

Constrained Adaptive Projection with Pretrained Features for Anomaly Detection

Xingtai Gui¹, Di Wu¹, Yang Chang^{1,2}, Shicai Fan^{1,2*}

¹School of Automation Engineering, University of Electronic Science and Technology of China(UESTC)

²Shenzhen Institute for Advanced Study, UESTC

{tabgui,wudi123,changyang}@std.uestc.edu.cn, shicaifan@uestc.edu.cn

Abstract

Anomaly detection aims to separate anomalies from normal samples, and the pretrained network is promising for anomaly detection. However, adapting the pretrained features would be confronted with the risk of pattern collapse when finetuning on one-class training data. In this paper, we propose an anomaly detection framework called constrained adaptive projection with pretrained features (CAP). Combined with pretrained features, a simple linear projection head applied on a specific input and its k most similar pretrained normal representations is designed for feature adaptation, and a reformed self-attention is leveraged to mine the inner-relationship among one-class semantic features. A loss function is proposed to avoid potential pattern collapse. Concretely, it considers the similarity between a specific data and its corresponding adaptive normal representation, and incorporates a constraint term slightly aligning pretrained and adaptive spaces. Our method achieves state-of-the-art anomaly detection performance on semantic anomaly detection and sensory anomaly detection benchmarks including 96.5% AUROC on CIFAR-100 dataset, 97.0% AUROC on CIFAR-10 dataset and 89.9% AUROC on MvTec dataset.

1 Introduction

Anomaly is the pattern that deviates considerably from the concept of normality and detecting such pattern is of key importance in science and industry. Anomaly detection (AD) is a specific task designed to learn a model that accurately detects anomalous test samples [Ruff *et al.*, 2021]. The anomaly may appear due to either covariate shift or semantic shift and these shifts lead to two sub-tasks: sensory AD or semantic AD respectively [Yang *et al.*, 2021].

In standard AD settings, labeled anomalous data are often nonexistent and only normal data are accessible. In such one-class training case, self-supervised methods based on auxiliary tasks, or unsupervised methods like autoencoder are widely used in AD. However, for self-supervised methods,

the features trained on auxiliary domains may not generalize well to the target domain and for autoencoder-based methods, their generalization ability could reconstruct the abnormal inputs well and lead to a misjudgment on anomalies.

Due to the restriction of one-class training data, the learned representation is indistinguishable to some extent resulting in the limited performance for anomaly detection. Recently, anomaly detection based on pretrained features has been widely studied. Some works [Bergmann *et al.*, 2020; Salehi *et al.*, 2021] only consider leveraging knowledge distillation to transfer the pretrained features of anomaly-free data to the student network. However, they do not adapt the features to the target data set. PANDA [Reiss *et al.*, 2021] proposes a baseline based on pretrained network, and implements the finetune by referring to Deep-SVDD [Ruff *et al.*, 2018]. But such finetune method will cause pattern collapse which means all features will shrink to the center point (See Appendix A for details about pattern collapse). PANDA leverages elastic weight consolidation to restrict the weight change in finetune stage. However, it still needs to further pretrain on an auxiliary task to obtain a Fisher information matrix.

In this work, we introduce a simple yet effective framework considering pretrained feature adaptation which is more suitable for anomaly detection. In order to avoid potential pattern collapse, the framework abandons the traditional optimization with a global parameter such as mean center but adopts a strategy paying attention to normal patterns locally. For a specific input image, its k nearest normal representations in pretrained feature space are traced. Instead of finetuning the layers in pretrained network, we propose a simple adaptive projection head that can play a role in adapting pretrained features properly and map the input and its k nearest normal pretrained features. Self-attention [Vaswani *et al.*, 2017] is reformed to obtain the weights of the projected normal features whose weighted average is regarded as an adaptive normal representation of the input. As a novel anomaly criterion, the similarity between the projected feature of input and the corresponding normal representation is optimized. In addition, based on such framework, we propose a constraint term considering an alignment between the pretrained and adaptive feature spaces to ensure obtaining nontrivial solutions. The joint optimization of the similarity and constraint can avoid pattern collapse and offer promising detection performance.

*The corresponding author

We summarize our contributions as follows: i) we design an anomaly detection framework called constrained adaptive projection with pretrained feature (CAP). Instead of finetuning the large pretrained network, the pretrained features adaptation is with a simple projection head. ii) under the one-class setting, self-attention is reformed to mine inner-relationship between projected k nearest normal features to bring semantic interpretability and enhance the quality of adaptive representation. iii) we propose a novel criterion for anomaly detection and an effective constraint avoiding learning trivial solution. Consequently, without global optimization goal, it can reduce the risk of pattern collapse in optimization. iv) extensive experiments and visualizations validate the effectiveness of our proposed framework for both sensory AD and semantic AD. The source code of CAP is released at <https://github.com/TabGuigui/CAP>.

2 Related work

Anomaly Detection. Anomaly detection aims to detect any anomalous samples that are deviated from the predefined normality during testing [Yang *et al.*, 2021]. Considering traditional anomaly detection with specific machine learning model, kNN [Angiulli and Pizzuti, 2002] was used to detect outliers. Combined with one-class classification, OC-SVM [Schölkopf *et al.*, 2001] was proposed to obtain a discriminative hyperplane for anomaly detection. Based on statistical methods, such as PCA [Ding and Kolaczyk, 2013] and GMM [Glodek *et al.*, 2013], some suitable models were established to fit the characteristic distribution of normal data. However, it is difficult for such classical methods to work satisfactorily when facing high-dimensional data, such as image anomaly detection tasks.

Deep-learning Based Anomaly Detection. Recently, deep learning has pushed the performance of computer vision systems to soaring heights on a broad array of high-level problems. Many deep learning methods have been introduced into anomaly detection community. Ruff *et al.* proposed Deep-SVDD [Ruff *et al.*, 2018] and Deep SAD [Ruff *et al.*, 2019] to tackle anomaly detection in unsupervised and semi-supervised way specifically. Gong *et al.* [Gong *et al.*, 2019] added memory module on the basis of autoencoders so as to enlarge the reconstruction errors of abnormal samples for detection. Perera *et al.* [Perera *et al.*, 2019] proposed OC-GAN method, which would find a potential space to make the reconstructed image of the generator similar to the normal sample. Combined with self-supervised learning, MHRot [Hendrycks *et al.*, 2019] was proposed to learn more effective representations leveraging auxiliary tasks.

Transfer-learning Based Anomaly Detection. Since anomaly detection is a data-poor task, representations learned on extensive datasets can be leveraged to tackle the limitation. The experiments of Bergman *et al.* [Bergman *et al.*, 2020] proved that the performance of many self-supervised learning methods applied in image anomaly detection in recent years is far inferior to the simple method based on pretrained strong feature extractor. Salehi *et al.* [Salehi *et al.*, 2021] designed an anomaly location method utilizing the different behaviors of the clone network distilled

by knowledge of pretrained expert network to detect and locate anomalies. Dong *et al.* [Dong *et al.*, 2021] designed a new unsupervised semantic transfer model to explore domain-invariant knowledge between different data distributions. Reiss *et al.* [Reiss *et al.*, 2021] proposed an approach to adapt pretraining features and suggested some training methods to mitigate pattern collapse.

3 Our Approach

CAP designs a simple adaptive projection head and reforms self-attention for feature adaptation. A novel loss function focusing on data and its adaptive normal representation, with an effective constraint term is proposed. The whole framework is illustrated in Figure 1 which includes a K-normal nearest neighbors module, an adaptive projection head module and a reformed self-attention module.

3.1 K-normal Nearest Neighbors.

For a specific input, CAP firstly traces its k most similar normal representations in pretrained space. In the specific implementation, instead of exhaustively computing pretrained representations every time, a memory bank is used to store the pretrained features of all normal training samples. The memory items are obtained through a pretrained large network G_P which is regarded as an encoder. A memory addressing module is leveraged to find the k -normal nearest neighbors.

Pretrained Memory Bank. The memory bank $\mathbf{M} \in \mathbb{R}^{N \times D}$ is a real-valued matrix containing N items of fixed dimension D and \mathbf{m}_i represents a memory item. The memory item \mathbf{m}_i is obtained from the encoder G_P , $\mathbf{m}_i = G_P(x_i^{train})$ where x_i^{train} is a training normal image.

Top-K Memory Addressing. For a specific image x_i , we first compute its pretrained encode $\mathbf{z}_i = G_P(x_i)$, and then compare it against all the pretrained memory items. With a cosine similarity $s_{j,i} = \frac{\mathbf{m}_j \cdot \mathbf{z}_i}{\|\mathbf{m}_j\| \|\mathbf{z}_i\|}$, the k -normal nearest neighbors in pretrained space are recorded in a subset $\mathbf{M}_i \in \mathbb{R}^{K \times D}$ and would then be used to construct a normal representation. The similarity ranking of memory items corresponding to \mathbf{z}_i is $[\mathbf{m}_{i1}, \mathbf{m}_{i2}, \mathbf{m}_{i3}, \dots, \mathbf{m}_{in}]$ (from the most to the least). It should be noted that in the training phase, $\mathbf{z}_i = \mathbf{m}_{i1}$, thus the appropriate K-normal nearest neighbors are $\mathbf{M}_i = [\mathbf{m}_{i2}, \mathbf{m}_{i3}, \dots, \mathbf{m}_{i(k+1)}]$ while in the test phase, $\mathbf{M}_i = [\mathbf{m}_{i1}, \mathbf{m}_{i2}, \dots, \mathbf{m}_{ik}]$.

3.2 Simple Adaptive Projection Head

Instead of finetuning some layers of the pretrained large network, we add an adaptive projection head G_θ to obtain adaptive representations. [Chen *et al.*, 2020; Tolstikhin *et al.*, 2021] proposed that a function including linear transformation and nonlinear activation function can effectively improve the representation ability. In CAP, a simple adaptive projection head merely including linear transformation is leveraged to finetune \mathbf{z}_i and \mathbf{M}_i simultaneously and get their corresponding mapping

$$G_\theta(\mathbf{z}; W_\theta) = W_\theta \mathbf{z} \quad W_\theta \in \mathbb{R}^{D \times D} \quad (1)$$

$$\hat{\mathbf{z}}_i = G_\theta(\mathbf{z}_i; W_\theta) \quad \hat{\mathbf{M}}_i = G_\theta(\mathbf{M}_i; W_\theta) \quad (2)$$

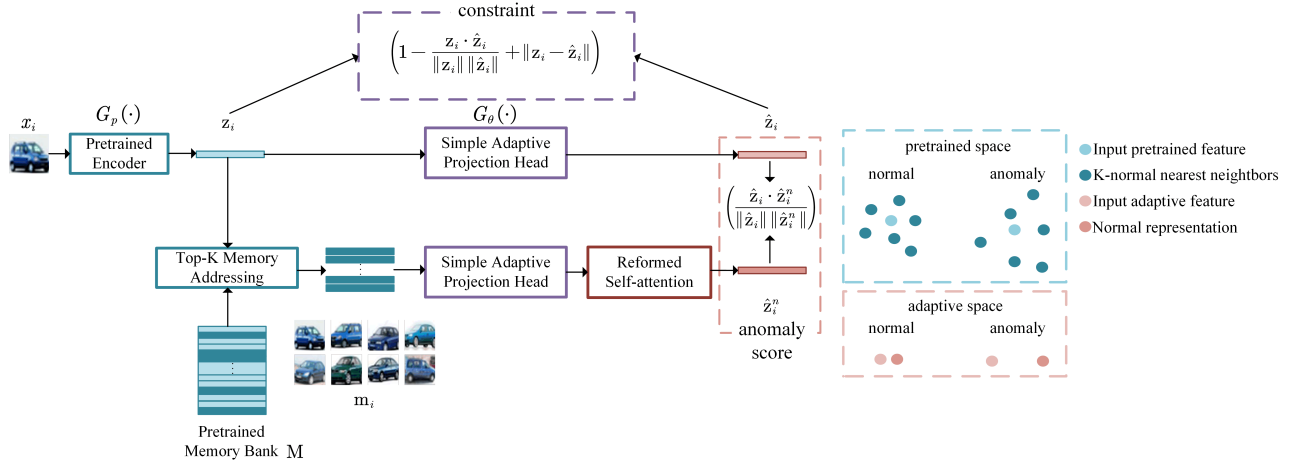


Figure 1: **Left:** The architecture of CAP contains a pretrained network, a K-normal nearest neighbors module, a simple adaptive projection head and a reformed self-attention module. Here, for a specific input, \mathbf{z}_i is the pretrained feature, $\hat{\mathbf{z}}_i$ is the adaptive feature and $\hat{\mathbf{z}}_i^n$ is its corresponding adaptive normal representation. It should be noted that the two projection heads are weight-shared. **Right:** A toy example of CAP. In the pretrained space, the dots with light color are the input, and the deep ones are the normal training data. In the adaptive space, the distance between the anomaly and its adaptive normal representation is large.

We have found and proved experimentally that in anomaly detection task, a simple one-layer mapping without activation function is more effective than multi-layer mapping. Moreover, for better adapting to target data, the structure of adaptive projection is designed as a square matrix. These contents will be described in detail later.

3.3 Reformed Self-attention

For K-normal nearest neighbors, their feature-level mixup can be used as a better normal representation. The general way is with mean value or considering the cosine similarity with a query. However, the relationship among these items is not considered. In CAP, we believe that the nearest k adaptive representations in $\hat{\mathbf{M}}_i$ can learn the weight considering their inner-relationship. [Vaswani *et al.*, 2017] proposed self-attention to learn the relationship in sequences paralleled and [Su *et al.*, 2019] applied a batch-attention module to capture the discriminative information from similar objects. We consider that the subset $\hat{\mathbf{M}}_i$ contains the information that can infer $\hat{\mathbf{z}}_i$, naturally, we reform the self-attention mechanism to adapt to the one-class scenario.

$$A(\hat{\mathbf{M}}_i; W_Q, W_K) = \text{Softmax}\left(\frac{(\hat{\mathbf{M}}_i W_Q)(\hat{\mathbf{M}}_i W_K)^T}{\sqrt{D}}\right) \hat{\mathbf{M}}_i \quad (3)$$

$W_Q, W_K \in \mathbb{R}^{D \times D}$ are linear transforms and add sufficient expressive power for normal features in adaptive space. In addition to such architecture, a residual connection and average operation are inserted and the output of such module is

$$\hat{\mathbf{z}}_i^n = \text{Mean}(\hat{\mathbf{M}}_i + A(\hat{\mathbf{M}}_i; W_Q, W_K)) \quad (4)$$

$\hat{\mathbf{z}}_i^n$ is the weighted average of subset $\hat{\mathbf{M}}_i$ whose weight is learned considering the inner-relationship. Compared with the basic self-attention structure, we remove W_V since the

data space should keep unchanged. Moreover, the layer normalization is removed because its anti-overfitting property is unfavorable to anomaly detection. Since $\hat{\mathbf{z}}_i^n$ is calculated with transformed normal subset, it can be regarded as an adaptive normal representation.

3.4 CAP Loss Function

We make a hypothesis that the nearest k items in the pretrained space are still the nearest ones in the adaptive space. Thus, we look upon $\hat{\mathbf{z}}_i^n$ as the normal representation corresponding to $\hat{\mathbf{z}}_i$. Naturally, like Figure 1 right, for a normal sample, $\hat{\mathbf{z}}_i$ would be similar to $\hat{\mathbf{z}}_i^n$ while for an anomaly, their dissimilarity is large. In the training phase, since all accessible samples are normal, we take the similarity between $\hat{\mathbf{z}}_i$ and $\hat{\mathbf{z}}_i^n$ as the optimization target.

$$L_S = \frac{1}{N} \sum_{i=1}^N \left(1 - \frac{\hat{\mathbf{z}}_i \cdot \hat{\mathbf{z}}_i^n}{\|\hat{\mathbf{z}}_i\| \|\hat{\mathbf{z}}_i^n\|}\right) \quad (5)$$

Unlike the loss function mentioned in Deep-SVDD and PANDA which reduces the distance between training data and a global target, such loss function optimizes the similarity between a specific input and its k-normal nearest neighbors in the adaptive space. Therefore, global pattern collapse can be avoided to some extent. Considering L_S only, CAP may tend to learn a trivial solution. However, based on such framework, a simple and practicable constraint is proposed to effectively avoid such potential risks. Here we theoretically demonstrate the potential trivial solution and the feasible method.

Proposition. Let G_θ be the all-zero network weights. For this choice of parameters, the adaptive projection head maps any input to the same output, i.e. $G_\theta(\mathbf{z}_i) = \mathbf{0}$.

Since G_θ is a simple linear projection, $G_\theta = \mathbf{0}$ is a potential solution. As the output of the all-zero-weights projection is zero for every input, L_S will become meaningless.

Method	CIFAR-10	CIFAR-100	FMNIST	MvTec
Deep-SVDD[Ruff <i>et al.</i> , 2018]	64.8	67.0	84.8	77.9
MHRot[Hendrycks <i>et al.</i> , 2019]	90.1	80.1	93.2	65.5
Distillation[Salehi <i>et al.</i> , 2021]	87.2	-	94.5	87.7
DN2[Reiss <i>et al.</i> , 2021]	92.5	94.1	94.5	86.5
PANDA[Reiss <i>et al.</i> , 2021]	96.2	94.1	95.6	86.5
CAP(no adaptation)	96.6	96.0	95.1	89.4
CAP	97.0	96.5	95.5	89.9

Table 1: Anomaly detection performance (Average AUROC %).

However, L_S would optimize G_θ into a sparse matrix tending to zero which greatly reduces the representation capacity of model, since with such solution, the loss error tends to be minimum.

To avoid such a trivial solution, based on such framework, the adaptation can be constrained by the pretrained space. Specifically, we consider an alignment between pretrained and adaptive spaces. We design G_θ as a square matrix and propose a constraint term

$$\Omega = \frac{1}{N} \sum_{i=1}^N \left(1 - \frac{\mathbf{z}_i \cdot \hat{\mathbf{z}}_i}{\|\mathbf{z}_i\| \|\hat{\mathbf{z}}_i\|} + \|\mathbf{z}_i - \hat{\mathbf{z}}_i\|_2^2 \right) \quad (6)$$

Where we consider the cosine similarity and Euclidean distance of \mathbf{z} and $\hat{\mathbf{z}}_i$ simultaneously. This encourages $\hat{\mathbf{z}}_i$ to be not only close to the \mathbf{z} in terms of Euclidean distance but also be in the same direction. Such constraint, on the one hand, prevents the solution from approaching zero matrix, on the other hand, retains the pretrained information to a certain extent. Combining the aforementioned loss and constraint, L_{total} is formulated as a weighted sum of L_S and Ω

$$L_{total} = L_S + \lambda \Omega \quad (7)$$

For a given test image, we can naturally define an anomaly score by the similarity between its adaptive mapping $\hat{\mathbf{z}}_i$ and adaptive normal representation $\hat{\mathbf{z}}_i^n$

$$s(\mathbf{x}_i) = 1 - \frac{\hat{\mathbf{z}}_i \cdot \hat{\mathbf{z}}_i^n}{\|\hat{\mathbf{z}}_i\| \|\hat{\mathbf{z}}_i^n\|} \quad (8)$$

4 Experiments

In this section, we experimentally evaluate the performance of CAP and use AUROC as an evaluation metric. Four datasets including CIFAR-10, CIFAR-100, FMNIST and MvTec are used as the benchmarks. ResNet152 and WideResNet50 pretrained on ImageNet is leveraged as the pretrained network which is the same as PANDA[Reiss *et al.*, 2021] and the output of its fourth block is regarded as the pretrained feature. Detailed experiment configurations are shown in Appendix B.

4.1 Comparison with State-of-the-art

We present anomaly detection performance of CAP compared to the state-of-the-art deep learning based methods: Deep-SVDD [Ruff *et al.*, 2018], MHRot [Hendrycks *et al.*, 2019], Distillation[Salehi *et al.*, 2021], DN2 and

PANDA[Reiss *et al.*, 2021]. Deep-SVDD is the widely accepted deep-learning based method, MHRot is a high-performance self-supervised method, Distillation, DN2 and PANDA are pretrain based methods. All the results that are available in the original papers are copied exactly. To verify the effectiveness of CAP, we additionally compare the results without adaptation whose anomaly score is the similarity between feature of a specific input and mean value of its K-normal nearest features in pretrained space.

The performances in different datasets are reported in Table 1. We can make the following observations from this table: i) Proposed anomaly detection criterion achieves a significant improvement. Specifically, without adaptation, the anomaly detection performance is better than DN2 which is also a baseline pretrained based method without adaptation. ii) Feature adaptation of a CAP improves the performance. For all datasets, the detection results are about 0.5% higher than method without adaptation. iii) CAP prominently outperforms the state-of-the-art pretrained based method PANDA (achieving around 0.8%, 2.4% and 3.4% higher on the CIFAR-10, CIFAR-100 and MvTec datasets) while for FMNIST, CAP performs basically the same. Therefore, it demonstrates the effectiveness of the proposed adaptation strategy and anomaly criterion both on semantic AD and sensory AD.

4.2 Design of Adaptive Projection Head

For structural design of the simple adaptive projection head, we refer to [Chen *et al.*, 2020]. Concretely, the layer number and nonlinear function would greatly affect the performance. We compare performance with different structure on CIFAR-10, CIFAR-100 and MvTec datasets.

The performances of different head designs are reported in Table 2. The result shows that for anomaly detection, the simpler the structure of projection is, the more effective the adaptation would be. In the traditional deep learning tasks, deep model has a certain generalization capability, while for anomaly detection, generalization is disadvantageous where we hope feature adaptation only works in normal data. Thus, in all experiments, the one-layer linear mapping is leveraged as the adaptation function.

4.3 Effect of Reformed Self-attention

Self-attention mechanism aims to make the normal representation taking inner-relationship into account. We visualize the effect of reformed self-attention module in Figure 2 where

Dataset	L+ReLU+L	L+ReLU	L	attention	AUROC
CIFAR-10	✓			✓	96.4
		✓		✓	96.8
			✓	✓	97.0
			✓		96.8
CIFAR-100	✓			✓	95.2
		✓		✓	96.3
			✓	✓	96.5
			✓		96.2
MvTec	✓			✓	89.5
		✓		✓	89.7
			✓	✓	89.9
			✓		89.8

Table 2: Anomaly detection performance with different CAP structure design.


 Figure 2: The visualization of effect of reformed self-attention. The left image is a specific input. The middle ones: **top** the K-normal nearest images in the pretrained space. **middle** the weight without reformed self-attention. **bottom** the weight with reformed self-attention.

a specific input image, its K-normal nearest images in pretrained space and the weights with or without self-attention are displayed. Moreover, performance with reformed self-attention or not are shown in Table 2.

Traditionally, the weights of different items are often obtained by calculating the similarity with the input. We believe the K-normal nearest neighbors contain inner-relationship in one-class setting and the input feature of a specific image should be inferred through such inner-relationship. In Figure 2, the first weight bar illustrates that weights without self-attention are approximately same while the second bar shows the weights have obvious discrimination. Moreover, the visualization result implies that CAP have considered semantic inner-information. Specifically, The items with larger weight are semantically closest to the input (black cars), while the items with smaller weight are different (cars with other colors). Appendix B provides more case studies.

Anomaly detection performance shown in Table 2 demonstrates that reformed self-attention module slightly improved adaptation effect. However, such module visually brings more interpretability. It can be concluded that the considering the inner-relationship is reasonable and this module is effective.

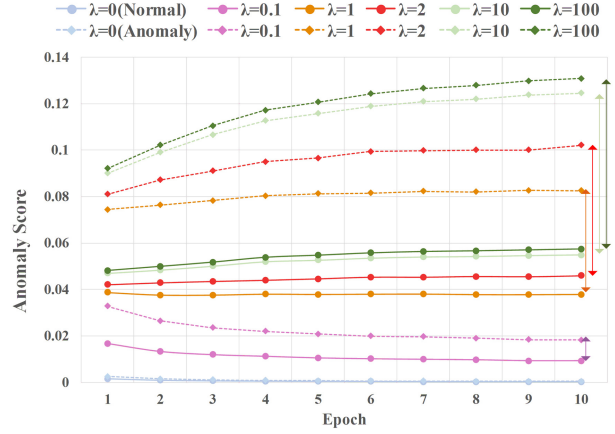
4.4 Comparison with Different k Number

K-normal nearest neighbors is an important module in CAP, in which the number of nearest normal pretrained features should be analyzed. We show performance with different k

K	1	4	8	16	32	64
CIFAR-10	96.6	96.7	96.8	96.9	97.0	97.0
CIFAR-100	96.0	96.1	96.3	96.4	96.5	96.4
MvTec	89.5	89.9	89.3	88.8	88.1	86.7

Table 3: Anomaly detection performance with different K-normal nearest neighbors (Average AUROC %).

λ	0	0.1	1	2	10	100	no ada
AUROC	95.1	95.6	96.9	97.0	97.0	97.0	96.6

 Table 4: Anomaly detection performance on CIFAR-10 with different weights λ of constraint term.

 Figure 3: Anomaly score of normal (solid line) and anomaly (dashed line) in test set when training on CIFAR-10 class 0 with different λ .

on CIFAR-10, CIFAR-100 and MvTec in Table 3.

It can be seen that the performance will be significantly improved when considering multiple pretrained features. For CIFAR-10 and CIFAR-100, a larger k is preferred while for MvTec, k = 4 has obvious advantages. The reason for that is the number of normal training data in MvTec is small where the normal representation of all inputs will be similar if k is large, and resulting in a decrease in the specificity of the adaptive normal representation for each input. There are enough training normal data in CIFAR-10 and CIFAR-100, thus naturally, a larger k promises better performance.

4.5 Benefit of Constraint Term

We have mentioned a potential trivial solution and proposed a constraint term to avoid such situation. To verify the necessity of the constraint term Ω , we compare the performance with different weights λ on CIFAR-10 which is shown in Table 4.

When $\lambda = 0$, the performance has decreased greatly which shows the potential risk indeed exists. when λ is relatively small, the performance has slightly improved, however, it is still inferior to performance without adaptation. While for larger λ , the model can adapt well, and we find that when the constraint weight reaches a certain value, the result tends to be stable.

To show the benefit of constraint more intuitively, we show detailed anomaly score of normal and anomaly when training

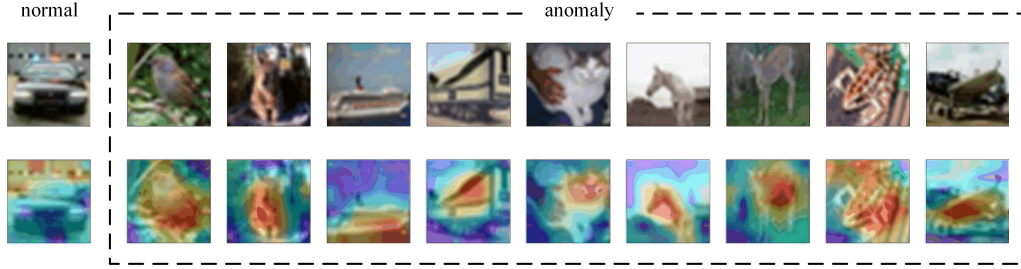


Figure 4: Visualization of CIFAR-10 (Semantic AD). The inferred anomaly region of normal image is in the inessential area and that of anomaly image is on the semantic target.

on class 0 in Figure 3. For $\lambda = 0$, the anomaly score of all samples are basically equal to 0 which implies the similarity between $\hat{\mathbf{z}}_i$ and $\hat{\mathbf{z}}_i^n$ is equal to 1 and the model would indeed tend to learn a trivial solution when there is no constraint. For $\lambda = 0.1$, there is a gap gradually being small between normal and anomaly which implies an unsatisfactory adaptation. For $\lambda = 1$ and $\lambda = 2$, it is shown that the gap of anomaly score between normal and anomaly is gradually growing and a gratifying adaptation appears. The anomaly score of normal is basically unchanged while the anomaly score of anomaly has increased. For larger λ like 10 and 100, even though the gap is distinct, the anomaly score of normal rises obviously signifying over-fitting happens on normal sample (See Appendix B for more case studies). These results prove the necessity of constraint term and a slight alignment is promising.

4.6 Visualization

To show the mechanism of CAP more intuitively, we visualize results referring to class activation mapping(CAM)[Zhou *et al.*, 2016]. In CAM, the product of a global average pooling feature \mathbf{z} and the last layer feature map \mathbf{f} reveals that through which pixels does the model know that the picture belongs to a category. In CAP, \mathbf{z}_i is the global average pooling, thus we compute the element-wise cosine similarity between \mathbf{z}_i and its corresponding $7*7$ pretrained feature map. We also compute the element-wise cosine similarity between $\hat{\mathbf{z}}_i^n$ and the same feature map. These two calculation results are upsampled to $224*224$. Naturally, the absolute difference between the two can reflect the anomaly region. We show the visualization results in Figure 4 and Figure 5. The dark yellow area represents the anomaly region CAP infers.

Semantic AD. Semantic AD only focuses on the semantic shift. We show the anomaly region on the normal class and anomaly classes. In the normal class, the anomaly region is some edge pixels and for anomaly classes, the anomaly region is accurately located on the semantic target. This result illustrates that adaptive normal representation obtained by CAP contains the semantic information of the normal class. Meanwhile, this result also shows that the proposed anomaly detection criterion is reasonable and convincing.

Sensory AD. Sensory AD only focuses on objects with the same or similar semantics and identifies the observational differences on their surface. We select some classes in MvTec and show the calculated anomaly region and the ground truth in Figure 5. The visualization results show that even only the

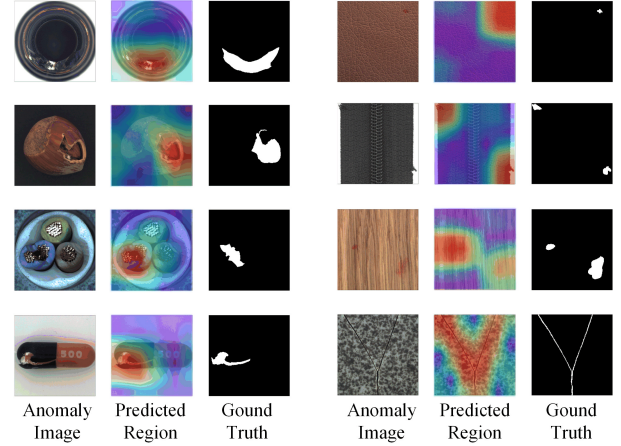


Figure 5: Visualization of MvTec (Sensory AD): With only semantic feature and low resolution feature map, CAP can locate the anomaly area.

difference between the two semantic features, \mathbf{z}_i and $\hat{\mathbf{z}}_i^n$, is considered, the anomaly region can be located. For anomaly with large region, the localization performance is better than that with smaller region, and the reason for this is that CAP does not leverage decoder mechanism to learn an original size output and the upsampling is through interpolation, thus the inferred anomaly region will be generally larger than authentic region. However, these anomaly regions can all be located more or less with only encodes which implies CAP is an effective and interpretative adaptation strategy for sensory AD task.

5 Conclusion

In this paper, we propose a simple yet effective anomaly detection framework. CAP abandons the global optimization goal in traditional one-class setting but focuses on a specific image and its corresponding normal representation. A novel loss function including a simple constraint term is proposed to effectively detect anomalies and avoid pattern collapse. We conduct experiments on multiple benchmarks to verify the effectiveness of CAP, and prove the rationality of each module in the framework through diverse ablation experiments. Adding decoder mechanism to CAP to locate anomaly more precisely is our future research interest.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (no. 61872063), Shenzhen Science and Technology Program, China (no. JCYJ20210324140407021)

References

- [Angiulli and Pizzuti, 2002] Fabrizio Angiulli and Clara Pizzuti. Fast outlier detection in high dimensional spaces. In *European conference on principles of data mining and knowledge discovery*, pages 15–27. Springer, 2002.
- [Bergman et al., 2020] Liron Bergman, Niv Cohen, and Yedid Hoshen. Deep nearest neighbor anomaly detection. *arXiv preprint arXiv:2002.10445*, 2020.
- [Bergmann et al., 2020] 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.
- [Chen et al., 2020] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [Ding and Kolaczyk, 2013] Qi Ding and Eric D Kolaczyk. A compressed pca subspace method for anomaly detection in high-dimensional data. *IEEE Transactions on Information Theory*, 59(11):7419–7433, 2013.
- [Dong et al., 2021] Jiahua Dong, Yang Cong, Gan Sun, Zhen Fang, and Zhengming Ding. Where and how to transfer: knowledge aggregation-induced transferability perception for unsupervised domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [Glodek et al., 2013] Michael Glodek, Martin Schels, and Friedhelm Schwenker. Ensemble gaussian mixture models for probability density estimation. *Computational statistics*, 28(1):127–138, 2013.
- [Gong et al., 2019] Dong Gong, Lingqiao Liu, Vuong Le, Budhaditya Saha, Moussa Reda Mansour, Svetha Venkatesh, and Anton van den Hengel. Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1705–1714, 2019.
- [Hendrycks et al., 2019] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. *arXiv preprint arXiv:1906.12340*, 2019.
- [Perera et al., 2019] Pramuditha Perera, Ramesh Nallapati, and Bing Xiang. Ocgan: One-class novelty detection using gans with constrained latent representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2898–2906, 2019.
- [Reiss et al., 2021] Tal Reiss, Niv Cohen, Liron Bergman, and Yedid Hoshen. Panda: Adapting pretrained features for anomaly detection and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2806–2814, 2021.
- [Ruff et al., 2018] Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *International conference on machine learning*, pages 4393–4402. PMLR, 2018.
- [Ruff et al., 2019] Lukas Ruff, Robert A Vandermeulen, Nico Görnitz, Alexander Binder, Emmanuel Müller, Klaus-Robert Müller, and Marius Kloft. Deep semi-supervised anomaly detection. *arXiv preprint arXiv:1906.02694*, 2019.
- [Ruff et al., 2021] Lukas Ruff, Jacob R Kauffmann, Robert A Vandermeulen, Grégoire Montavon, Wojciech Samek, Marius Kloft, Thomas G Dietterich, and Klaus-Robert Müller. A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, 2021.
- [Salehi et al., 2021] Mohammadreza Salehi, Niousha Sadjadi, Soroosh Baselizadeh, 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.
- [Schölkopf et al., 2001] Bernhard Schölkopf, John C Platt, John Shawe-Taylor, Alex J Smola, and Robert C Williamson. Estimating the support of a high-dimensional distribution. *Neural computation*, 13(7):1443–1471, 2001.
- [Su et al., 2019] Yanzhou Su, Yongjian Wu, Min Wang, Feng Wang, and Jian Cheng. Semantic segmentation of high resolution remote sensing image based on batch-attention mechanism. In *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, pages 3856–3859. IEEE, 2019.
- [Tolstikhin et al., 2021] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Peter Steiner, Daniel Keysers, Jakob Uszkoreit, et al. Mlp-mixer: An all-mlp architecture for vision. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [Vaswani et al., 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [Yang et al., 2021] Jingkan Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*, 2021.
- [Zhou et al., 2016] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.