved by memory bank. In addition, both RD4AD [6] and STPM [26] use the teacher-student structure of knowledge distillation. And CutPaste [14] is a method of one-class classification. In reconstruction-based methods, we selected FAVAE [5] and DRAEM [32] for our experiments. For comparison purposes, these methods are used as baselines to test the effectiveness of our proposed few-shot IAD.

#### 3.3. Our Data Augmentation

It is challenging to select a suitable data augmentation method for unsupervised IAD. Since the training samples are all normal, we cannot use data augmentation to introduce anomalies or make the samples more similar to the abnormal ones. Therefore, some common data augmentation methods such as adding noise, mix-up [35], etc. are not applicable. In this paper, we experiment with six data aug-

mentation methods that can potentially enhance the performance of unsupervised IAD. They are rotation, flip, scale, translate, color jitter, and perspective.

**Rotation.** Rotate the image evenly within  $[0, 360)$  degrees. For example, to augment an image to 4 images, we would rotate it by 0, 90, 180, and 240 degrees.

**Flip.** For an image, the transformation is performed in the order of keeping the original image, flipping horizontally, flipping vertically, and flipping both horizontally and vertically. If we need to generate more pictures, we set the probability of horizontal and vertical flips to 0.5, respectively.

**Scale.** The side length of the original image is reduced, and the value is uniformly set at  $(0.5, 1]$  times the side length of the original image. Enlarging the image may cause the object to be detected beyond the screen, which will cause adverse effects. So only scale the image down a little bit.

**Translate.** Translate the content of the image horizontally and vertically by  $[0, 0.5)$  times the side length of the image. The distance is uniformly set.

**Color Jitter.** Adjust the brightness and contrast of the image. Except for the original image, the adjustment range is uniformly set at  $[0.5, 1.5)$  times, and the original image is always kept at the first image for the use of subsequent mixed data augmentation.

**Perspective.** It simulates the perspective change of human eyes looking at pictures. The change range is uniform in  $[0, 0.5)$ .

In this paper, for each augmentation, we augment one training image to four. In other words, in addition to keeping the unmodified image, there are three corresponding training images augmented with this image.

### 4. Experimental Results and Analysis

In this section, we aim to conduct a comprehensive and systematic analysis of the impact of data augmentation on few-shot IAD. We explore several research questions, such as: How does data augmentation affect the performance of different types of IAD algorithms? Can data augmentation methods that achieve better results individually also enhance the accuracy when combined? Do all data augmentation methods exhibit consistent effects across different settings and scenarios? To answer these questions, we use the 6 data augments described in Section 3.3, the 3 datasets introduced in Section 3.1, and the 11 IAD algorithms mentioned in Section 3.2, and thousands of experiments have been conducted.

**Key Takeaways** (i) There is no universal winner among data augmentation methods for all IAD methods. (ii) We summarize the best single data augmentation for each category of IAD methods in Table 3. (iii) As for Memory Bank-based methods, we observe that the mixed data augmenta-

tion significantly outperforms the single data augmentation for image-level metrics.




Figure 2. An image from the MvTec AD training set, as well as rotated and flipped images. If you look at the whole image, the data-augmented images are not consistent with the test set. However, if we only look at the patches one by one, the patches after data augmentation can still correspond to the normal toothbrush.

#### 4.1. Single Data Augmentation

As shown in Table 1(a), 1(b), 1(c), a large amount of data is obtained through the experiments. The best and second-best data augmentation for this IAD algorithm in this dataset is marked in red and blue, respectively.

Ideally, we would expect to see a column of the table dominated by either blue or red cells, indicating that the corresponding augmentation method consistently outperforms others. However, this is not the case in our results. Instead, we observe that different augmentation methods have varying effects on different IAD methods. This suggests that there is no single augmentation method that can be universally optimal for all IAD methods. One possible reason for this phenomenon is that different IAD methods have distinct technical designs and mechanisms, which may interact differently with different types of data transformations. Therefore, we further divide the IAD methods into several categories based on their main characteristics and investigate whether similar methods share similar preferences or sensitivities to certain augmentation methods.

As we discussed in Section 3.2, we classify the 11 IAD methods that we evaluate in our experiments into five categories: Normalizing Flow, Memory Bank, Student-Teacher, One-Class, and Reconstruction. These categories are based on the main techniques or principles that each method employs, as shown in Table 2. For each IAD method and each data augmentation method, we calculate their average improvement of image-level AUC-ROC over all datasets and all shot numbers compared to the baseline of no data augmentation. We observe that within each category of IAD methods, there is some degree of consistency in how data

augmentation affects their performance. For example, for Normalizing Flow methods, all data augmentation methods except Perspective can improve their results significantly. For Memory Bank methods, although the overall improvement is modest, Rotation and Flip can consistently boost all four methods in this category. For Student-Teacher methods, all six data augmentation methods have positive effects, especially Flip and Color Jitter. However, for One-Class methods, the impacts of different data augmentation methods vary widely across different models and datasets. Rotation and Flip seem to work well for most models in this category, but other data augmentation methods may have detrimental effects. The only exception is the Reconstruction-based category, where FAVAE [5] suffers from a decrease in accuracy with any data augmentation method applied while DRAEM [32] benefits from Perspective and Rotation.

#### 4.2. Mixed Data Augmentation

We have shown that different data augmentation methods have consistent effects on the performance of IAD methods within the same category. However, we also notice that the optimal data augmentation method may vary depending on the specific model. For instance, for the Student-Teacher category, Flip is the best data augmentation method for RD4AD [6], but Color Jitter achieves the highest improvement for STPM [26]. This motivates us to investigate whether combining two data augmentation methods can further enhance the results compared to using only one. To combine two data augmentation methods, we apply them sequentially to an image. For example, we first rotate an image by 90 degrees and then translate it by a certain amount. In this way, we can generate more diverse and complex variations of the original image.

We choose different combinations of data augmentation methods for different categories of IAD methods based on their performance in the previous experiments. For Normalizing Flow-based methods, we combine Translate and Scale, which are both geometric transformations that can preserve the shape and appearance of the objects in the image. For Memory Bank-based and One-Class-based methods, we combine Rotation and Flip, which are both simple and effective data augmentation methods that can introduce some degree of rotation invariance and symmetry to the models. For Student-Teacher-based methods, we combine Translate, Color Jitter, and Flip, which can generate more diverse and challenging variations of the image by changing its position, color, and orientation. We do not conduct experiments with mixed data augmentation for Reconstruction-based methods because we have observed that data augmentation does not improve their performance and may even degrade it. This may be because data augmentation interferes with the reconstruction objective of these methods and makes it harder for them to learn mean-

Classes	Normalizing Flow		Memory Bank				Student-Teacher		One-Class	Reconstruction	
Methods	CSFlow	FastFlow	CFA	SPADE	PaDiM	PatchCore	RD4AD	STPM	CutPaste	FAVAE	DRAEM
Rotation	0.020	0.080	0.007	0.010	0.043	-0.001	0.036	0.000	0.067	-0.018	0.042
Flip	0.011	0.082	0.007	0.006	0.036	0.002	0.045	0.008	0.086	0.000	0.028
Scale	0.045	0.081	-0.034	0.001	0.001	0.002	0.021	0.041	-0.052	-0.056	0.001
Translate	0.037	0.093	-0.012	0.010	0.015	0.007	0.030	0.033	-0.030	-0.066	-0.012
Color_Jitter	0.023	0.066	0.005	-0.001	0.017	-0.004	0.022	0.047	-0.035	-0.047	0.012
Perspective	-0.006	0.038	-0.001	0.007	0.014	0.000	0.023	0.030	-0.029	-0.035	0.046

Table 2. The average improvement of image-level AUC-ROC of each data augmentation over each IAD method with respect to no data augmentation.




Figure 3. Input all training data of one class of Mvtec AD data set, 4 randomly selected few-shot training images, and images after data augmentation of these four into PatchCore. The distribution of embedding of each patch after dimensionality reduction is shown in the figure.

Method	Best Single Augmentation
CSFlow	Scale
FastFlow	Translate
CFA	Rotation
SPADE	Rotation
PaDiM	Rotation
PatchCore	Translate
RD4AD	Flip
STMP	Color Jitter
CutPaste	Flip

Table 3. The best single data augmentation corresponding to various IAD methods.

ingful features from the images.

We take the best single data augmentation corresponding to various IAD methods as the baseline for mixed data augmentation. As shown in Table 3, the best single data augmentation corresponding to each method varies. Next, we will investigate whether mixed data augmentation can further improve accuracy.

Table 4 demonstrates the improvement of the mixed data augmentation over the best single data augmentation in the image-level AUC-ROC metric. One can see that mixed data augmentation only has a more obvious improvement on PatchCore [17], and even causes a decline in accuracy in other methods.

### 4.3. Analysis

The main purpose of applying data augmentation to the training data is to increase its diversity and make it more




Figure 4. Using PatchCore to perform mixed data augmentation of 4-shot IAD in the Mvtec AD dataset can get good results in all classes.

similar to the distribution of the test data. However, as we can see from Figure 2, some data augmentation methods may change the orientation or appearance of the objects in the training images in ways that do not match the test im-

Shot	Dataset	Normalizing Flow		Memory Bank				S-T		One-class
		CSFlow	FastFlow	CFA	SPADE	PaDiM	PatchCore	RD4AD	STPM	CutPaste
1	MVTec AD	-0.029	0.029	0.006	-	0.004	0.004	-0.039	-0.012	0.049
	MTD	0.1	-0.026	0.007	-	-0.031	0.04	-0.075	-0.025	-0.094
	BTAD	0.001	0.009	-0.004	-	0	0.004	-0.227	-0.056	-0.003
2	MVTec AD	-0.047	-0.005	-0.005	0.003	-0.007	0.009	-0.087	-0.045	-0.029
	MTD	0.011	0.028	-0.001	-0.002	-0.007	0.001	-0.1	0.054	-0.011
	BTAD	-0.026	-0.06	-0.003	0	0.001	0.007	-0.154	-0.035	-0.013
4	MVTec AD	-0.123	-0.06	0.002	0.002	0	0.018	-0.104	-0.074	0.052
	MTD	0.011	-0.101	0.001	0	0.004	0.023	-0.069	0.028	0.012
	BTAD	-0.029	-0.103	-0.003	0.001	0.002	-0.008	-0.044	-0.073	-0.006
8	MVTec AD	-0.21	-0.137	0.007	-0.002	-0.012	0.008	-0.245	-0.142	0.004
	MTD	-0.152	-0.08	0	-0.001	0.01	-0.004	-0.133	0.072	0.047
	BTAD	-0.029	-0.15	-0.007	0.002	0	0.024	-0.361	-0.079	-0.036

Table 4. The improvement of the mixed data augmentation over the best single data augmentation of the image-level AUC-ROC metric.

ages. For example, after applying Rotation and Flip to a toothbrush image in the training set, the toothbrush may appear upside down or sideways, which is unlikely to occur in the test set. Therefore, for most IAD methods, using multiple data augmentation methods together may not be beneficial and may even harm their performance. However, there are two exceptions: PaDiM [4] and PatchCore [17]. These two methods do not rely on the global shape or structure of the image but rather on its local patches, as shown in Figure 2. Even if the image is transformed by data augmentation, each patch still retains some similarity to the normal image patches. However, PaDiM performs Gaussian fitting for patches at the same spatial location across different images and expects them to have low variance. Data augmentation introduces more diversity and noise to these patches and thus degrades PaDiM’s performance. On the other hand, PatchCore does not care about the spatial location or global changes of the patches and only focuses on their local features. Therefore, data augmentation has little negative impact on PatchCore’s performance and may even improve it further by generating more challenging variations of normal images.

To gain more insight into how PatchCore [17] works and to verify our previous claim, we conduct an experiment using the Mvtec AD dataset. We feed all the training images of one class, four randomly selected few-shot training images, and their augmented versions into PatchCore and extract the embeddings of all the patches. We then use t-SNE [25] to reduce the dimensionality of all the embeddings to two and visualize them in Figure 3. We can observe that the embeddings from the few-shot images are very sparse and only cover a small region of the feature space spanned by the full training data. However, after applying data augmentation to these images, their embeddings become dense and cover a larger region of the feature space. This suggests that data augmentation can effectively generate more diverse and realistic patches and thus improve the performance of few-shot IAD with PatchCore.

Figure 4 shows some examples of PatchCore’s [17] performance on different objects with mixed data augmenta-

tion using only 4-shot training images. We can see that PatchCore can not only correctly classify the images as normal or abnormal but also precisely locate the abnormal regions on the objects. This is because the patch embeddings obtained from data augmentation are similar to those from normal images and form a tight cluster in the feature space. Therefore, PatchCore can easily distinguish them from the abnormal patches that deviate from this cluster. As a result, PatchCore does not suffer from false positives, which means that it does not mistakenly label normal regions as abnormal.

## 5. Conclusions

In this paper, we conduct a comprehensive study on the role of data augmentation in few-shot IAD. We evaluate 11 algorithms on 3 datasets using 6 types of data augmentation methods and analyze their performance and characteristics. The experimental results show that data augmentation has different effects on different IAD algorithms depending on their underlying principles and assumptions. However, for algorithms that share similar technical approaches, such as embedding-based methods, data augmentation plays a similar role. Moreover, we also discover that while some data augmentation methods may improve the performance of a certain algorithm individually, combining them together may not be beneficial and may even degrade the anomaly detection accuracy. This is because some data augmentation methods may introduce unrealistic or inconsistent variations to the images that do not match the test images. However, PatchCore is an exception to this rule because it only focuses on the local features of patches and ignores their spatial location or global changes. Mixed data augmentation does not affect its accuracy but rather enhances it by generating more diverse and challenging normal images.

## 6. Acknowledgments

This work is supported by the National Natural Science Foundation of China under Grant No. 62122035, 62206122, and 61972188.



## References

- [1] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad — a comprehensive real-world dataset for unsupervised anomaly detection. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9584–9592, 2019. 5
- [2] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4183–4192, 2020. 1, 2
- [3] Niv Cohen and Yedid Hoshen. Sub-image anomaly detection with deep pyramid correspondences. *arXiv preprint arXiv:2005.02357*, 2020. 1, 2, 4, 5
- [4] Thomas Defard, Aleksandr Setkov, Angélique Loesch, and Romaric Audigier. Padim: a patch distribution modeling framework for anomaly detection and localization. In *International Conference on Pattern Recognition*, pages 475–489. Springer, 2021. 1, 4, 5, 8
- [5] David Dehaene and Pierre Eline. Anomaly localization by modeling perceptual features. *ArXiv*, abs/2008.05369, 2020. 1, 4, 5, 6
- [6] Hanqiu Deng and Xingyu Li. Anomaly detection via reverse distillation from one-class embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9737–9746, 2022. 2, 5, 6
- [7] Denis Gudovskiy, Shun Ishizaka, and Kazuki Kozuka. Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 98–107, 2022. 1
- [8] Chuanfei Hu, Kai Chen, and Hang Shao. A semantic-enhanced method based on deep svdd for pixel-wise anomaly detection. In *2021 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2021. 1
- [9] Chaoqin Huang, Haoyan Guan, Aofan Jiang, Ya Zhang, Michael Spratling, and Yan-Feng Wang. Registration based few-shot anomaly detection. *arXiv preprint arXiv:2207.07361*, 2022. 1, 2, 4
- [10] Yibin Huang, Congying Qiu, and Kui Yuan. Surface defect saliency of magnetic tile. *The Visual Computer*, 36(1):85–96, 2020. 5
- [11] Jieliang Jiang, Jiale Zhu, Muhammad Bilal, Yan Cui, Neeraj Kumar, Ruihan Dou, Feng Su, and Xiaolong Xu. Masked swin transformer unet for industrial anomaly detection. *IEEE Transactions on Industrial Informatics*, 2022. 1
- [12] Jin-Hwa Kim, Do-Hyeong Kim, Saehoon Yi, and Taehoon Lee. Semi-orthogonal embedding for efficient unsupervised anomaly segmentation. *arXiv preprint arXiv:2105.14737*, 2021. 1
- [13] Sungwook Lee, Seunghyun Lee, and Byung Cheol Song. Cfa: Coupled-hypersphere-based feature adaptation for target-oriented anomaly localization. *arXiv preprint arXiv:2206.04325*, 2022. 4, 5
- [14] Chun-Liang Li, Kihyuk Sohn, Jinsung Yoon, and Tomas Pfister. Cutpaste: Self-supervised learning for anomaly detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9664–9674, 2021. 2, 5
- [15] Ning Li, Kaitao Jiang, Zhiheng Ma, Xing Wei, Xiaopeng Hong, and Yihong Gong. Anomaly detection via self-organizing map. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 974–978. IEEE, 2021. 1
- [16] Pankaj Mishra, Riccardo Verk, Daniele Fornasier, Claudio Piciarelli, and Gian Luca Foresti. Vt-adl: A vision transformer network for image anomaly detection and localization. In *2021 IEEE 30th International Symposium on Industrial Electronics (ISIE)*, pages 01–06. IEEE, 2021. 5
- [17] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14318–14328, 2022. 1, 4, 5, 7, 8
- [18] Marco Rudolph, Bastian Wandt, and Bodo Rosenhahn. Same same but different: Semi-supervised defect detection with normalizing flows. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1907–1916, 2021. 1, 2, 4
- [19] Marco Rudolph, Tom Wehrbein, Bodo Rosenhahn, and Bastian Wandt. Fully convolutional cross-scale-flows for image-based defect detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1088–1097, 2022. 1, 2, 5
- [20] Mohammadreza Salehi, Niousha Sadjadi, Soroosh Baselzadeh, Mohammad H Rohban, and Hamid R Rabiee. Multiresolution knowledge distillation for anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14902–14912, 2021. 1
- [21] Hannah M Schlüter, Jeremy Tan, Benjamin Hou, and Bernhard Kainz. Self-supervised out-of-distribution detection and localization with natural synthetic anomalies (nsa). *arXiv preprint arXiv:2109.15222*, 2021. 1
- [22] Eli Schwartz, Assaf Arbelle, Leonid Karlinsky, Sivan Harary, Florian Scheidegger, Sivan Dovel, and Raja Giryes. Maeday: Mae for few and zero shot anomaly-detection. *arXiv preprint arXiv:2211.14307*, 2022. 2
- [23] Shelly Sheynin, Sagie Benaim, and Lior Wolf. A hierarchical transformation-discriminating generative model for few shot anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8495–8504, 2021. 1, 2
- [24] Kihyuk Sohn, Chun-Liang Li, Jinsung Yoon, Minh Jin, and Tomas Pfister. Learning and evaluating representations for deep one-class classification. In *International Conference on Learning Representations*, 2020. 2
- [25] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 8
- [26] Guodong Wang, Shumin Han, Errui Ding, and Di Huang. Student-teacher feature pyramid matching for anomaly detection. In *BMVC*, 2021. 1, 2, 5, 6

- [27] Shinji Yamada and Kazuhiro Hotta. Reconstruction student with attention for student-teacher pyramid matching. *arXiv preprint arXiv:2111.15376*, 2021. [1](#)
- [28] Jie Yang, Yong Shi, and Zhiqian Qi. Dfr: Deep feature reconstruction for unsupervised anomaly segmentation. *arXiv preprint arXiv:2012.07122*, 2020. [2](#)
- [29] Suorong Yang, Weikang Xiao, Mengcheng Zhang, Suhan Guo, Jian Zhao, and Furao Shen. Image data augmentation for deep learning: A survey. *arXiv preprint arXiv:2204.08610*, 2022. [4](#)
- [30] Jihun Yi and Sungroh Yoon. Patch svdd: Patch-level svdd for anomaly detection and segmentation. In *Proceedings of the Asian Conference on Computer Vision*, 2020. [1](#)
- [31] 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. [1](#), [2](#), [5](#)
- [32] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Draem—a discriminatively trained reconstruction embedding for surface anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8330–8339, 2021. [1](#), [2](#), [4](#), [5](#), [6](#)
- [33] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition*, 112:107706, 2021. [1](#)
- [34] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Dsr—a dual subspace re-projection network for surface anomaly detection. *arXiv preprint arXiv:2208.01521*, 2022. [1](#)
- [35] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. [5](#)