



A Prototype-Based Neural Network for Image Anomaly Detection and Localization

Chao Huang¹ · Zhao Kang¹ · Hong Wu¹

Accepted: 25 November 2023
© The Author(s) 2024

Abstract

Image anomaly detection and localization perform not only image-level anomaly classification but also locate pixel-level anomaly regions. Recently, it has received much research attention due to its wide application in various fields. This paper proposes ProtoAD, a prototype-based neural network for image anomaly detection and localization. First, the patch features of normal images are extracted by a deep network pre-trained on nature images. Then, the prototypes of the normal patch features are learned by non-parametric clustering. Finally, we construct an image anomaly localization network (ProtoAD) by appending the feature extraction network with $L2$ feature normalization, a 1×1 convolutional layer, a channel max-pooling, and a subtraction operation. We use the prototypes as the kernels of the 1×1 convolutional layer; therefore, our neural network does not need a training phase and can conduct anomaly detection and localization in an end-to-end manner. Extensive experiments on two challenging industrial anomaly detection datasets, MVTec AD and BTAD, demonstrate that ProtoAD achieves competitive performance compared to the state-of-the-art methods with a higher inference speed. The code and pre-trained models are publicly available at <https://github.com/98chao/ProtoAD>.

Keywords Image anomaly detection · Image anomaly localization · Non-parametric clustering · Prototype-based network

1 Introduction

Anomaly detection (AD) [1, 2] aims to detect anomalous samples that are deviated from a set of normal samples predefined during training. Traditional image anomaly detection adopts a semantic AD setting [3–6], where anomaly samples are from unknown semantic classes

✉ Hong Wu
hwu@uestc.edu.cn

Chao Huang
chaohuang@std.uestc.edu.cn

Zhao Kang
zkang@uestc.edu.cn

¹ School of Computer Science and Engineering, University of Electronic Science and Technology of China, Chengdu 611731, Sichuan, China




Fig. 1 Examples from the MVTEC benchmark datasets. From top to bottom: anomaly samples, anomaly mask, and anomaly score maps predicted by our method

different from the one normal samples belong to. Recently, detecting and localizing subtle image anomalies has become an important task in computer vision with various applications, such as anomaly or defect detection in industrial optical inspection [7, 8], anomaly detection and localization in video surveillance [9–11], or anomaly detection in medical images [12, 13]. In this setting, anomaly detection determines whether an image contains any anomaly, and anomaly localization, aka anomaly segmentation, localizes the anomalies at the pixel level. This paper focuses on the second setting, especially industrial anomaly detection and localization. Some examples from the MVTEC AD dataset [8] along with predictions by our method are shown in Fig. 1.

In the above applications, anomalous samples are scarce and hard to collect. Therefore, image anomaly detection and localization are often solved with only normal samples. In addition, anomalous regions within images are often subtle (see Fig. 1), making image anomaly localization a more challenging task that has not been thoroughly studied compared to image anomaly detection. Recent anomaly localization methods can be roughly categorized into two classes: reconstruction-based methods and OOD-based (out-of-distribution based) methods.

Reconstruction-based methods are mainly based on the assumption that a model trained only on normal images can not reconstruct anomalous images accurately. They reconstruct image as a whole [8, 12, 14–21], or reconstruct in the feature space [22–24]. Then anomaly detection and localization can be performed by measuring the difference between the reconstructed and original ones. This kind of method always needs cumbersome network training.

OOD-based methods evaluate the degree of abnormality for a patch feature by measuring its deviation from a set of normal patch features, which is intrinsically a patch-wise OOD detecting task. Some methods such as PatchSVDD [25] and CutPaste [26] learn feature representation by self-supervised learning. On the contrary, some other methods [27–30] simply extract features by deep networks pre-trained on natural image datasets such as ImageNet [31], and achieve promising and even better performances. Since the number of training patches is much larger than that of training images, the inference time and storage increase remarkably.

Different strategies have been proposed to tackle this problem. Napoletano et al. [27] used k-means to learn the dictionary/prototypes for normal patch features, but they evaluated each test patch independently, resulting in high inference time. SPADE [28] selects k-nearest normal images for patch-wise evaluation based on the global image features, limiting anomaly localization performance. PaDiM [29] models the normal patches at each position by a multidimensional Gaussian distribution and measures the anomaly by the Mahalanobis distance between a test patch feature and the Gaussian at the same position. However, both SPADE [28] and PaDiM [29] are reliant on image alignment. The current state-of-the-art method, PatchCore [30], uses greedy coresets subsampling to reduce the inference time and storage significantly.

This paper proposes ProtoAD, a prototype-based neural network for image anomaly detection and localization, to improve OOD-based methods' inference speed. We assume that all normal patch features can be grouped into some prototypes, and abnormal patch features cannot be properly assigned to any of them. Therefore, image anomaly localization can be performed by measuring the deviation of test patch features from the prototypes of normal patch features. First, the patch features of normal images are extracted by a deep network pre-trained on nature images and are L_2 -normalized. Then the prototypes of the normalized normal patch features are learned by a non-parametric clustering algorithm. The cosine similarity between two L_2 -normalized vectors is equivalent to the dot product between them. Therefore the cosine similarity between a normalized patch feature and a prototype can be implemented by a 1×1 convolution. Based on this equivalence, we construct an image anomaly localization network (ProtoAD) by appending the feature extraction network with the L_2 feature normalization, a 1×1 convolutional layer, a channel max-pooling, and a subtraction operation. We use the prototypes as the kernels of the 1×1 convolutional layer; therefore, our neural network does not need a training phase. Compared with previous OOD-based methods [27–30], ProtoAD can perform the anomaly detection and localization in an end-to-end manner, which is more elegant and efficient. Extensive experiments on two challenging industrial anomaly detection datasets, MVTec AD [8] and BTAD [32], demonstrate that ProtoAD achieves competitive performance compared to the state-of-the-art methods with a higher inference speed. This advantage of ProtoAD makes it better match the needs of real-world industrial applications.

2 Related Works

2.1 Image Anomaly Localization

Anomaly detection is an image-level task to determine whether an image contains any anomaly. On the other hand, anomaly localization is more complex to locate anomalies at the pixel level. Here, we only introduce the methods that can be directly applied to image anomaly localization and roughly categorize current methods into two types: reconstruction-based and OOD-based.

Reconstruction-based methods are mainly based on the assumption that a model trained only on normal images can not reconstruct anomalous images accurately, and anomaly detection and localization can be performed by measuring the difference between the reconstructed and original images. Early reconstruction-based methods [8, 12, 14, 15, 17] reconstruct image by auto-encoders (AE), variational autoencoders (VAE) or generative adversarial networks (GAN). However, the neural networks have high generalization capacities and can reconstruct anomalies well. Later, different strategies have been proposed to tackle this problem. Differ-

ent memory-based auto-encoders [16, 18, 20] have been proposed to reconstruct images with features from memory bank to limit the generalization ability. Student-teacher models [22, 23] have been used to reconstruct pre-trained deep features. RIAD [19] randomly removes partial image regions and reconstructs the image by image in-painting. Glance [24] trains a Global-Net to regress the deep features of cropped patches based on their context. DRAEM [21] combines a reconstructive sub-network and a discriminative network and trains them in an end-to-end manner on synthetically generated just-out-of-distribution images.

OOD-based methods evaluate the degree of abnormality for a patch feature by measuring its deviation from a set of normal patch features, which is intrinsically a patch-wise OOD detecting task. Some methods such as PatchSVDD [25] and CutPaste [26] learn feature representation by self-supervised learning. On the contrary, some other methods [27–30] simply extract features by deep networks pre-trained on natural image datasets such as ImageNet [31], and achieve promising and even better performances. Since the number of training patches is much larger than that of training images, the inference time and storage increase remarkably. Different strategies such as clustering, density estimation, and sampling have been proposed to tackle this problem. Napoletano et al. [27] learned a dictionary of normal patches from the training set by k-means, and evaluated each patch of a test image by measuring its visual similarity with the k-nearest neighbors in the dictionary. SPADE [28] compares patch features of a test image with the patch features at the same position of k-nearest normal images selected based on global image features. However, this oversimplified pre-selection strategy will limit the localization performance. PaDiM [29] models the normal patches at each position by a multidimensional Gaussian distribution and detect anomaly by the Mahalanobis distance between a test patch feature and the Gaussian at the same position. Both SPADE [28] and PaDiM [29] are reliant on image alignment. Recently, PatchCore [30] constructs the memory bank of locally aware patch features by greedy coresets subsampling, and localizes anomaly by measuring the distances of test patch features to their nearest normal patch features in the bank. As a result, PatchCore achieves a new state-of-the-art and significantly reduces the inference time and storage.

Our method is also an OOD-based method with pre-trained deep features but has several differences from the previous works. Our method uses non-parametric clustering instead of k-means in [27] to learn the prototypes for normal patch features. More importantly, our method can perform anomaly detection and localization by a network in an end-to-end manner, which is more elegant and efficient than the previous methods. Compared to reconstruction-based methods, our network do not need a cumbersome network training phase.

2.2 Clustering Algorithms

Clustering is a type of unsupervised learning task of dividing a set of unlabeled data points into a number of groups such that the data points in the same groups are more similar to each other than they are to the data points in other groups. Clustering provides an abstraction from data points to the clusters, and each cluster can be characterized by a cluster prototype, such as the centroid of a cluster, for further analysis. Clustering algorithms can be roughly divided into four categories: Partition-based cluster, Density-based clustering, Spectral Clustering, and Hierarchical-based clustering.

Partition-based clustering algorithms divide the data into k groups, where k is the pre-defined number of cluster. The classical algorithms are k-means [33] and its variations. Although these algorithms are very fast, they need the number of clusters as a parameter and are sensitive to the selection of the initial k centroids.

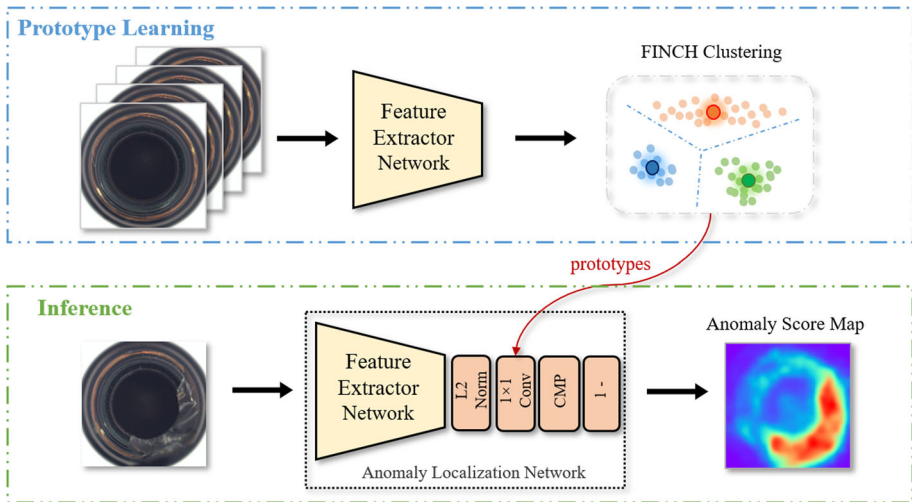


Fig. 2 An overview of the proposed method. First, the patch features of normal images are extracted by a deep network pre-trained on nature images. Then, the prototypes of the normal patch features are learned by FINCH clustering. For inference, an image anomaly localization network (ProtoAD) is constructed by appending the feature extraction network with the L_2 feature normalization, a 1×1 convolutional layer, a channel max-pooling (CMP), and a subtraction operation, and anomaly localization is performed in an end-to-end manner

Density-based clustering defines a cluster as the largest set of densely connected points and can find clusters of arbitrary shapes. DBSCAN [34] is the most representative algorithm of this class. It has two parameters, radius length ϵ and a parameter *MinPts*. If there are *MinPts* points in the radius of ϵ of a point, it is regarded as a high-density point.

Spectral Clustering [35] has recently attracted much attention. Most spectral clustering algorithms need to compute the full similarity graph Laplacian matrix and have quadratic complexities, thus severely restricting their application to large data sets.

Hierarchical clustering [36] is of two types: bottom-up and top-down approaches. In the bottom-up approach (aka agglomerative clustering), each data point starts as a cluster, and the most similar cluster pairs are iteratively merged according to the chosen similarity measure until some stopping criteria are met. In the top-down approach (aka divisive clustering), the clustering begins with a large cluster including all data and recursively breaks down into smaller clusters. Hierarchical clustering produces a clustering tree that provides meaningful ways to interpret data at different levels of granularity. Recently, Sarfraz et al. [37] proposed FINCH, a high-speed, scalable, and fully parameter-free hierarchical agglomerative clustering algorithm.

In [27], k-means is used to learn the prototypes from normal patch features. To avoid choosing the number of clusters ahead, we adopt FINCH to learn the prototypes for normal patch features.

3 Method

Our method consists of three steps: patch feature extraction, prototype learning, and anomaly detection and localization. An overview of our method is given in Figure. 2. We describe them sequentially in the following subsection.

3.1 Patch Feature Extraction

Since the features extracted by pre-trained networks have shown their effectiveness for various visual applications including anomaly detection [22, 23, 27–30], we also adopt deep networks pre-trained on ImageNet dataset [31] as the feature extractor, and choose the backbone of Wide-ResNet [38] as the feature extractor following the previous works [28–30].

ResNet-like deep networks [38, 39] include several convolutional stages. The features become more abstract when the stage goes deeper, but their resolution gets lower. Thus, the feature maps from different stages form a feature hierarchy for an input image. Each spatial position of a feature map has a receptive field and corresponds to a patch/region in an input image; therefore, the feature vector at a spatial position of feature maps can be considered as a feature representation for the corresponding image patch. If the feature maps of a stage have a resolution of $H \times W$, they contain $H \times W$ patch features. The deep and abstract features from the ImageNet pre-trained networks are biased towards the ImageNet classification task and are less relevant to the anomaly detection and localization task. Therefore, we adopt the low- and mid-level (stage 1–3) feature representations and combine them as the patch features. Concretely, the feature maps at the higher-level are bilinearly re-scaled to have the same resolution as the lowest level, then the feature maps at different levels are concatenated together for handling multi-scale anomalies. The extracted features are then L_2 -normalized where each feature vector is divided by its L_2 norm.

3.2 Prototype Learning

After feature extraction, the prototypes of the L_2 -normalized patch features are learned by a clustering algorithm. Then, the prototypes are used in anomaly detection and localization instead of all the normal patch features to reduce the inference time and storage. There are mainly two concerns in choosing a clustering algorithm. First, the number of patch features is much larger than that of training images. For example, each category of MVTEC AD dataset has several hundreds of images, while it has several hundreds of thousands of patch features in our implementation. Therefore, the clustering algorithm should be efficient and scalable to large-scale data. Second, most clustering algorithms have some parameters, e.g., the number of clusters or distance thresholds, which can not be well set without a priori knowledge of the data distribution. Thus, these algorithms demand a tedious parameter tuning process to achieve good performance. To meet the requirements of real applications, we adopt FINCH [37], a high-speed, scalable, and fully parameter-free hierarchical agglomerative clustering algorithm.

The core idea of FINCH is to use the nearest neighbor information of each data point for clustering, which does not need to specify any parameters and has a low computational overhead. Given the integer indices of the first neighbor of each data point, an adjacency matrix is defined according to the following rules:

$$A(i, j) = \begin{cases} 1, & \text{if } j = \kappa_i^1 \text{ or } \kappa_j^1 = i \text{ or } \kappa_i^1 = \kappa_j^1 \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where κ_i^1 symbolizes the first neighbor of data point i . This sparse adjacency matrix specifies a graph where connected data points form clusters. It directly provides clusters without solving a graph segmentation problem. After computing the first partition, FINCH merges the clusters recursively by using cluster means to compute the first neighbor of each cluster until all data points are included in a single cluster or until some stopping criteria is met. In

this work, we define the stopping criteria as the number of cluster is less than a threshold and set the threshold to 10,000 to get good results in our experiments. We choose the last partition as the clustering result, and use the mean vectors of clusters as the prototypes of normal patch features.

When the features are L_2 -normalized (making the length of a vector to 1), cosine similarity and Euclidean distance between the normalized features are equivalent in the sense of nearest neighbor searching:

$$\frac{1}{2}L_2(\mathbf{x}_a, \mathbf{x}_b)^2 = \frac{1}{2}(\mathbf{x}_a - \mathbf{x}_b) \cdot (\mathbf{x}_a - \mathbf{x}_b) = 1 - \mathbf{x}_a \cdot \mathbf{x}_b = 1 - \cos(\mathbf{x}_a, \mathbf{x}_b) \tag{2}$$

where $L_2()$ is Euclidean distance, \mathbf{x}_a and \mathbf{x}_b are two L_2 -normalized feature vectors, and \cos is cosine similarity. Therefore, we use cosine similarity for clustering and measuring the deviation of test patch features from norm patch features in the next subsection.

3.3 Neural Network for Anomaly Detection and Localization

When a test image passes through the feature extraction network, $H \times W$ patch features have been extracted. The anomaly score of each patch feature can be computed by measuring its deviation from the prototypes of normal patch features. We compute the anomaly score of a test patch as one minus the cosine similarity between the normalized test patch feature and its nearest prototype. Formally, the anomaly score for the patch at position (i, j) can be calculated as

$$s_{ij} = 1 - \max_{1 \leq k \leq K} \cos(\mathbf{x}_{ij}, \mathbf{m}_k) \tag{3}$$

where \mathbf{x}_{ij} is the normalized patch feature at position (i, j) , \mathbf{m}_k is the k -th prototype, and \cos is cosine similarity. In addition, the image-level anomaly score for a test image can be simply computed by maximizing the anomaly scores of all its patch features.

$$S = \max_{1 \leq i \leq H, 1 \leq j \leq W} s_{ij} \tag{4}$$

The cosine similarities between a normalized patch feature and a prototype can be computed by a 1×1 convolution (dot product) between them. Based on this equivalence, we construct a neural network (ProtoAD) for anomaly detection and localization. First, the L_2 feature normalization and a 1×1 convolutional layer are appended to the feature extraction network, and outputs feature maps of size $H \times W \times K$, including the cosine similarities between the $H \times W$ normalized patch features and all K prototypes. Then, channel max-pooling (CMP) is applied to the feature maps to get the normal score map of $H \times W$, including the cosine similarities between the $H \times W$ normalized patch features and their nearest prototypes. The anomaly score map can be further obtained by computing one minus the normal score map. This process is illustrated by Fig. 3. Since the spatial resolution of feature maps is lower than that of an input image, we resize the anomaly score map to the resolution of the input image and use a Gaussian filter to smooth it. Finally, anomaly localization can be achieved by thresholding the anomaly score map, and the anomaly score for the test image can be obtained by maximizing the anomaly score map.

We use the prototypes of normal patch features as the kernels of the 1×1 convolutional layer. Therefore the proposed neural network does not need a training phase. Compared to previous works [27–30], our method can perform the anomaly detection and localization in an end-to-end manner, which is more elegant and efficient.




Fig. 3 Anomaly detection and localization process of ProtoAD

4 Experiments

4.1 Datasets and Metrics

4.1.1 Dataset

MVTec AD dataset [8] is a real-world industrial defect detection dataset which has become a standard benchmark for evaluating image anomaly detection and localization methods. It has 5354 high-resolution images belonging to 10 objects and 5 texture categories. The images of each category are split into a training and a testing set. Totally, the training set has 3629 normal images, and the test set has 1725 normal and abnormal images of various defects. The ground truth of the test set contains anomaly labels for image-level evaluation and anomaly masks for pixel-level evaluation.

BTAD (BeanTech Anomaly Detection dataset) is a real-world industrial dataset recently released by [32]. It contains a total of 2830 real-world images of 3 industrial products. The images of each category are split into a defect-free training set and a testing set, supporting evaluation of both anomaly detection and localization.

We follow the split of the two datasets for training and testing.

4.1.2 Evaluation Metrics

AUROC (Area Under the Receiver Operating Characteristic curve) is the most commonly used metric for anomaly detection, which is independent of the threshold. We use image-level AUROC for evaluating the performance of anomaly detection, pixel-level AUROC for anomaly localization. Since the pixel-level AUROC is biased in favor of large anomalies, we also use PRO-score (per-region-overlap) [22] to evaluate anomaly localization, which weights ground-truth regions of different sizes equally.

4.2 Experimental Setup

We normalize the size of images from all categories of MVTEC AD and BTAD dataset to 256×256 , center crop images to 224×224 , and do not apply any data augmentation. The backbone of Wide-ResNet50 pre-trained on ImageNet is employed as the feature extractor in our method as in [28–30]. We define the stopping criteria for FINCH clustering algorithm as the number of clusters is less than 10,000 and choose the last generated partition as the clustering result. For inference, we up-sample the anomaly score map to image size using

bilinear interpolation and smooth it with the Gaussian filter with parameter $\delta = 4$ as in [29]. We implemented our models in Python 3.7 [40] and PyTorch [41], and run experiments on NVIDIA GeForce RTX 2080 Ti.

4.3 Results on MVTec AD

4.3.1 Comparison with the State-of-the-art

We compare ProtoAD with the state-of-the-art methods including both the reconstruction and OOD-based methods. The compared reconstruction-based methods include Uninformed students (U-Student) [22], RIAD [19], MKD [23], Glance [24], DAAD [20] and DREAM [21]. And the compared OOD-based methods include SPADE [28], PatchSVDD (P-SVDD) [25], CutPaste [26], PaDiM [29], and PatchCore (P-Core) [30]. We directly use their evaluation results if they have been provided.

We report the evaluation results (pixel-level AUROC and PRO-score) for pixel-level anomaly localization on MVTec AD dataset in Tables 1 and 2 respectively. From table 1, we can see that the OOD-based methods generally achieve better pixel-level AUROC than the reconstruct-based methods. Among the OOD-based methods, the methods using the pre-trained deep features achieve better pixel-level AUROC than the methods based on self-supervised learning. PatchCore achieves the best pixel-level AUROC, PaDiM the second, and the reconstruct-based method DREAM the third. The pixel-level AUROC of our method is very close to those of PaDiM and DREAM. We also notice that our method is more effective on the texture category and achieves the second best AUROC. Table 2 gives the PRO-score results for methods which have used this metric. Among them, Glance achieves the best result, our method is the second best and outperform other OOD-based methods. After all, our method achieves competitive anomaly localization performance to the state-of-the-art methods.

Figure 4 gives qualitative anomaly localization results of our method on MVTec AD dataset. We can see that our method can give accurate pixel-level localization regardless of anomaly region size and type (see supplementary for more qualitative results).

We also report the image-level AUROC results for anomaly detection in Table 3. PatchCore achieves the best AUROC again, DREAM the second. Our method remains competitive and achieves the third-best AUROC, which is very close to that of DREAM.

4.3.2 Inference Efficiency

Anomaly detection and localization algorithms need high precision and inference speed to match the requirements of real-world applications. Thus, we also report the inference speed of our method and previous OOD-based methods using pre-trained deep features [28–30] in the table 4. In the experiments, all the methods adopt Wide-ResNet50 pre-trained on ImageNet as the feature extractor, center-cropped 224×224 image as input, and run on the same machine with a NVIDIA GeForce RTX 2080 Ti. For PatchCore, we use the implementation provided by the authors, which downsamples the normal patch features via greedy coreset subsampling (PatchCore- $x\%$ denotes the percentage x of normal patch features are used in inference) and uses *faiss* [42] for nearest neighbor retrieval and distance computations. For PaDiM, we make extensive optimization via GPU acceleration. Compared with the previous methods, our model achieves the highest speed, which is 1.2x, 2.7x, and 9.5x faster than PaDiM, PatchCore, and SPADE, respectively. The high inference speed is mainly because

Table 1 Anomaly localization performance on MVTec AD (Pixel-level AUROC)

Category	Reconstruction-based				OOD-based				P-Core	ProtoAD
	MKD	Glance	RIAD	DRAEM	P-SVDD	CutPaste	SPADE	PaDiM		
Carpet	95.6	96.0	96.3	95.5	92.6	98.3	97.5	99.1	99.0	99.2
Grid	91.8	78.0	98.8	99.7	96.2	97.5	93.7	97.3	98.7	98.0
Leather	98.1	90.0	99.4	98.6	97.4	99.5	97.6	99.2	99.3	99.4
Tile	82.8	80.0	89.1	99.2	91.4	90.5	87.4	94.1	95.6	95.2
Wood	84.8	81.0	85.8	96.4	90.8	95.5	88.5	94.9	95.0	95.6
Texture	90.6	85.0	93.9	97.9	93.7	96.3	92.9	96.9	97.5	97.5
Bottle	96.3	93.0	98.4	99.1	98.1	97.6	98.4	98.3	98.6	98.3
Cable	82.4	94.0	84.2	94.7	96.8	90.0	97.2	96.7	98.4	97.5
Capsule	95.9	90.0	92.8	94.3	95.8	97.4	99.0	98.5	98.8	98.2
Hazelnut	94.6	84.0	96.1	99.7	97.5	97.3	99.1	98.2	98.7	98.8
Metal Nut	86.4	91.0	92.5	99.5	98.0	93.1	98.1	97.2	98.4	96.8
Pill	89.6	93.0	95.7	97.6	95.1	95.7	96.5	95.7	97.4	94.2
Screw	96.0	96.0	98.8	97.6	95.7	96.7	98.9	98.5	99.4	98.9
Toothbrush	96.1	96.0	98.9	98.1	98.1	98.1	97.9	98.8	98.7	98.8
Transistor	76.5	100	87.7	90.9	97.0	93.0	94.1	97.5	96.3	92.5
Zipper	93.9	99.0	97.8	98.8	95.1	99.3	96.5	98.5	98.8	96.7
Object	90.8	93.6	94.3	97.0	96.7	95.8	97.6	97.8	98.4	97.1
All	90.7	90.7	94.2	97.3	95.7	96.0	96.0	97.5	98.1	97.2

The best results of the two classes of methods are bold-faced respectively

Table 2 Anomaly localization performance on MVTec AD (PRO-score)

Category	Reconstruction-based		OOD-based			
	U-Student	Glance	SPADE	PaDiM	P-Core	ProtoAD
Carpet	87.9	97.7	94.7	96.2	96.6	97.0
Grid	95.2	93.2	86.7	94.6	96.0	93.9
Leather	94.5	90.9	97.2	97.8	98.9	98.1
Tile	94.6	88.3	75.6	86.0	87.3	87.2
Wood	91.1	94.1	87.4	91.1	89.4	93.2
Texture	92.7	92.8	88.3	93.1	93.6	93.9
Bottle	93.1	96.8	95.5	94.8	96.2	95.8
Cable	81.8	98.0	90.9	88.8	92.5	93.8
Capsule	96.8	96.0	93.7	93.5	95.5	93.7
Hazelnut	96.5	96.2	95.4	92.6	93.8	95.3
Metal Nut	94.2	96.7	94.4	85.6	91.4	94.2
Pill	96.1	97.8	94.6	92.7	93.2	94.7
Screw	94.2	100	96.0	94.4	97.9	94.7
Toothbrush	93.3	96.1	93.5	93.0	91.5	91.2
Transistor	66.6	99.9	87.4	84.5	83.7	87.9
Zipper	95.1	99.2	92.6	95.9	97.1	93.3
Object	90.8	97.7	93.4	91.6	93.3	93.4
All	91.4	96.1	91.7	92.1	93.4	93.6

The best results of the two classes of methods are bold-faced respectively




Fig. 4 Qualitative anomaly localization results of our method. From top to bottom: abnormal images, ground-truth, and anomaly score maps produced by our method

our model performs inference in an end-to-end manner, and the main computation added to the feature extraction network is the 1×1 convolutional layer. Compared to the reconstruction-based methods, our method does not need a cumbersome network training process.

4.4 Ablation Study

We report ablation studies on the MVTec AD dataset to evaluate the impact of different components of our method on the performance.

Table 3 Anomaly detection performance on MVTec AD (Image-level AUROC)

Category	Reconstruction-based				OOD-based				P-Core	ProtoAD
	MKD	DAAD	RIAD	DRAEM	SPADE	FaDiM	P-SVDD	CutPaste		
Carpet	79.3	86.6	84.2	97.0	-	-	92.6	98.3	98.7	99.5
Grid	78.0	95.7	99.6	99.9	-	-	96.2	97.5	98.2	93.7
Leather	95.1	86.2	100	100	-	-	97.4	99.5	100	100
Tile	91.6	88.2	98.7	99.6	-	-	91.4	90.5	98.7	99.2
Wood	94.3	98.2	93.0	99.1	-	-	90.8	95.5	99.2	99.1
Texture	87.6	91.0	95.1	99.1	-	98.8	93.7	96.3	99.0	98.3
Bottle	99.4	97.6	99.9	99.2	-	-	98.1	97.6	100	99.9
Cable	89.2	84.4	81.9	91.8	-	-	96.8	90.0	99.5	98.3
Capsule	80.5	76.7	88.4	98.5	-	-	95.8	97.4	98.1	93.1
Hazelnut	98.4	92.1	83.3	100	-	-	97.5	97.3	100	100
Metal Nut	73.6	75.8	88.5	98.7	-	-	98.0	93.1	100	99.9
Pill	82.7	90.0	83.8	98.9	-	-	95.1	95.7	96.6	95.8
Screw	83.3	98.7	84.5	93.9	-	-	95.7	96.7	98.1	94.9
Toothbrush	92.2	99.2	100	100	-	-	98.1	98.1	100	99.7
Transistor	85.6	87.6	90.9	93.1	-	-	97.0	93.0	100	97.7
Zipper	93.2	85.9	98.1	100	-	-	95.1	99.3	99.4	94.6
Object	87.8	88.8	89.9	97.4	-	93.6	96.7	95.8	99.2	97.4
All	87.7	89.5	91.7	98.0	85.5	95.3	95.7	96.0	99.1	97.7

The best results of the two classes of methods are bold-faced respectively

Table 4 Comparison of inference speed. Scores includes image-level AUROC, pixel-level AUROC, and PRO-score

Method	Scores	Inference Speed (FPS)
SPADE	(85.5, 96.0, 91.7)	7.58
PatchCore(25%)	(99.1, 98.1, 93.4)	22.30
PatchCore(10%)	(99.0, 98.1 , 93.5)	24.45
PatchCore(1%)	(99.0, 98.0, 93.1)	26.17
PaDiM	(95.3, 97.5, 92.1)	60.32
ProtoAD	(97.7, 97.2, 93.6)	72.45

The best results are bold-faced

Table 5 Anomaly detection and localization performance of ProtoAD with features at different levels. Each tuple shows image-level AUROC and pixel-level AUROC

Feature Level	Texture	Object	All
Level 1	(96.2, 96.6)	(85.5, 92.7)	(89.0, 94.0)
Level 2	(97.9, 97.3)	(97.2, 95.5)	(97.4, 96.1)
Level 3	(98.0, 96.7)	(95.7, 96.0)	(96.5, 96.2)
Level 2+3	(97.9, 97.3)	(96.8, 96.9)	(97.1, 96.9)
Level 1+2+3	(98.3, 97.5)	(97.4, 97.1)	(97.7, 97.2)

4.4.1 Feature Layer Selection

ResNet-like deep networks [38, 39] include several convolutional stages. The feature maps from different stages can compose a feature hierarchy for an image. Since the deepest feature maps in the hierarchy are biased towards the ImageNet classification task, we only adopt the features at the low and middle hierarchy levels (stage 1–3) for anomaly detection and localization. Table 5 gives the performance achieved with the features from different levels and their combination. It can be observed that the features from hierarchy level 2 can achieve the best performance among the first three levels, and a combination of the three levels can further improve the performance. Therefore, our method uses the combination of the first three feature levels as the patch feature.

4.4.2 Partition Selection from Clustering Hierarchy

FINCH is a hierarchical agglomerative clustering algorithm. It recursively merges clusters from the bottom up and provides a set of partitions in a hierarchical structure. Each successive partition is a super-set of its preceding partitions, and the number of clusters in it is smaller than those in the preceding partitions. Thus, we need select a partition from the clustering hierarchy as the clustering result.

We report the performance of our method with different partitions, from the second (P2) to the 6-th (P6) partition of FINCH, in Table 6 (see Table 1 in supplementary for more detailed results). We do not include the first partition because it has a huge number of clusters. The results in Table 6 indicate the average performance decreases along with the merging process. This may be because, when the number of clusters gets smaller, clusters are less compact and unsuitable for anomaly detection. On the other hand, if the number of clusters is too large, there are too many prototypes, and the inference time and storage would increase rapidly. We also give the “Best” performance, which FINCH can achieve by selecting the best partition

Table 6 Anomaly detection and localization performance of ProtoAD with different FINCH partitions. Each tuple shows image-level AUROC, pixel-level AUROC, and average cluster numbers

Partition	Texture	Object	All
P2	(98.9, 97.7, 48132)	(98.2, 97.2, 92904)	(98.5, 97.3, 77980)
P3	(98.5, 97.6, 7802)	(97.8, 97.3, 20616)	(98.0, 97.4, 16345)
P4	(98.4, 97.4, 1166)	(97.4, 97.1, 3397)	(97.7, 97.2, 2653)
P5	(98.1, 97.1, 234)	(95.2, 96.6, 902)	(96.1, 96.8, 679)
P6	(95.8, 96.4, 56)	(92.2, 95.9, 190)	(93.4, 96.1, 146)
Best	(99.0, 97.7, 17559)	(98.2, 97.2, 52844)	(98.5, 97.4, 41083)
Ours	(98.3, 97.5, 1787)	(97.4, 97.1, 4626)	(97.7, 97.2, 3680)

Table 7 Anomaly detection and localization performance of ProtoAD with different clustering methods. Each tuple shows image-Level AUROC and pixel-level AUROC

Category	K-Means (L2)	K-Means (Norm L2)	FINCH (Coisne)
Texture	(95.0, 95.3)	(97.2, 96.6)	(98.3, 97.5)
Object	(92.9, 95.9)	(95.4, 96.1)	(97.4, 97.1)
All	(93.6, 95.7)	(96.0, 96.2)	(97.7, 97.2)

Table 8 Anomaly detection and localization performance on BTAD (Image-level and Pixel-level AUROC). The best results are bold-faced

Category	AE(MSE)	AE(MSE+SSIM)	VT-ADL	SPADE	PaDiM	PatchCore	ProtoAD
01	(-, 49.0)	(-, 53.0)	(-, 99.0)	(93.2, 90.2)	(100.0, 97.2)	(95.4, 96.2)	(97.0, 95.5)
02	(-, 92.0)	(-, 96.0)	(-, 94.0)	(74.8, 93.5)	(81.5, 95.4)	(85.1, 95.2)	(85.2, 96.5)
03	(-, 95.0)	(-, 89.0)	(-, 77.0)	(99.4, 96.3)	(98.6, 99.6)	(99.7, 99.5)	(99.8, 99.0)
All	(-, 78.7)	(-, 79.3)	(-, 90.0)	(89.1, 93.3)	(93.4, 97.4)	(93.4, 97.0)	(94.0, 97.0)

for each category respectively. This best performance is the upper bound that our method can achieve. However, selecting partition based on the average performance (from P2 to P6) or performance for each category (Best) is time-consuming and not suitable for real applications. In our method, we stop FINCH when the number of cluster is less than 10,000 and use the final partition as the clustering result, and give its results in the last line of Table 6. Our partition selection rule can achieve performance very close to the best one with only a tenth of clusters. Therefore, our method can reach a good trade-off between effectiveness and efficiency.

4.4.3 FINCH vs. K-Means

We compare FINCH clustering algorithm with k-means for the prototype-based anomaly detection. In our method, we choose the partition generated so far by FINCH which having less than 10,000 clusters as the clustering result. For a fair comparison, we set k to 10,000 for k-means. The results in table 7 indicate that the method based on FINCH (the third column) achieves better performance than that based on k-means (the first column). Although it may achieve better performance for k-means by tuning k, it is time-consuming and not feasible for real applications.

4.4.4 Feature Normalization and Cosine Similarity

We also explore the importance of feature normalization for the prototype-based anomaly detection. As shown in Table 7, k-means with Euclidean distance on the L_2 -normalized features (Norm L2) outperforms k-means with Euclidean distance on the original features (L_2) in both anomaly detection and anomaly localization and achieves greater improvements in anomaly detection.

When the features are L_2 -normalized, cosine similarity and Euclidean distance are equivalent in the sense of nearest neighbor searching. Therefore, we use cosine similarity for clustering and measuring the deviation of test patch features from norm patch features. We further implement cosine similarity with a 1×1 convolution and append it to the feature extraction network. Therefore inference can be performed in an end-to-end manner.

4.5 Results on BTAD

In Table 8, we report the results of our method on the BTAD dataset and compare them with those of the SOTA OOD-based method (SPADE, PaDiM, and ProtoAD) and the approaches adopted in [32]. In [32], three reconstruction-based methods have been evaluated, auto-encoder (AE) with MSE loss, auto-encoder with MSE and SSIM loss, and Vision-Transformer-based image anomaly detection and localization (VT-ADL). We report the image-level and pixel-level AUROC for each category and their average for all categories. For anomaly detection, ProtoAD achieved the best image-level AUROC. For anomaly localization, ProtoAD achieved the second-best pixel-level AUROC (97.0), very close to the best one (97.4) achieved by PaDiM. These results show our method's potential to generalize to new anomalous scenarios.

5 Conclusion

We propose ProtoAD, a new OOD-based image anomaly detection and localization method. First, a pre-trained neural network is used to extract features for image patches. Then, a non-parametric clustering algorithm learns the prototypes for normal patch features. Finally, an image anomaly detection and localization network is constructed by appending the feature extraction network with the L_2 feature normalization, a 1×1 convolutional layer, a channel max-pooling, and a subtraction operation. As a result, ProtoAD does not need a network training process and can conduct anomaly detection and localization in an end-to-end manner. Experimental results on the MVTec AD dataset and the BTAD dataset show that ProtoAD can achieve competitive performance compared to state-of-the-art methods. Furthermore, compared to other OOD-based methods, ProtoAD is more elegant and efficient. And compared to the reconstruct-based methods, ProtoAD does not need a cumbersome network training process. Therefore, it can better meet the requirements of real applications.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s11063-024-11466-7>.

Author Contributions All authors contributed to the study's conception and design. CH performed material preparation, data collection, and analysis. CH and prof. HW wrote the first draft of the manuscript, and all authors commented on previous versions of the manuscript. All authors read and approved the final manuscript.

Funding This research was supported by the National Defense Basic Scientific Research Program of China under Grant JCKY2020903B002.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Chandola V, Banerjee A, Kumar V (2009) Anomaly detection: a survey. *ACM Comput Surv (CSUR)* 41(3):1–58
2. Salehi M, Mirzaei H, Hendrycks D, Li Y, Rohban M, Sabokrou M, et al (2022) A unified survey on anomaly, novelty, open-set, and out of-distribution detection: Solutions and future challenges. *Trans Mach Learn Res* (234)
3. Ruff L, Vandermeulen R, Goernitz N, Deecke L, Siddiqui SA, Binder A, Müller E, Kloft M (2018) Deep one-class classification. In: *International Conference on Machine Learning*, pp. 4393–4402
4. Sabokrou M, Khalooei M, Fathy M, Adeli E (2018) Adversarially learned one-class classifier for novelty detection. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3379–3388
5. Golan I, El-Yaniv R (2018) Deep anomaly detection using geometric transformations. *Adv Neural Inf Proc Syst* 31
6. Bergman L, Hoshen Y (2020) Classification-based anomaly detection for general data. [arXiv:2005.02359](https://arxiv.org/abs/2005.02359)
7. Mei S, Wang Y, Wen G (2018) Automatic fabric defect detection with a multi-scale convolutional denoising autoencoder network model. *Sensors* 18(4):1064
8. Bergmann P, Fauser M, Sattlegger D, Steger C (2019) MVTec AD—A comprehensive real-world dataset for unsupervised anomaly detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9592–9600
9. Sabokrou M, Fayyaz M, Fathy M, Klette R (2017) Deep-cascade: cascading 3D deep neural networks for fast anomaly detection and localization in crowded scenes. *IEEE Trans Image Process* 26(4):1992–2004
10. Sabokrou M, Fayyaz M, Fathy M, Moayed Z, Klette R (2018) Deep-anomaly: fully convolutional neural network for fast anomaly detection in crowded scenes. *Comput Vis Image Underst* 172:88–97
11. Sabokrou M, Pourreza M, Fayyaz M, Entezari R, Fathy M, Gall J, Adeli E (2018) Avid: adversarial visual irregularity detection. In: *Asian Conference on Computer Vision*, pp. 488–505
12. Schlegl T, Seeböck P, Waldstein SM, Schmidt-Erfurth U, Langs G (2017) Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In: *International Conference on Information Processing in Medical Imaging*, pp. 146–157
13. Li Z, Wang C, Han M, Xue Y, Wei W, Li L-J, Fei-Fei L (2018) Thoracic disease identification and localization with limited supervision. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8290–8299
14. An J, Cho S (2015) Variational autoencoder based anomaly detection using reconstruction probability. *Spec Lect IE 2(1)*:1–18
15. Bergmann P, Löwe S, Fauser M, Sattlegger D, Steger C (2018) Improving unsupervised defect segmentation by applying structural similarity to autoencoders. [arXiv:1807.02011](https://arxiv.org/abs/1807.02011)
16. Gong D, Liu L, Le V, Saha B, Mansour MR, Venkatesh S, Hengel A (2019) Memorizing normality to detect anomaly: memory-augmented deep autoencoder for unsupervised anomaly detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1705–1714
17. Liu W, Li R, Zheng M, Karanam S, Wu Z, Bhanu B, Radke RJ, Camps O (2020) Towards visually explaining variational autoencoders. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8642–8651
18. Park H, Noh J, Ham B (2020) Learning memory-guided normality for anomaly detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14372–14381

19. Zavrtnik V, Kristan M, Skočaj D (2021) Reconstruction by inpainting for visual anomaly detection. *Pattern Recogn* 112:107706
20. Hou J, Zhang Y, Zhong Q, Xie D, Pu S, Zhou H (2021) Divide-and-assemble: learning block-wise memory for unsupervised anomaly detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8791–8800
21. Zavrtnik V, Kristan M, Skočaj D (2021) Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8330–8339
22. Bergmann P, Fauser M, Sattlegger D, Steger C (2020) Uninformed students: student-teacher anomaly detection with discriminative latent embeddings. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4183–4192
23. Salehi M, Sadjadi N, Baselizadeh S, Rohban MH, Rabiee HR (2021) Multiresolution knowledge distillation for anomaly detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14902–14912
24. Wang S, Wu L, Cui L, Shen Y (2021) Glancing at the patch: anomaly localization with global and local feature comparison. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 254–263
25. Yi J, Yoon S (2021) Patch svdd: Patch-level svdd for anomaly detection and segmentation. In: *Computer Vision-ACCV 2020: 15th Asian Conference on Computer Vision, Kyoto, Japan, November 30–December 4, 2020, Revised Selected Papers, Part VI 15*, pp. 375–390
26. Li C-L, Sohn K, Yoon J, Pfister T (2021) Cutpaste: self-supervised learning for anomaly detection and localization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9664–9674
27. Napoletano P, Piccoli F, Schettini R (2018) Anomaly detection in nanofibrous materials by CNN-based self-similarity. *Sensors* 18(1):209
28. Cohen N, Hoshen Y (2020) Sub-image anomaly detection with deep pyramid correspondences. [arXiv:2005.02357](https://arxiv.org/abs/2005.02357)
29. Defard T, Setkov A, Loesch A, Audigier R (2021) Padim: a patch distribution modeling framework for anomaly detection and localization. In: *International Conference on Pattern Recognition*, pp. 475–489
30. Roth K, Pemula L, Zepeda J, Schölkopf B, Brox T, Gehler P (2022) Towards total recall in industrial anomaly detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14318–14328
31. Deng J, Dong W, Socher R, Li L-J, Li K, Fei-Fei L (2009) Imagenet: a large-scale hierarchical image database. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255
32. Mishra P, Verk R, Fornasier D, Piciarelli C, Foresti GL (2021) Vt-adl: a vision transformer network for image anomaly detection and localization. In: *2021 IEEE 30th International Symposium on Industrial Electronics (ISIE)*, pp. 01–06
33. MacQueen J (1965) Some methods for classification and analysis of multivariate observations. In: *Proceedings of Berkeley Symposium on Mathematical Statistics & Probability*
34. Ester M, Kriegel H-P, Sander J, Xu X, et al (1996) A density-based algorithm for discovering clusters in large spatial databases with noise. In: *Kdd*, 96, 226–231
35. Von Luxburg U (2007) A tutorial on spectral clustering. *Stat Comput* 17(4):395–416
36. Karypis G, Han E-H, Kumar V (1999) Chameleon: hierarchical clustering using dynamic modeling. *Computer* 32(8):68–75
37. Sarfraz S, Sharma V, Stiefelwagen R (2019) Efficient parameter-free clustering using first neighbor relations. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8934–8943
38. Zagoruyko S, Komodakis N (2016) Wide residual networks. In: *Proceedings of the British Machine Vision Conference*, 2016
39. He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778
40. Van RG, Drake F (2009) Python 3 reference manual. CreateSpace, Scotts Valley, p 1593511
41. Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, et al (2019) Pytorch: an imperative style, high-performance deep learning library. *Adv Neural Inf Proc Syst* 32
42. Johnson J, Douze M, Jégou H (2019) Billion-scale similarity search with gpus. *IEEE Trans Big Data* 7(3):535–547