

Mean-Shifted Contrastive Loss for Anomaly Detection

Tal Reiss, Yedid Hoshen

The Hebrew University of Jerusalem

Abstract

Deep anomaly detection methods learn representations that separate between normal and anomalous images. Although self-supervised representation learning is commonly used, small dataset sizes limit its effectiveness. It was previously shown that utilizing external, generic datasets (e.g. ImageNet classification) can significantly improve anomaly detection performance. One approach is outlier exposure, which fails when the external datasets do not resemble the anomalies. We take the approach of transferring representations pre-trained on external datasets for anomaly detection. Anomaly detection performance can be significantly improved by fine-tuning the pre-trained representations on the normal training images. In this paper, we first demonstrate and analyze that contrastive learning, the most popular self-supervised learning paradigm cannot be naively applied to pre-trained features. The reason is that pre-trained feature initialization causes poor conditioning for standard contrastive objectives, resulting in bad optimization dynamics. Based on our analysis, we provide a modified contrastive objective, the Mean-Shifted Contrastive Loss. Our method is highly effective and achieves a new state-of-the-art anomaly detection performance including 98.6% ROC-AUC on the CIFAR-10 dataset.

1 Introduction

Anomaly detection is a fundamental task for intelligent agents and has broad applications for scientific and industrial tasks. Due to the significance of the task, many efforts have been focused on automatic anomaly detection, particularly on statistical and machine learning methods. A common paradigm used by many anomaly detection methods is estimating the probability density function (PDF) of the data, given a training set of normal samples. New samples are then predicted to be normal if they are likely under the estimator, while low-likelihood samples are predicted to be anomalous. The quality of the density estimators is closely related to the quality of features used to represent the data. Classical methods used k -means, k -nearest-neighbors (k NN), or Gaussian mixture models (GMMs) on raw features, however, this often results in poor estimators on high-dimensional data such as images.

One way to improve PDF estimators for high-dimensional data is by representing them using effective features. Many

recent methods use self-supervised features to detect anomalies. Unfortunately, anomaly detection datasets are typically small and do not include labeled anomalous samples, resulting in weak features. To overcome this issue, many previous methods suggested incorporating the use of generic, external datasets (e.g. ImageNet). As such supervision is often found off-the-shelf and does not require any further annotation, this does not present extra costs. External datasets have been used in two major ways. The first is outlier exposure (OE), which amounts to using an external dataset to simulate anomalies. Supervised techniques can then be used for anomaly detection. Although very effective, OE can completely fail when the anomalies are more similar to the normal data than to the OE dataset.

The alternative approach, that we will follow here, advocates for pre-training representations on generic, external datasets and transferring them to the anomaly detection task. (Reiss et al. 2021) found that even a simple k NN anomaly detection classifier based on ImageNet pre-trained representation already outperforms nearly all self-supervised methods. Furthermore, fine-tuning the pre-trained features on the normal training data can result in significant performance improvements. Although it may appear natural that this can simply be done by initializing standard anomaly detection techniques with the pre-trained features, it is quite challenging. PANDA (Reiss et al. 2021) combined the DeepSVDD objective (Ruff et al. 2018) with pre-trained features. However, as contrastive learning approaches typically perform better than the SVDD objective, we hypothesize that combining pre-trained features with contrastive methods would achieve the best of both worlds.

We begin with the surprising result that standard contrastive methods, initialized with pre-trained weights, do not improve anomaly detection accuracy at all. An analysis of the learning dynamics reveals that this occurs as the standard contrastive loss is poorly suited for data that are concentrated in a compact subspace (which the normal data under strong pre-trained features are). We propose an alternative objective, the Mean-Shifted Contrastive (MSC) loss. The MSC loss is found to achieve better One-Class Classification (OCC) performance than the center loss (used in DeepSVDD and PANDA) and sets a new anomaly detection state-of-the-art. **Our contributions:**

1. We analyze the standard contrastive loss for fine-tuning

pre-trained representations for OCC and show that it is poorly initialized and achieves poor performance.

2. Proposing an alternative objective, the *Mean-Shifted Contrastive Loss*, and analyzing its importance for adapting features for anomaly detection.
3. Extensive experiments demonstrating the state-of-the-art anomaly detection performance of our method (e.g. 98.6% ROC-AUC on CIFAR-10).

2 Related Work

Classical anomaly detection methods: Traditional AD methods follow three main paradigms: (i) Reconstruction e.g. principal component analysis (PCA) (Jolliffe 2011) and K nearest neighbors (k NN) (Eskin et al. 2002). (ii) Density estimation e.g. Gaussian mixture models (EGMM) (Glodek, Schels, and Schwenker 2013), and kernel density estimation (Latecki, Lazarevic, and Pokrajac 2007). (iii) One-class classification e.g. one-class support vector machine (OC-SVM) (Scholkopf et al. 2000) and support vector data description (SVDD) (Tax and Duin 2004).

Self-supervised deep learning methods: Much research was performed on learning from unlabeled data. Methods typically operate by devising an auxiliary task with automatically generated artificial labels for each image and then training a deep neural network using standard supervised techniques. Tasks include: video frame prediction (Mathieu, Couprie, and LeCun 2016), image colorization (Zhang, Isola, and Efros 2016; Larsson, Maire, and Shakhnarovich 2016), puzzle solving (Noroozi and Favaro 2016), rotation prediction (Gidaris, Singh, and Komodakis 2018). The latter was adapted by (Golan and El-Yaniv 2018; Hendrycks et al. 2019; Bergman and Hoshen 2020) for anomaly detection in image and tabular data. Most relevant to this work is contrastive learning (Chen et al. 2020), which learns representations by distinguishing similar views of the same sample from other data samples. Variants of contrastive learning were also introduced to OCC. CSI (Tack et al. 2020) treats augmented input as positive samples and the distributionally-shifted input as negative samples. DROC (Sohn et al. 2020) shares a similar technical formulation as CSI without ensembling or test-time augmentation.

Feature adaptation for one-class classification: Hand-crafted or externally-learned representations are often sub-optimal for AD. Instead, OCC methods can be initialized using pre-trained features but then adapted on OCC objectives to improve their AD accuracy. DeepSVDD (Ruff et al. 2018) pre-trains a representation encoder by autoencoding on the normal data. Several works use ImageNet pre-trained features e.g. (Hendrycks et al. 2019; Perera and Patel 2019; Reiss et al. 2021), achieving much better results. The features are then adapted to OCC - often by the SVDD objective (e.g. DeepSVDD (Ruff et al. 2018) and PANDA (Reiss et al. 2021)). Adaptation often encounters catastrophic collapse. DeepSVDD tackles this by incorporating architectural constraints. PANDA proposed a simple early stopping approach or a more principled continual learning approach based on EWC (Kirkpatrick et al. 2017).

3 Background: Learning Representations for One-Class Classification

3.1 Preliminaries

In the one-class classification task, we are given a set of training samples that are all normal (and contain no anomalies) $x_1, x_2, \dots, x_N \in \mathcal{X}_{train}$. The objective is to classify a new sample x as being normal or anomalous. The methods considered here learn a deep representation of a sample parameterized by the neural network function $\phi : \mathcal{X} \rightarrow \mathbb{R}^d$, where $d \in \mathbb{N}$ is the feature dimension. In several methods, ϕ is initialized by pre-trained weights ϕ_0 , which can be learned either using external datasets (e.g. ImageNet classification) or using self-supervised tasks on the training set. The representation is further tuned on the training data to form the final representation ϕ . Finally, an anomaly scoring function $s(\phi(x))$ determines the anomaly score of a sample x . The binary anomaly classification can be predicted by applying a threshold on $s(x)$. In Sec. 3.2 and Sec. 3.3, we review the most relevant methods for learning the representation ϕ .

3.2 Self-supervised Objectives for OCC

We review two relevant self-supervised objectives:

Center Loss: This loss uses the simple idea, that features should be learned so that normal data lie within a compact region of feature space, whereas anomalous data lie outside it. As we focus on the OCC setting, there are no examples of anomalies in training. Instead, the center loss encourages the features of the normal samples to lie as near as possible to a predetermined center. Specifically, the center loss for an input sample $x \in \mathcal{X}_{train}$ can be written as follows:

$$\mathcal{L}_{center}(x) = \|\phi(x) - c\|^2 \quad (1)$$

This objective suffers from a trivial solution - the features $\phi(x)$ collapse to a singular point c for all samples, normal and anomalous. This is often called "catastrophic collapse". Such a collapsed representation cannot, of course, discriminate between normal and anomalous samples.

Contrastive Loss: Recently, contrastive learning was responsible for much progress in self-supervised representation learning (Chen et al. 2020). In the contrastive training procedure a mini-batch of size B is randomly sampled and the contrastive prediction task is defined on pairs of augmented examples derived from the mini-batch, resulting in $2B$ data points. For OCC, the contrastive objective simply states that: (i) the angular distance between the features of any positive pair (x_i, x_{i+B}) should be small (ii) the distance between the features of normal samples $x_i, x_m \in \mathcal{X}_{train}$, where $i \neq m$, should be large. The typical contrastive loss for a positive pair (x_i, x_{i+B}) where x_i and x_{i+B} are augmentations of some $x \in \mathcal{X}_{train}$, denoted by $\mathcal{L}_{con}(x_i, x_{i+B})$, is written below:

$$-\log \frac{\exp(\text{sim}(\phi(x_i), \phi(x_{i+B}))/\tau)}{\sum_{m=1}^{2B} \mathbb{1}[i \neq m] \cdot \exp(\text{sim}(\phi(x_i), \phi(x_m))/\tau)} \quad (2)$$

where $\forall m \in [1, 2B] : x_m$ is an augmented view of some $x \in \mathcal{X}_{train}$, τ denotes a temperature hyper-parameter and sim is the cosine similarity. Contrastive methods currently achieve the top performance for OCC without the utilization of externally trained network weights.

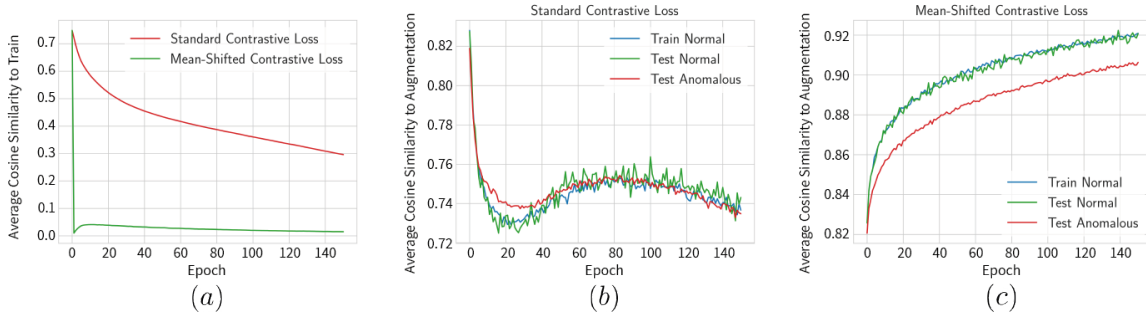


Figure 1: CIFAR-10 "Airplane" class. Average cosine similarity between features on training set vs. training epoch. (a) Similarity between pairs of images. Similarity between images and their augmentation for (b) Contrastive objective (c) MSC objective.

3.3 Initialization with Pre-trained Weights

Self-supervised representation learning methods have high sample complexity and in many cases do not outperform supervised representation learning methods. It is common practice in deep learning to transfer the weights of classifiers pre-trained on large, somewhat related, labeled datasets to the task of interest. Previous methods used pre-trained weights for anomaly detection (Perera and Patel 2019; Reiss et al. 2021). It was found that fine-tuning the pre-trained weights of ϕ_0 on the normal data, results in a stronger feature extractor ϕ . The latest approach, PANDA, used the center loss (Eq. 1) for fine-tuning the pre-trained weights. Several attractive properties of methods based on ImageNet pre-trained features were established: (i) they outperform self-supervised anomaly detection methods by a wide margin, without using any labeled examples of anomalies or OE. (ii) they generalize to datasets that are very different from ImageNet including aerial and medical.

As contrastive objectives typically perform better than the center loss, it is natural to assume that replacing PANDA's center loss with the contrastive loss would be advantageous. Unfortunately, the representation collapses immediately and this modification achieves poor OCC results. In Sec. 4 we will analyze this phenomenon and present an alternative objective that overcomes this issue.

4 Modifying the Contrastive Loss for Anomaly Detection

In this section, we introduce our new approach for OCC feature adaptation. In Sec. 4.1 we analyze the mechanism that prevents standard contrastive objectives from benefiting from pre-trained weights for OCC. In Sec. 4.2 we present our new objective function, the mean-shifted contrastive (MSC) loss. In Sec. 4.3 we analyze the proposed mean-shifted contrastive loss for OCC transfer learning.

4.1 Adaptation Failure of The Contrastive Loss

While contrastive methods have achieved state-of-the-art performance on visual recognition tasks, they were not designed for feature adaptation for OCC. Here, we report and analyze a surprising result: when optimizing a contrastive

loss for OCC using ImageNet pre-trained features, the representations not only fail to improve but degrade quickly.

To understand this phenomenon, we present in Fig. 1 plots of two metrics as a function of training epoch: (i) uniformity: the average cosine similarity between the features of pairs of examples in the training set (more uniform = close to zero) (ii) augmentation distance: the average cosine similarity between features of train samples and their augmentation (higher generally means better ordering of feature space). (Wang and Isola 2020) showed the contrastive loss optimizes two properties (i) uniform distribution of $\{\phi(x)\}_{x \in \mathcal{X}_{train}}$ across the unit sphere. (ii) different augmentations of the same image map to the same representation. We can see that contrastive training improved the uniformity of the distribution but failed to increase the similarity between the features of two views of the same image. Results for other temperatures are in the SM. This shows that contrastive training did not make features more discriminative, suggesting the training objective is not well specified.

We provide an intuitive explanation for the empirical observation. Commonly, the normal data occupy a compact region in the ImageNet pre-trained feature space. When viewed in the spherical coordinate system having its center at the origin, normal images span only a small, bounded region of the sphere. As one of the objectives of contrastive learning is to have features that occupy the entire sphere, the optimization would be focused on changing the features accordingly, putting far less emphasis on improving the features so that they are invariant to augmentations. This is bad for anomaly detection as this uniformity makes anomalies harder to detect (as they become less likely to occupy a sparse region of the feature space). Additionally, such drastic changes in the features cause the loss of the useful properties of the pre-trained feature space. This is counter to the objective of transferring strong auxiliary features.

4.2 The Mean-Shifted Contrastive Loss for Better Adaptation

To overcome the limitations of contrastive learning explained above, we propose a simple modification of its objective for OCC feature adaptation. In our modified objective, we compute the angles between the features of images with respect to the **center** of the normal features rather than




Figure 2: *Top*: The angular representation in relation to the origin. \mathcal{L}_{con} enlarging the angles between positive and negative samples, thus increasing their Euclidean distance to c . *Bottom*: The mean-shifted representation. \mathcal{L}_{msc} does not affect the Euclidean distance between c and the mean-shifted representations while maximizing the angles between the negative pairs.

the origin (as done in the original contrastive loss). Although this can be seen as a simple shift of the original objective, we will show that it resolves the critical issues highlighted above and allows contrastive learning to benefit from the powerful, pre-trained feature initialization (See Sec. 4.3). We name this new objective, the *Mean-Shifted Contrastive*.

Let us denote the center of the normalized feature representations of the training set by c :

$$c = \mathbb{E}_{x \in \mathcal{X}_{train}} \left[\frac{\phi_0(x)}{\|\phi_0(x)\|} \right] \quad (3)$$

where ϕ_0 is the initialized pre-trained model. For each image x , we create two different augmentations of the image, denoted x_i, x_{i+B} . All the augmented images are first passed through a feature extractor ϕ . They are then scaled to the unit sphere by ℓ_2 normalization (see Sec. 5.2 for the motivation of using ℓ_2 normalization). We mean-shift each representation, by subtracting the center c from each normalized feature representation. The MSC loss for two augmentations (x_i, x_{i+B}) of some image $x \in \mathcal{X}_{train}$ from an augmented mini-batch of size $2B$, is defined as follows:

$$\mathcal{L}_{msc}(x_i, x_{i+B}) = -\log \left(\frac{\exp(\text{sim}(\frac{\phi(x_i)}{\|\phi(x_i)\|} - c, \frac{\phi(x_{i+B})}{\|\phi(x_{i+B})\|} - c)/\tau)}{\sum_{i=1}^{2B} \mathbb{1}_{[i \neq m]} \cdot \exp(\text{sim}(\frac{\phi(x_i)}{\|\phi(x_i)\|} - c, \frac{\phi(x_m)}{\|\phi(x_m)\|} - c)/\tau)} \right) \quad (4)$$

where τ denotes a temperature hyper-parameter and sim is the cosine similarity.

Anomaly criterion: To classify a sample as normal or anomalous, we use the cosine similarity from a set of K suitably selected training exemplars $N_k(x)$. The set $N_k(x)$ can be selected by k nearest-neighbors (more accurate) or k -means (faster). We compute the cosine similarity between the features of the target image x and the k exemplars $N_k(x)$. The anomaly score is given by:

$$s(x) = \sum_{\phi(y) \in N_k(x)} 1 - \text{sim}(\phi(x), \phi(y)) \quad (5)$$

where sim is the cosine similarity. By checking if the anomaly score $s(x)$ is larger than a threshold, we determine if the image x is normal or anomalous. See Sec. 5.2 for a comparison between different exemplar selection methods.

4.3 Understanding the Mean-Shifted Loss

Uniformity: Optimizing pre-trained weights with the standard contrastive loss focuses on optimizing uniformity around the origin-centered sphere but hurts feature semantic similarity (Sec. 4.1). The MSC loss proposes a simple but very effective solution - evaluating uniformity in the coordinate frame around the data center. As the features in this frame are already roughly uniform, the optimization focuses on improving their semantic similarity. As shown in Fig. 1, the features are uniform right from initialization with our objective (low cosine similarity between normal examples). Thus, the optimization can focus on improving the features.

Compactness around center: The standard contrastive loss maximizes the angles between representations of negative pairs even when they are both normal training images. By maximizing these angles, the distance to the center increases as well, as illustrated in Fig. 2 (top). This behavior is in contrast to the optimization of the center loss (Eq. 1), which learns representations by minimizing the Euclidean distance between normal representations and the center. (Reiss et al. 2021) showed that optimizing the center loss results in high anomaly detection performance. Our proposed loss does not suffer from this issue. Instead of measuring the angular distance between samples in relation to the origin, we measure the angular distance in relation to the center of the normal features. As can be seen in Fig. 2 (bottom), our proposed loss maximizes the angles between the negative pairs while preserving their distance to the center.

In Fig. 3.b-c we present histograms of the angular distance between images to the center measured around the origin with (i) standard contrastive loss and (ii) our MSC loss. The distributions of normal and anomalous features overlap in standard contrastive, but not in our MSC loss.

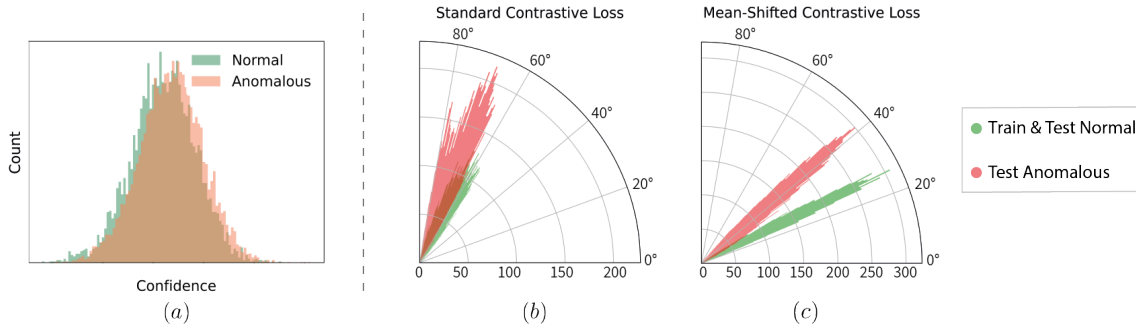


Figure 3: CIFAR-10 "Bird" class. *Left:* (a) Confidence histogram. The ℓ_2 norm confidence of the extracted features derived by ϕ does not differentiate between normal and anomalous samples. *Right:* Histograms of the angular distance between images to the center measured around the origin for (b) Standard Contrastive objective (c) The Mean-shifted contrastive objective.

5 Experiments

In this section, we extensively evaluate our method. In Sec. 5.1, we report our OCC results with a comparison to previous works on the standard benchmark datasets. In Sec. 5.2 we further analyze our objective and we present an ablation study. Building upon the framework suggested in (Reiss et al. 2021), we use ResNet152 pre-trained on ImageNet classification task as ϕ_0 (unless specified otherwise), and adding an additional final ℓ_2 normalization layer - this is our initialized feature extractor ϕ . By default, we fine-tune our model with \mathcal{L}_{msc} (as in Eq. 4). For inference, we use the criterion described in Sec. 4.2. We adopt the ROC-AUC metric as the detection performance score. Full training and implementation details are in the Supplementary Material¹.

5.1 Main Results

We evaluated our approach on a wide range of anomaly detection benchmarks. Following (Golan and El-Yaniv 2018; Hendrycks et al. 2019), we run our experiments on commonly used datasets: CIFAR-10 (Krizhevsky, Hinton et al. 2009), CIFAR-100 coarse-grained version that consists of 20 classes (Krizhevsky, Hinton et al. 2009), and CatsVsDogs (Elson et al. 2007). Following standard protocol, multi-class datasets are converted to anomaly detection by setting a class as normal and all other classes as anomalies. This is performed for all classes, in practice turning a single dataset with C classes into C datasets. We compare our approach with the top current self-supervised and pre-trained feature adaptation methods (Ruff et al. 2018; Hendrycks et al. 2019; Tack et al. 2020; Sohn et al. 2020; Reiss et al. 2021). Full dataset descriptions and baselines are in the SM.

Tab. 1 shows that our proposed approach surpasses the previous state-of-the-art on the common OCC benchmarks. This establishes the superiority of our approach, resulted from our new objective, over previous self-supervised and pre-trained methods. Full class-wise results are in the SM.

5.2 Further Analysis & Ablation Study

Small datasets. To demonstrate different challenges in image anomaly detection, we further extend our results on

small datasets following the standard protocol. We tested our method on: MVTEC (Bergmann et al. 2019) and DIOR (Li et al. 2020) and the fine-grained version of CIFAR-100 (100 classes). Furthermore, we used the CIFAR-10 dataset with different amounts of training data. In Tab. 2 we present a comparison between (i) the top self-supervised contrastive-learning based method - CSI (ii) the top OCC feature adaptation method - PANDA (iii) our method. We see that the self-supervised method does not perform well on such small datasets, whereas our method achieves very strong performance. The reason for the poor performance of self-supervised methods on small datasets is that they only see the small dataset for training, and cannot learn strong features using such a small sample size. This is particularly severe for contrastive methods (but is also the case for all other self-supervised methods). As pre-trained methods transfer features from external datasets, they perform better.

The Angular Representation. Our initial feature extractor ϕ_0 is pre-trained on a classification task (ImageNet classification). To obtain class probabilities, the features $\phi_0(x)$ are subsequently multiplied by classifier matrix C and passed through a softmax layer. The logits are therefore given by $C\phi_0(x)$. As softmax is a monotonic function, scaling of the logits does not change the order of probabilities. However, scaling does determine the degree of confidence in the decision. We propose to disambiguate the representation $\phi_0(x)$ into two components: (i) the semantic class $\frac{\phi_0(x)}{\|\phi_0(x)\|}$, and (ii) the confidence $\|\phi_0(x)\|$. The confidence acts as a per-sample temperature that determines how confident the discrimination between the classes is. A thorough investigation that we conducted, showed that the confidence of an ImageNet pre-trained feature representation did not help the anomaly detection performance. In Fig. 3.a, we compare the histogram of confidence values between the normal and anomalous values on a particular class of the CIFAR-10 dataset. We observe that confidence does not discriminate between normal and anomalous images in this dataset. In Fig. 4 we demonstrate the sensitivity of the mean-shifted representation to class confidence. This emphasizes the importance of confidence normalization for MSC optimization.

We thus propose to use the angular center loss. The

¹<https://github.com/talreiss/Mean-Shifted-Anomaly-Detection>

Architecture	Self-supervised					Pre-trained	
	LeNet	ResNet18			ResNet152	ResNet152	
Method	DeepSVDD	MRot	DROC	CSI	\mathcal{L}_{msc}	PANDA	\mathcal{L}_{msc}
CIFAR-10	64.8	90.1	92.5	94.3	94.8	96.2	97.2
CIFAR-100	67.0	80.1	86.5	89.6	94.4	94.1	96.4
CatsVsDogs	50.5	86.0	89.6	86.3	98.4	97.3	99.3

Table 1: Anomaly detection performance (mean ROC-AUC%). Best in bold.

Dataset	CSI	PANDA	\mathcal{L}_{msc} (ResNet18)	\mathcal{L}_{msc} (ResNet152)
DIOR	78.5	94.3	97.5	97.7
MvTec	63.6	86.5	85.0	87.2
CIFAR-100 (Fine-grained)	90.1	97.1	92.0	97.6
CIFAR-10 (200 samples)	81.3	95.4	93.1	96.5
CIFAR-10 (500 samples)	88.1	95.6	93.8	96.7

Table 2: Anomaly detection accuracy (mean ROC-AUC%) on small dataset. Best in bold.

DN2		PANDA		\mathcal{L}_{msc}	$\mathcal{L}_{msc} + \mathcal{L}_{ang}$
Raw	Ang.	\mathcal{L}_{center}	\mathcal{L}_{ang}		
92.5	95.8	96.2	96.8	97.2	97.5

Table 3: Ablation study (CIFAR-10, mean ROC-AUC%).

angular center loss encourages the angular distance between each sample and the center to be minimal. This contrasts with the standard center loss (used by PANDA and DeepSVDD), which uses the Euclidean distance. Although a simple change, the angular center loss achieves much better results than the regular center loss (see Tab. 3).

$$\mathcal{L}_{ang} = -\phi(x) \cdot c \quad (6)$$

Training objective. An ablation of the objectives and of DN2 (k NN on unadapted ImageNet pre-trained ResNet features) is presented in Tab. 3. Note that both the confidence-invariant form of DN2 and PANDA outperform their Euclidean versions. We further notice that the MSC loss outperforms the rest, and combining it with the angular center loss results in further improvements.

Optimization from scratch. The mean-shifted objective assumes that the relative distance to the center of the features is correlated with high detection performance. When initializing the center as a random Gaussian vector we lose this strong prior, as a result, the detection capabilities are drastically degraded. Therefore when training a model from scratch without any strong initialization that comes from a pre-trained model, our objective does not improve over standard contrastive losses. The MSC loss is therefore a directed contribution to anomaly detection from pre-trained features.

Do ImageNet pre-trained features extend to distant domains? Our results on DIOR and MVTEC which are significantly different from ImageNet provide such evidence. This was also extensively established by (Reiss et al. 2021).

Method	ResNet	EfficientNet	DenseNet	ViT
DN2	92.5	89.3	85.6	95.7
PANDA	96.2	95.3	82.4	95.8
\mathcal{L}_{msc}	97.2	97.0	95.7	98.6

Table 4: Performance gains with different network architectures (CIFAR-10, mean ROC-AUC%). Best in bold.

Catastrophic collapse & Early Stopping. Like other OCC pre-trained feature adaptation methods (e.g. PANDA), our method suffers from catastrophic collapse after a very large number of training epochs. However, our method is less sensitive than PANDA, as we dominate PANDA at any point in the curve and collapse much more slowly. See SM.

Why do self-supervised OCC models not suffer from catastrophic collapse? Pre-trained methods start from highly discriminative features and can therefore lose accuracy whereas self-supervised features start from random features and therefore have nothing to forget.

Rotation prediction methods do not benefit from pre-trained features. The best performing self-supervised approaches (Tack et al. 2020; Sohn et al. 2020) use a combination of contrastive and rotation prediction objectives. Although initially counter-intuitive, AD by rotation prediction does not benefit from pre-trained features. The reason is that features that generalize better (e.g. ImageNet pre-trained), achieve better performance on rotation prediction for *both* normal and anomalous data. Pre-training, therefore, decreases the gap between the performance of normal and anomalous images on rotation prediction in comparison to randomly-initialized networks. As this gap is used for discriminating between normal and anomalous samples, its decrease leads to worse anomaly detection performance. Specifically, we found that CSI with ImageNet pre-trained features achieves 89.5% on CIFAR-10 compared to the stan-

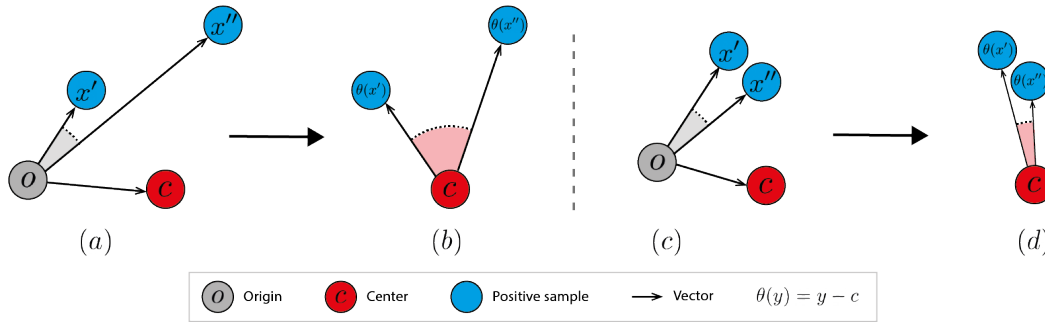


Figure 4: Sensitivity of \mathcal{L}_{msc} to class confidence. (a) The angular representation in relation to the origin without confidence normalization. (b) The mean-shifted representation enlarges the angle between the positive samples. (c) The angular representation after confidence normalization. (d) The angle between the positive samples is approximately preserved after mean-shifting.

DN2		PANDA		SSL		Ours
Raw	Ang.	\mathcal{L}_{center}	\mathcal{L}_{ang}	MRot	CSI	\mathcal{L}_{msc}
76.2	80.4	78.5	78.0	76.7	79.0	85.3

Table 5: Multi-Class AD (CIFAR-10, mean ROC-AUC%).

standard version which results in 94.3%.

Self-supervised methods do not benefit from large architectures. Pre-trained models can use large deep networks, a quality that self-supervised OCC methods lack as OCC benchmarks are not large. We tested this by evaluating CSI with different ResNet backbone sizes (ResNet18, 50, 152). The CSI results were the same for all backbone sizes 94.3% ROC-AUC on CIFAR-10. This contrasts with pre-trained feature adaptation in our method which benefits from bigger backbones and also outperforms CSI on the same-sized backbone (ResNet18).

Performance on different network architectures. In Tab. 4 we provide the CIFAR-10 ROC-AUC% results of DN2, PANDA and our method on leading CNN and ViT architectures pre-trained on ImageNet. PANDA is sensitive to the choice of architecture and did not improve results on ViT and collapsed on DenseNet. Our MSC loss generalizes across architectures and results in significant performance gains, including 98.6% ROC-AUC on CIFAR-10.

Multi-Class Anomaly Detection. We evaluate the setting introduced by (Ahmed and Courville 2020) where the normal data contains multiple semantic classes. Note that *no class labels are provided* for the normal data. This setting is more challenging than the single-class setting as the normal data has a multi-modal distribution. For each experiment, we denoted a single CIFAR-10 class as anomalous and other CIFAR-10 classes as normal. We report the mean ROC-AUC% over the 10 experiments in Tab. 5. PANDA provides no improvement over DN2 (with cosine distance) as the data is no longer uni-modal and therefore not compact. In contrast, our MSC loss does not rely on the uni-modal assumption to the same extent and produces better results than previous pre-trained and self-supervised methods.

Pre-training and Outlier Exposure (OE). AD methods

$k = 1$	$k = 5$	$k = 10$	$k = 100$	Full train set
94.2	95.8	96.1	97.0	97.2

Table 6: CIFAR-10 OCC with k -means (mean ROC-AUC%)

employ different levels of supervision. The most extensive supervision is used by OE (Deecke et al. 2021; Ruff et al. 2020), which requires a large external dataset at training time, and requires the dataset to be from a similar domain to the anomalies. OE methods fail when the OE is more similar to the normal than the anomalous data have different distributions, as anomalous data are more likely to be classified as normal than OE. This can be overcome by using supervision to select an OE dataset that simulates the anomalies, but this contradicts the objective of anomaly detection, the detection of any type of anomaly. Pre-training, like OE, is also achieved through an external labeled dataset, but differently from OE, the labels are not required to be related to the anomalies and the external dataset is only required once - at the pre-training stage and is never used again.

Detection scoring functions. k NN has well-established approximations that mitigate its inference time complexity. A simple, but effective solution is reducing the set of gallery samples via k -means. In Tab. 6 we present a comparison of the performance of our method and its k -means approximations with the features of the normal training images compressed using different numbers of means (k). We use the sum of the distances to the nearest neighbor means as the anomaly score. We can see that significant inference time improvement can be achieved for a small loss in accuracy.

6 Conclusion

This paper investigated feature adaptation methods for anomaly detection. We first analyzed the standard contrastive loss and found that it provides a poor initialization for OCC feature adaptation. We then introduced the Mean-Shifted Contrastive loss, which overcomes the limitations of the standard contrastive loss. Extensive experiments verified the outstanding performance of our method.

Acknowledgments

We are grateful to Shira Bar-On for making our lovely figures. Tal Reiss was funded by grants from the Israeli Cyber Directorate, the Israeli Council for Higher Education, and a Facebook Research Award.

References

- Ahmed, F.; and Courville, A. 2020. Detecting semantic anomalies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 3154–3162.
- Bergman, L.; and Hoshen, Y. 2020. Classification-Based Anomaly Detection for General Data. In *ICLR*.
- Bergmann, P.; Fauser, M.; Sattlegger, D.; and Steger, C. 2019. MVTEC AD—A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9592–9600.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*.
- Deecke, L.; Ruff, L.; Vandermeulen, R. A.; and Bilen, H. 2021. Transfer-based semantic anomaly detection. In *International Conference on Machine Learning*, 2546–2558. PMLR.
- Elson, J.; Douceur, J. R.; Howell, J.; and Saul, J. 2007. Asirra: a CAPTCHA that exploits interest-aligned manual image categorization. In *ACM Conference on Computer and Communications Security*, volume 7, 366–374.
- Eskin, E.; Arnold, A.; Prerau, M.; Portnoy, L.; and Stolfo, S. 2002. A geometric framework for unsupervised anomaly detection. In *Applications of data mining in computer security*, 77–101. Springer.
- Gidaris, S.; Singh, P.; and Komodakis, N. 2018. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*.
- Glodek, M.; Schels, M.; and Schwenker, F. 2013. Ensemble Gaussian mixture models for probability density estimation. *Computational Statistics*, 28(1): 127–138.
- Golan, I.; and El-Yaniv, R. 2018. Deep Anomaly Detection Using Geometric Transformations. In *NeurIPS*.
- Hendrycks, D.; Mazeika, M.; Kadavath, S.; and Song, D. 2019. Using self-supervised learning can improve model robustness and uncertainty. In *NeurIPS*.
- Jolliffe, I. 2011. *Principal component analysis*. Springer.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13): 3521–3526.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. *Citeseer*.
- Larsson, G.; Maire, M.; and Shakhnarovich, G. 2016. Learning representations for automatic colorization. In *ECCV*.
- Latecki, L. J.; Lazarevic, A.; and Pokrajac, D. 2007. Outlier detection with kernel density functions. In *International Workshop on Machine Learning and Data Mining in Pattern Recognition*, 61–75. Springer.
- Li, K.; Wan, G.; Cheng, G.; Meng, L.; and Han, J. 2020. Object detection in optical remote sensing images: A survey and a new benchmark. *ISPRS Journal of Photogrammetry and Remote Sensing*, 159: 296–307.
- Mathieu, M.; Couprie, C.; and LeCun, Y. 2016. Deep multi-scale video prediction beyond mean square error. *ICLR*.
- Noroozi, M.; and Favaro, P. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*.
- Perera, P.; and Patel, V. M. 2019. Learning deep features for one-class classification. *IEEE Transactions on Image Processing*, 28(11): 5450–5463.
- Reiss, T.; Cohen, N.; Bergman, L.; and Hoshen, Y. 2021. PANDA: Adapting Pretrained Features for Anomaly Detection and Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2806–2814.
- Ruff, L.; Gornitz, N.; Deecke, L.; Siddiqui, S. A.; Vandermeulen, R.; Binder, A.; Müller, E.; and Kloft, M. 2018. Deep one-class classification. In *ICML*.
- Ruff, L.; Vandermeulen, R. A.; Franks, B. J.; Müller, K.-R.; and Kloft, M. 2020. Rethinking assumptions in deep anomaly detection. *arXiv preprint arXiv:2006.00339*.
- Scholkopf, B.; Williamson, R. C.; Smola, A. J.; Shawe-Taylor, J.; and Platt, J. C. 2000. Support vector method for novelty detection. In *NIPS*.
- Sohn, K.; Li, C.-L.; Yoon, J.; Jin, M.; and Pfister, T. 2020. Learning and Evaluating Representations for Deep One-class Classification. *arXiv preprint arXiv:2011.02578*.
- Tack, J.; Mo, S.; Jeong, J.; and Shin, J. 2020. Csi: Novelty detection via contrastive learning on distributionally shifted instances. *NeurIPS*.
- Tax, D. M.; and Duin, R. P. 2004. Support vector data description. *Machine learning*.
- Wang, T.; and Isola, P. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, 9929–9939. PMLR.
- Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful image colorization. In *ECCV*.