

Removing Geometric Bias in One-Class Anomaly Detection with Adaptive Feature Perturbation

Romain Hermary Vincent Gaudillière Abd El Rahman Shabayek Djamila Aouada

University of Luxembourg, Esch-sur-Alzette, Luxembourg

{romain.hermary, vincent.gaudilliere, abdelrahman.shabayek, djamila.aouada}@uni.lu

Abstract

One-class anomaly detection aims to detect objects that do not belong to a predefined normal class. In practice, training data lack those anomalous samples; hence, state-of-the-art methods are trained to discriminate between normal and synthetically-generated pseudo-anomalous data. Most methods use data augmentation techniques on normal images to simulate anomalies. However, the best-performing ones implicitly leverage a geometric bias present in the benchmarking datasets. This limits their usability in more general conditions. Others are relying on basic noising schemes that may be suboptimal in capturing the underlying structure of normal data. In addition, most still favour the image domain to generate pseudo-anomalies, training models end-to-end from only the normal class and overlooking richer representations of the information. To overcome these limitations, we consider frozen, yet rich feature spaces given by pretrained models, and create pseudo-anomalous features with a novel adaptive linear feature perturbation technique. It adapts the noise distribution to each sample, applies decaying linear perturbations to feature vectors, and further guides the classification process using a contrastive learning objective. Experimental evaluation conducted on both standard and geometric bias-free datasets demonstrates the superiority of our approach with respect to comparable baselines. The codebase is accessible via our [public repository](#).

1. Introduction

Anomaly detection (AD), also known as novelty or out-of-distribution detection, is a widely investigated research topic, with applications ranging from machine faults detection [31, 40], to malicious transactions in banking [6] and hazardous environmental situations in autonomous driving [3]. In most cases, abnormal samples are too costly to obtain, and only normal samples are available. This AD problem is therefore unsupervised and is also referred to as

unlabelled one-class anomaly detection [45, 47], one-class novelty detection [29], or semantic outlier detection [26]. In this paper, we tackle the problem from an image-level perspective, *i.e.*, given images from a single semantic class, being able to classify unseen images as belonging to that class (*normal* samples) or not (*anomalies*). This differs conceptually from anomaly segmentation, addressed in fields such as medical diagnosis [12] or industrial quality control [2] and concerned with *localisation* of pixel-level abnormalities within one semantic class, requiring dedicated heuristics and generally heavier computations.

State-of-the-art (SotA) one-class anomaly detection approaches can be classified into three main categories: reconstruction-based methods [8, 10, 34, 37, 48, 51], embedding-based methods [16, 20, 35, 38, 46] and synthetic anomaly-based methods [1, 7, 15, 18, 22, 39, 44, 45, 47]. Reconstruction-based methods consist of training a deep neural network to reconstruct normal images and using the reconstruction error as a metric to differentiate between normal and abnormal samples. However, such methods rely on the assumption that the network can accurately reconstruct training data but cannot generalise beyond it, which is not necessarily true in practice. Embedding-based methods aim to embed normal feature distribution into a compressed space to facilitate the distinction with abnormal data. However, these methods are subject to feature collapse and require dedicated heuristics to prevent this from happening.

Synthetic anomaly-based methods consist mostly in generating artificial anomalies, *i.e.* *pseudo-anomalies*, from normal data during training to learn discriminative features able to separate normal and abnormal data. These methods, whose development has been driven by the rise of self-supervised image representation learning [9, 14], have achieved among the highest performances on standard benchmarks such as CIFAR-10 [24]. In practice, these methods generate pseudo-anomalous images using data augmentation techniques, and results from the literature show that the most effective type of augmentation is image rotation. This is to the extent that, to the best of our




Figure 1. Illustration of the geometric bias. *first row*: original images sampled from 2 public datasets. *left column*: representative of CIFAR-10 [24] classes, second row is a manually rotated version of the top image and differs significantly from it. Rotated images can be treated as anomalies, since the original data is *biased towards specific orientations*. *right column*: second row contains another original image found in the dataset, which *exhibits natural rotations* compared to the top image. In the absence of orientation bias, rotations cannot be used to create pseudo-anomalies.

knowledge, all of the best-performing methods take advantage of it [1, 15, 18, 22, 32, 39, 44, 45, 47]. However, we believe that the effectiveness of such an image transform comes from the fact that original images in standard benchmarks are biased towards specific orientations (animals or cars are rarely pictured upside down). This is supported by Goyal *et al.* [16] who noticed that these types of methods are “*heavily domain dependent*”, and that even within the same dataset the “*suitability of a transformation varies based on the structure of the typical points*”, making some transformations very powerful in some cases but ineffective in others. In addition, Gidaris *et al.* [14] found that human-captured image are subject to *well-posedness*: always being in the same orientation, making “*the rotation recognition task well defined*”. To the misfortune of SotA methods, this geometric bias is not omnipresent and therefore cannot be leveraged, for example, in the case of spacecraft datasets [33]. This intuition is illustrated in Figure 1.

Instead of leveraging geometric bias in the images, SimpleNet [29] creates pseudo-anomalies at the feature level; we argue that such an approach mitigates the risk of not being generalisable beyond standard anomaly detection datasets. Nevertheless, this method is from the other branch of AD, localisation, it depends on locally applied perturbations that do not allow for relevant semantic anomalies to

be generated, and, especially, relies on fixed Gaussian noise for every sample, which refrains from any adaptation to the likely complex structure of the representation space. The latter limitation was anticipated with the method PLAD [7]. It proposes to automatically adapt the noise distribution parameters to every sample using a Variational Auto-Encoder (VAE) [23]. However, such a method applies multiplicative and additive noise at the image level, and we argue that (i) deep learning features obtained from pretrained models represent a more optimal encoding of useful semantic information in images, making it easier to bound normality while also limiting the risks of overfitting to the normal class, and that (ii) applying a *linear* noising process is better suited to perturb feature vectors.

Therefore, we propose PLUME, *Perturbation Learning with Unified inforMation Embeddings*, an unsupervised anomaly detection method that generates adaptive pseudo-anomalies within the feature space during training. This is achieved by introducing a linear disturbance to the normal representations, automatically modulating this disturbance with an adaptive noise level. Then, a multilayer perceptron is used to learn a decision boundary between normal and pseudo-anomalous features. It is further guided by a contrastive learning objective to aggregate normal features in a more unified representation and repel abnormal ones. In a nutshell, our contributions are three-fold:

- we propose to work in a frozen feature space and develop an adaptive linear feature perturbation technique to create pseudo-anomalies from normal samples without exploiting any dataset-specific geometric bias,
- we propose PLUME, a model that integrates the aforementioned feature perturbator to learn to detect anomalies in an unsupervised manner with the help of a contrastive learning objective,
- we provide an experimental validation that demonstrates the superiority of our method with respect to comparable baselines on both standard and geometric bias-free datasets.

The rest of the paper is organised as follows. Related work is discussed in Section 2, and our approach is described in Section 3. Section 4 presents experimental validation against state-of-the-art baselines, along with an ablation study. Specific aspects of our method are further discussed in Section 5, and Section 6 concludes the paper.

2. Related Work

2.1. Reconstruction & Embedding-based Methods

Reconstruction-based methods [8, 10, 34, 37, 48, 51] exploit the notion that anomalous image regions deviate significantly from the patterns observed in training data, making their faithful reconstruction challenging. These methods leverage generative models such as auto-encoders (AEs) or generative-adversarial networks (GANs). They learn

a compressed representation and reconstruct normal data from it. Deviations from this learnt representation are flagged as anomalies. However, learning to generate the entire normal data distribution from a finite training set can be challenging and inaccurate in practice (see Table 1a).

Embedding-based methods for AD embed features extracted from normal data into a lower-dimensional space. This compression allows for easier identification of anomalies, which supposedly lie far away from the cluster formed by normal features. Historically, OCSVM [38] uses kernel SVM to separate the normal data from the origin, considering it as the only negative data point. Support Vector Data Description (SVDD) [46] also uses kernel SVM to find an hypersphere that encloses the normal data. However, such shallow methods struggle on complex domains like images, where feature-engineering is quite challenging. Deep-SVDD [35] is a deep learning-based version of SVDD; it minimizes the volume of the hypersphere that encloses the normal features. However, it may suffer from representation collapse. HRN [20] tackles this issue by a holistic regularization method. It constrains the model training to consider the normal features holistically. Such a heuristic is, however, insufficient in practice (Table 1a).

2.2. Synthetic Anomaly-based Methods

Synthesising-based methods adopt a different approach, with the aim of learning how to differentiate normal data from synthesised anomalies.

Several works [1, 15, 18] aim to learn the observed geometric characteristic of the normal data (*e.g.*, orientation of the object) by applying specific transformations (*e.g.*, rotations) to the input image and learning to predict the parameters of the applied transformation. The intuition is that, at test time, a failure to accurately predict the transform likely comes from geometric image properties that are different from those of normal data, *i.e.* are from anomalies. DROC [44] and CSI [45] leverage distributional augmentation (*e.g.*, rotation) to simulate real-world outliers and model the inlier distribution by contrasting original samples with these simulated outliers. SSD [39] also leverages contrastive learning. However, the learnt representation is uniformly distributed on the hypersphere, contradicting the core principle of AD, which suggests that the inlier distribution should remain compact against outliers. To avoid this, UniConHA [47] proposes a dedicated “unilaterally aggregated” contrastive learning objective.

Nevertheless, to be successful all these previous works critically rely on side-information in the form of appropriate transformations to generate pseudo-anomalies, which are dataset-dependent. As the major SotA methods in one-class AD suggest, rotation transformations are relevant to create pseudo-anomalies on the standard benchmark.

In PLAD [7], pseudo-anomalies are generated by applying perturbations on the normal images, using multiplica-

tive and additive noise whose parameters are learnt to be sample-specific. A discriminator made of fully convolutional layers is then simultaneously trained to learn a tight boundary around normal data. This generic method is therefore applicable to any dataset, since it does not rely on any specific bias such as uniform object orientation repeated across normal images.

However, in the era of foundation models [4, 21, 36], where machines can perceive and understand the visual world with unprecedented accuracy, we argue that pre-trained features are more compact and structured representations of image content than raw pixel information and are therefore more suitable supports for perturbation-based anomaly synthesis. Nevertheless, we demonstrate in what follows that additive and multiplicative noise, even adaptive, is not effective on deep features (Section 4.5); therefore, we introduce with PLUME a new adaptive linear perturbation methodology that makes better use of this vector space, achieving new SotA results. In addition, our method enables the reduction of the trained discriminator to a small multilayer perceptron, which can be seamlessly integrated into pre-existing architectures –sharing the feature extraction costs–, while maintaining strong AD performance through the use of a contrastive learning objective.

3. Methodology

In this section, we detail our approach towards unsupervised semantic anomaly detection. We consider a general data space $\mathbb{X} \subset \mathbb{R}^D$, where D is any dimension, in which data points can either follow or defy an agreed normality. We denote by \mathbb{X}^+ the subspace of normality and by $\mathbb{X}^- = \mathbb{X} \setminus \mathbb{X}^+$ the subspace of anomalies. The training set $\mathcal{X} \subset \mathbb{X}$ is partitioned into $\mathcal{X}^+ \subset \mathbb{X}^+$ (normal samples) and $\mathcal{X}^- \subset \mathbb{X}^-$ (anomalies). In this paper, we place ourselves in the extreme case where no anomalous sample is available during training, meaning our solution is designed for applications where $\mathcal{X}^- = \emptyset$ (unsupervised anomaly detection).

3.1. Overview

To detect divergent samples, we train a classifier to predict whether the sample comes from \mathbb{X}^+ or \mathbb{X}^- . We focus on overcoming the lack of negative samples during the classifier training and improving existing fitting techniques. More precisely, we leverage the use of a pre-trained feature extractor to work on an improved vectorised representation of the data, *i.e.* features, and adapt our method to be effective within the considered vector space. In this work, we are therefore considering a model $\mathfrak{h} : \mathbb{R}^{J \times H \times W} \rightarrow \mathbb{X}$ to go from the image space $\mathbb{R}^{J \times H \times W}$ (with J, H, W the number of channels, height and width of the images) to a feature space considered as our data space \mathbb{X} .

The role of the classifier is to derive the complex decision boundary to establish the origin of the samples. With only the positive class available during training, it would be

unnecessary for the network to find a well-fitting boundary around \mathcal{X}^+ to satisfy the classification objective; instead, the modelled solution would ignore the data information, converging towards an infinitely large boundary radius and 100% false negatives when encountering anomalies.

To avoid this degenerate behaviour, we base our method on the union of \mathcal{X}^+ with generated negative samples. To cope with the lack of anomalous samples in the training set \mathbb{X}^+ , we follow the standard paradigm that consists in applying a set of transformations \mathcal{T} on elements of \mathcal{X}^+ to simulate anomalies. The difficulty lies in determining what would be a valid \mathcal{T} for any type of data. Indeed, we have no guarantee that a transformation applied to a normal sample gives a legitimate anomaly, except if we have prior knowledge of both normal and abnormal data.

Following Cai and Fan [7] (PLAD), we propose to adaptively noise the normal data, with the rationale that there is a need for the noise to be adapted to each sample, since they all lie in a different positions in the data space, and that the noise level needs to be small for the perturbed samples to be valid and allow sane convergence. Aiming for an effective generation process, we adopt their adversarial strategy with two models trained jointly.

Figure 2 depicts an overview of our architecture. A first model, called the *perturbator* (purple), generates the parameters for the transformations applied to the normal samples. Eventually, its objective is to model the parameter distribution which will allow correct transformations to be applied. In PLUME, we design \mathcal{T} to be effective in a feature space.

The second model is the classifier (green), for which the training objective is to distinguish between normal samples and pseudo-anomalies. Inspired by contrastive learning strategies, we make the model more efficient in separating samples from different sources. By forcing the classifier to project normal data into a unified representation, we aim at easing the task of deriving the decision boundary.

Section 3.2 further describes the perturbator module and Section 3.3 the optimisations in place for the classifier.

3.2. Feature Perturbator

Inspired by PLAD, a set of pseudo-anomalies $\tilde{\mathcal{X}}$ is dynamically generated from \mathcal{X}^+ using a trainable neural network $\mathbf{g} : \mathbb{R}^D \rightarrow \mathbb{R}^D \times \mathbb{R}^D$, the perturbator. In PLAD, the perturbator generates a couple of vectors $(\alpha_i, \beta_i) = \mathbf{g}(\mathbf{x}_i)$ specific to each normal sample $\mathbf{x}_i \in \mathcal{X}^+$ (vectorised image) and interpreted as element-wise multiplicative and additive noise maps $(\alpha_i \odot \mathbf{x}_i + \beta_i)$. These elements are restrained not to produce pseudo-anomalies lying far away from normal samples –implied by the adversarial objective of the classifier–, by a simple yet effective constraint $\mathcal{L}_n^{(i)}$, defined for the i^{th} sample in the current batch \mathcal{B}^+ as:

$$\mathcal{L}_n^{(i)} = \|\alpha_i - \mathbf{1}\|^2 + \|\beta_i - \mathbf{0}\|^2, \quad (1)$$

where $\mathbf{1} = [1, 1, \dots, 1]^\top$ and $\mathbf{0} = [0, 0, \dots, 0]^\top$ are vectors of dimension D , same as \mathbf{x}_i , α_i and β_i . $\|\cdot\|$ is the ℓ^2 norm.

Although shown to perform properly on images, additive and multiplicative noises are mostly used on raw signal data and we argue that such transformations are not optimal in a feature vector space. To make a more sensible use of the features produced by \mathbf{h} , better preserve the geometry, structure and integrity of the data space while enabling more complex transformations, we propose to redefine the perturbation as a linear map $\mathbf{A}_i \in \mathbb{R}^{D \times D}$. For the perturbation to produce cohesive and non aberrant data, we want this map to represent a small perturbation, hence be close to the identity map. Instead of modifying the perturbator \mathbf{g} to generate directly a matrix in $\mathbb{R}^{D \times D}$, which would heavily impact its number of parameters and disturb its convergence, we propose to keep α_i , β_i (and their constraint $\mathcal{L}_n^{(i)}$), using them ingeniously to instead noise an identity matrix $\mathbf{I} \in \mathbb{R}^{D \times D}$:

$$\mathbf{A}_i = \mathbf{I} + \alpha_i \beta_i^\top. \quad (2)$$

During the training, the perturbation \mathbf{A}_i then converges towards the desired identity map \mathbf{I} , and pseudo-anomalies can be generated by applying linear perturbations:

$$\tilde{\mathbf{x}}_i = \mathbf{A}_i \mathbf{x}_i. \quad (3)$$

The batch \mathcal{B} given to the classifier in any training step is therefore the union of a subset $\mathcal{B}^+ = \{\mathbf{x}_i\}_{i=1}^N \subset \mathcal{X}^+$ of N normal samples and the generated set of pseudo-anomalies $\tilde{\mathcal{B}}^- = \{\tilde{\mathbf{x}}_i\}_{i=1}^N$.

To generate α_i and β_i , we use a Variational Auto-Encoder (VAE) [23] since shown to perform well as the generator \mathbf{g} [7]. It is decomposed into two parts: an encoder and a decoder. The first aims to fit a noise distribution, generating the set of intermediary parameters $\mathcal{P} = \{(\mu_i, \sigma_i)\}_{i=1}^N$, with μ_i and $\sigma_i \in \mathbb{R}^D$ the encoded mean and variance for the i^{th} sample. The second generates the set of transformation parameters $\mathcal{G} = \{(\alpha_i, \beta_i)\}_{i=1}^N$, also respective to each sample. As in the original work of Kingma and Welling [23], the encoder is primarily optimised with a Kullback-Leibler divergence $\mathcal{D}_{KL}^{(i)}$:

$$\mathcal{D}_{KL}^{(i)} = \frac{1}{2} \sum_{d=1}^D [\sigma_i^2 + \mu_i^2 - 1 - \log(\sigma_i^2)]_d, \quad (4)$$

where $[\cdot]_d$ denotes the value at the d^{th} element of the vector. The objective of the decoder parameters is modified and the latter are optimised owing to $\mathcal{L}_n^{(i)}$ (Eq. 1).

3.3. Classifier

We use a neural network $\mathbf{f} : \mathbb{R}^D \rightarrow \mathbb{R}$ to distinguish between normal data and anomalies. As a training objective, we define a target value $y \in \{0, 1\}$ equal to 1 for all $\mathbf{x} \in \mathcal{X}^+$ and 0 for all generated $\tilde{\mathbf{x}}$ passed to \mathbf{f} . Subject to a logistic activation function, any prediction \hat{y} of the classifier lies in the range $[0, 1]$ and a binary cross-entropy loss $\mathcal{L}_{CE}^{(i)}$ is used as a training objective. It is defined as:

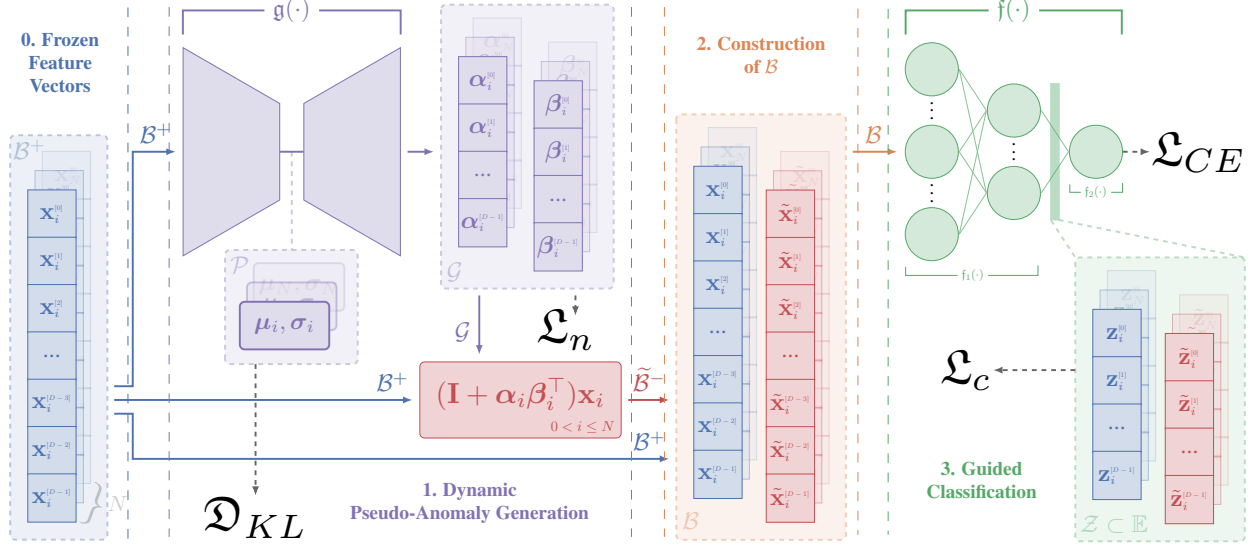


Figure 2. Diagram of our architecture, with the links between the different parts for a batch \mathcal{B}^+ of N vectors of dimension D (features extracted with the backbone h). In purple, the perturbator g generating (α_i, β_i) couples for each sample. In red, the application of linear perturbation to generate $\tilde{\mathcal{B}}$ and construction of \mathcal{B} . In green, the classifier f separated in two parts to access the classifier embeddings, in \mathbb{E} .

$$\mathcal{L}_{CE}^{(i)} = -[y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)], \quad (5)$$

for the i^{th} sample in the current batch \mathcal{B} , and with \hat{y}_i and y_i the predicted and groundtruth sample values, respectively. To facilitate the derivation of the decision boundary, our intuition is that the classifier f should map the normal data to a very dense space, *far* from where anomalies would be projected.

To develop the idea, let us decompose f into f_1 and f_2 so that $f = f_2 \circ f_1$. In our neural network setup, $f_1 : \mathbb{R}^D \rightarrow \mathbb{E}$ represents the first layers of the network that condense the information from the input into an embedding vector $\mathbf{z} \in \mathbb{E}$. Ideally, all vectors from the classifier embedding space \mathbb{E} should be very informative about the origin of the data. On the other hand, we have $f_2 : \mathbb{E} \rightarrow \mathbb{R}$ representing the last layer that combines all the information extracted into the final decision and is thus the boundary estimator. The more f_1 is discriminative, the easier it is for f_2 to model the decision boundary.

In order to increase the separability of the embeddings in \mathbb{E} , depending on their origin, we propose to maximise the similarity of every \mathbf{z} extracted from a normal sample while, at the same time, maximise their dissimilarity with the other embeddings extracted from anomalous samples. As a measure of similarity between the vectors, we follow the literature [9, 19] and leverage the use of temperature-scaled cosine similarity $\mathfrak{s} : \mathbb{E} \times \mathbb{E} \rightarrow [-\frac{1}{\tau}, \frac{1}{\tau}]$:

$$\mathfrak{s}(\mathbf{v}_1, \mathbf{v}_2) = \frac{1}{\tau} \frac{\mathbf{v}_1 \cdot \mathbf{v}_2}{\|\mathbf{v}_1\| \|\mathbf{v}_2\|}, \quad (6)$$

where $\tau \in \mathbb{R}$ helps controlling the concentration level of the similarities distribution [49]. To simultaneously optimise both objectives, we make use of the contrastive loss

defined in Equation (7), where \mathbf{z}_i denotes the embedding vectors derived from normal samples, $\tilde{\mathbf{z}}_i$ the ones from the generated pseudo-anomalies, and $\mathbb{1}_{a \neq b}$ is an indicator function evaluating to 1 if $a \neq b$, 0 otherwise:

$$\mathcal{L}_c^{(i)} = \frac{1}{N-1} \sum_{j=1}^N -\mathbb{1}_{j \neq i} \log \left[\frac{\exp(\mathfrak{s}(\mathbf{z}_i, \mathbf{z}_j))}{\sum_{l=1}^N [\exp(\mathfrak{s}(\mathbf{z}_i, \tilde{\mathbf{z}}_l)) + \mathbb{1}_{l \neq i} \exp(\mathfrak{s}(\mathbf{z}_i, \mathbf{z}_l))]} \right]. \quad (7)$$

Finally, our loss \mathcal{L} , defined in Equation (8), is an average over the batch of the summed sub-losses. λ, ν and $\gamma \in \mathbb{R}$ are tunable hyperparameters used to balance the different losses with respect to \mathcal{L}_{CE} :

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \left(\mathcal{L}_{CE}^{(i)} + \lambda \mathcal{L}_n^{(i)} + \nu \mathcal{D}_{KL}^{(i)} + \gamma \mathcal{L}_c^{(i)} \right). \quad (8)$$

4. Experimental Evaluation

In this section, we delve into the evaluation protocol used and present the results obtained from our experiments. We compare our approach with SotA methods and provide additional experiments to highlight its added value.

4.1. Datasets and Evaluation

Following the main benchmark in unsupervised one-class AD [7, 20, 35], we evaluate our method on CIFAR-10 [24] and report the best AUCs. We produce and give additional results on CIFAR-100 [24] and SPARK [33], the latter being used as a non-geometrically-biased and application-oriented dataset. To the best of our knowledge, our work is the first to exploit SPARK in the context of AD.

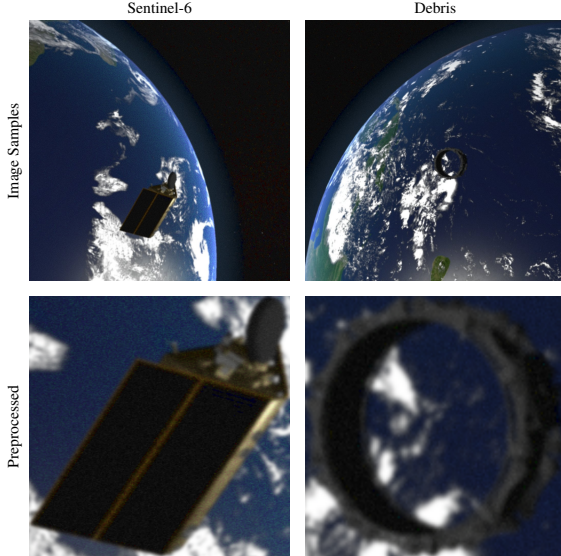


Figure 3. SPARK dataset [33] samples: satellite (left) and debris (right); original (top) and cropped images (bottom).

CIFAR-10 & CIFAR-100 [24]. CIFAR-10 dataset is composed of 10 object classes. For each of them, there are 5,000 training samples. The validation set is composed of 10,000 samples balanced between all classes. Size of images is 32×32 . We evaluate our method following the one-class classification paradigm, *i.e.* each class is alternately considered as the normal class, while the 9 remaining classes form the set of anomalies. CIFAR-100, on the other hand, has 100 object classes grouped into 20 meta-classes. Each meta-class has 2500 samples of 5 different objects. The test set is also made up of 10000 images, and we alternatively consider each meta-class as a normal class and the remaining ones as anomalies.

SPARK [33]. SPARK is a synthetic dataset composed of images of satellites and debris orbiting the Earth. The 10 satellite classes (75,000 training images) are combined as one normal class, while the debris class (5 different debris models) is used during validation as anomaly class. The validation dataset is then composed of 5,000 debris images and 25,000 satellite images. We only use the *RGB* images and discard depth maps. As we focus on the semantics of the objects, we also crop the images to the ground-truth bounding boxes and reshape the resulting images into a square format (512×512) using bilinear interpolation. Sample images are provided in Figure 3.

Metrics. To avoid relying on a specific threshold, the standard evaluation procedure consists in reporting the best Area Under ROC-Curve (AUC) [5] during the validation stage. We thus use this metric in our experiments.

4.2. Baselines

We compare PLUME on CIFAR-10 with 9 baseline methods. To the best of our knowledge, this selection covers




Figure 4. Performance of SotA methods on CIFAR-10 dataset. Methods on the left (blue background) achieve among the best results by exploiting a geometric bias in the dataset, *i.e.* typical object orientations. On the contrary, methods on the right (grey background - green dots) are less dataset-specific and are therefore considered as our baselines. Detailed results of all methods can be found in the supplementary.

most SotA approaches and contains the most accurate ones. We do not consider rotation-based methods, that exploit a bias present in the dataset but not necessarily in real-life scenarios. However, for the sake of transparency and completeness, we have depicted all the results in Figure 4, also including older well-known methods. We also ran PLUME and our principal baseline PLAD [7] on CIFAR-100 and SPARK, and reported the results in Table 1b.

4.3. Implementation Details

We use the Cyclical Learning Rate (CLR) [42, 43] for optimal training phases and setting initial learning rate, in combination with the AdamW [30] optimiser.

As our feature extractor h and unless otherwise specified, we use the widely adopted ResNet50 [17] model pre-trained on ImageNet [11]. This backbone is used as reference, but can be replaced by any other state-of-the-art model. We emphasise the fact that we do not retrain h for any of the experiments. The other elements are trained from scratch. An adaptive pooling layer is then used to bring every feature vector to a size of $D = 3072$, corresponding to the input size of the perturbator. Subsequently, a batch normalisation layer is used to regularise the input space. The batch size is set to $N = 32$. For the classifier, we use a simple multi-layer perceptron of 3 layers. It is important noting that the perturbator is only used during training, and discarded when evaluating the model. More details on the implementation can be found in the supplementary material.

We run each experiment 5 times for 100 epochs and, following the standard protocol in one-class AD [7, 16, 20, 29] which avoids epoch fine-tuning alongside benchmarking, save the best-performing model throughout the training. Following PLAD, we set $\nu = 1$ and tested different values for λ (5, 10 and 20 in our experiments), and found that PLUME is very stable and achieves SotA results whatever

Method	Plane	Car	Bird	Cat	Deer	Dog	Frog	Horse	Ship	Truck	Mean
ADGAN [10]	63.2	52.9	58.0	60.6	60.7	65.9	61.1	63.0	74.4	64.2	62.4
OCGAN [34]	75.7	53.1	64.0	62.0	72.3	62.0	72.3	57.5	82.0	55.4	65.6
TQM [48]	40.7	53.1	41.7	58.2	39.2	62.6	55.1	63.1	48.6	58.7	52.1
DROCC [16]	79.2	74.9	68.3	62.3	70.3	66.1	68.1	71.3	62.3	76.6	69.9
HRN-L2 [20]	80.6	48.2	64.9	57.4	73.3	61.0	74.1	55.5	79.9	71.6	66.7
HRN [20]	77.3	69.9	60.6	64.4	71.5	67.4	77.4	64.9	82.5	77.3	71.3
DPAD [13]	78.0±0.3	75.0±0.2	68.1±0.5	66.7±0.4	77.9±0.8	68.6±0.3	81.2±0.4	74.8±0.2	79.1±1.0	76.1±0.2	74.6
DO2HSC [50]	81.3±0.2	82.7±0.3	71.3±0.4	71.2±1.3	72.9±2.1	72.8±0.2	83.0±0.6	75.5±0.4	84.4±0.5	82.0±0.9	77.7
PLAD [7]	82.5±0.4	80.8±0.9	68.8±1.2	65.2±1.2	71.6±1.1	71.2±1.6	76.4±1.9	73.5±1.0	80.6±1.8	80.5±0.3	75.1
PLUME	89.5±1.0	85.7±1.7	74.5±2.3	78.3±3.6	87.7±1.3	79.5±3.7	87.5±0.9	84.6±3.7	87.8±3.1	90.0±2.7	84.5
PLUME (Max.)	91.0	90.1	80.2	80.8	90.6	83.7	88.4	89.5	92.6	94.0	88.1

(a) Anomaly detection results on CIFAR-10. We report average results over 5 runs for each class, with the best results per class in bold. Maximum AUC values for each class are reported in the last row (not considered in the comparison). The best performing value of λ was selected for each class.

Method	Average AUC	Perturbation	Guidance	Plane	Car	Bird	Cat	Deer	Dog	Frog	Horse	Ship	Truck	Mean
Gaussian	-	-	-	83.3±2.6	75.7±10.0	69.8±3.2	72.8±9.1	80.6±3.4	72.0±3.3	84.5±1.7	79.6±5.7	79.9±3.6	76.9±9.3	77.5
PLAD†	63.8	Gaussian	✓	75.1±3.2	79.6±3.4	68.4±1.3	70.4±1.6	74.5±3.7	69.1±2.6	78.5±2.1	69.1±3.2	75.2±5.0	77.6±5.3	73.7
PLUME	80.3	AddMult	-	66.8±3.5	70.5±9.5	67.4±1.5	64.9±3.6	67.5±5.4	67.0±3.5	70.6±7.8	68.1±7.7	72.5±4.6	68.7±5.2	68.4
		AddMult	✓	59.2±2.6	62.6±2.6	61.2±1.5	58.7±2.0	63.0±1.9	62.6±1.0	68.6±3.0	59.5±4.9	61.1±1.2	57.5±4.6	61.4
LinearMap	-	LinearMap	-	69.5±9.6	75.7±10.1	76.2±1.9	58.8±8.7	74.5±6.7	78.4±2.0	85.1±1.0	78.0±5.4	85.6±3.4	86.5±2.5	76.8
LinearMap	✓ (Mean)	LinearMap	✓	81.4±6.0	72.4±7.2	74.8±3.5	60.3±7.9	74.5±5.5	77.9±1.8	86.4±1.8	82.6±2.7	84.0±3.1	74.3±11.9	76.9
LinearMap	✓	LinearMap	✓	86.8±1.3	85.7±1.7	74.5±2.3	75.0±6.8	87.4±1.7	77.0±2.7	86.1±1.7	84.0±2.6	87.8±3.1	89.6±3.2	83.4

(b) Average AD score on CIFAR-100. We report average results over the 20 meta-classes and 5 runs ($\lambda = 5$). We report our results (†) with the available implementation of PLAD [7].

(c) Ablation study over the perturbation method and the contrastive guidance. Experiments on CIFAR-10 (†): ResNet50, $\lambda = 5$, 5 runs). We report results with adaptive additive and multiplicative noise (AddMult), adaptive linear mapping (LinearMap) and random Gaussian noise (Gaussian); with (✓) and without (-) the contrastive loss, and results with the contrastive loss with respect to the mean embedding vectors of normal and pseudo-anomalous data (Mean).

Table 1. Results on (CIFAR-10 Table 1a), CIFAR-100 (Table 1b) and ablation study (Table 1c).

the value. Detailed performance for each λ can be found in the supplementary. We experimentally found $\tau = 0.5$ and $\gamma = 1$ providing sufficient performance.

4.4. Anomaly Detection on CIFAR-10 & 100

Table 1a presents a comprehensive evaluation of our proposed AD method PLUME, along with comparable existing approaches on the CIFAR-10 dataset. We first notice that PLUME consistently outperforms all methods in terms of mean AUC, showcasing its superior AD capabilities. For example, while OCGAN achieves a mean AUC of 65.6%, DROCC 69.9%, DO2HSC 77.7% and PLAD 75.1%, PLUME significantly surpasses them with a mean AUC of 84.5%. Furthermore, our method also demonstrates consistent superiority in individual categories, achieving highest AUC on each class. Additionally, the results on CIFAR-100 (Table 1b) expose the higher scalability of PLUME to more heterogeneous normal classes and fewer training samples, without modification or additional training cost, compared to PLAD. The detailed results can be found in the supplementary material. These results underscore the substantial advancement in AD performance offered by PLUME.

4.5. Ablation Study

We report the impact of PLUME components in Table 1c. ResNet50 features were used and hyperparameters are the same for the entire study ($\lambda = 5$). We first replaced the linear mapping of PLUME by random Gaussian

noise (Gaussian). Without contrastive loss, the model achieves an average AUC comparable to the current best unbiased baseline, DO2HSC [50]. These results emphasise on the importance of better representations, here given by a ResNet50 trained on ImageNet. However, the decrease in performance with the addition of contrastive guidance (almost 4pp of average AUC) reinforces the idea of unfitting pseudo-anomaly generation, apparently conflicting with the structure of the vector space. Although the use of perturbations in the feature space seems very efficient, the simplicity and randomness do not allow to create sensible anomalous vectors that can be efficiently contrasted with respect to normal vectors.

Adaptive perturbations would allow to create meaningful pseudo-anomalous vectors, which in turn would profit the classifier and derive a better understanding of the normal data, and a better decision boundary. However, basic additive and multiplicative noise (AddMult) performs poorly on vector representations of the data and only achieves 68.4% average AUC. In addition, it also shows a decrease in performance with the contrastive guidance: AddMult also seems not enough efficient at preserving structure coherence in the vector space and rather confuses the guided network. Finally, our proposed linear mapping (LinearMap) is used and also achieves a performance comparable to PLAD [7] and DO2HSC [50]. The major difference from the other tested perturbation methods is that, thanks to its greater expressiveness and richer transformation set, adding contrastive guidance in this case is highly beneficial, offer-

	PLAD [†]	PLUME		
Classifier	LeNet5 [25]	3 Fully-Connected Layers		
Input Data	Raw Images	Features		
		ResNet50 [17]	VGG16 [41]	ConvNeXt [28]
AUC	0.69	0.77	0.74	0.76

Table 2. Anomaly detection results on SPARK. We report the best AUC reached by PLAD, which we ran (†), and PLUME with different frozen backbone architectures. We set $\lambda = 5$. ResNet50, VGG16 and ConvNeXt networks were pretrained on ImageNet.

ing an increase of 6.6pp of average AUC and reaching SotA results even without tuning λ . While using a simpler version of the contrastive loss (Mean), *i.e.* contrasting with respect to the mean embedding vectors of normal and pseudo-anomalous data, could be more computationally efficient (equation provided in the supplementary), it does not bring enough additional knowledge and guidance to the classifier. A completely contrastive guidance (Eq. 7) is necessary.

Overall, the ablation study indicates that (i) the Add-Mult perturbation performs worse than simply adding non-adaptive Gaussian noise in a frozen feature space and is inferior to the LinearMap; (ii) the latter greatly profits from the contrastive guidance and both combined allow PLUME to be successful in deriving a more precise boundary in the feature space, achieving SotA performance; (iii) using better representations, PLUME determines this decision boundary by only training a 3-layers MLP. These results underscore the effectiveness of the proposed methodology. Further qualitative analysis on this structural difference can be found in the supplementary.

4.6. Space Debris Detection on SPARK

SPARK’s complexity is significantly higher compared to CIFAR-10, due to the similarities between normal and abnormal samples (space debris are often parts of spacecrafts) and to the common background information shared by all classes. Moreover, it is important noting that 10 different spacecraft classes are combined in the normal class, adding an additional layer of intricacy to the data. This makes the unsupervised AD task particularly challenging.

The results in Table 2 illustrate that, despite this complexity, our method demonstrates a notable level of accuracy. This is particularly impressive given that our classifier head is relatively simple and can easily fail to derive a proper decision boundary. Moreover, none of the backbones was fine-tuned on datasets similar to SPARK (only pretrained on ImageNet). Therefore, their ability to extract optimal information from SPARK images is not guaranteed.

We trained PLAD, the main baseline, and obtained an AUC of 69%. Comparatively, PLUME performs consistently better with the different feature extractors, reaching a maximum AUC of 77% with ResNet50. The accuracy achieved by PLUME underlines the robustness of our method in dealing with complex datasets and the potential it holds for practical applications.

5. Discussion & Future Work

PLUME has shown versatility and robustness, performing well with standard pre-trained backbones and complex data. The ability to select different backbone architectures enables customisation for specific applications, while its universal design supports its scalability across different domains. The key strengths of our method are its efficiency, seamlessly integrating with existing backbones already used for other tasks like pose or trajectory estimation, and its independence *w.r.t.* data modality or special attributes. This is especially valuable in resource-limited environments, since adding a small classifier for AD significantly optimises resource usage.

While the introduction of linear mapping increases memory costs during training, this is mitigated when pre-extracting features, as PLUME operates on a frozen feature space, also inferring a reduced processing time. The computational overhead at test time is also minimal, since the AD network is a small MLP. However, pre-trained backbones on specialised datasets may be needed for optimal performance in specific cases.

Future improvements could focus on refining our loss function, as the current noise constraint was adapted from previous work and not tailored to our perturbation method. Maximising pairwise similarity between normal and generated samples, instead of constraining noise, could boost performance. Additionally, reducing the current four learning objectives may improve network stability. Investigating extensions to time-series, facing time dependencies, offers another promising direction for future research.

6. Conclusion

We have presented PLUME, a novel approach for one-class anomaly detection. It focuses on generating pseudo-anomalies within the feature space, leveraging a novel adaptive linear feature perturbation technique, without taking advantage of the geometric biases present in some datasets. To guide the classification process, we introduced a contrastive learning objective that enhances the aggregation of normal features and repels abnormal ones. This combination of adaptive feature perturbation and contrastive learning objective contributes to the effectiveness of our approach in distinguishing between normal and abnormal samples. Our extensive experimental evaluation demonstrates the superiority of PLUME over comparable baselines on both standard and geometric bias-free datasets. By outperforming existing methods on both, our approach showcases its potential for real-world applications. We believe it has the potential to contribute significantly to various domains, providing valuable insights and practical solutions for AD challenges.

Acknowledgments. This research was funded in whole or in part by the Luxembourg National Research Fund (FNR), grant reference DEFENCE22/17813724/AUREA.

References

- [1] Liron Bergman and Yedid Hoshen. Classification-based anomaly detection for general data. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. [1](#), [2](#), [3](#)
- [2] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec AD - A comprehensive real-world dataset for unsupervised anomaly detection. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 9592–9600. Computer Vision Foundation / IEEE, 2019. [1](#)
- [3] Daniel Bogdoll, Maximilian Nitsche, and J. Marius Zöllner. Anomaly detection in autonomous driving: A survey. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2022, New Orleans, LA, USA, June 19-20, 2022*, pages 4487–4498. IEEE, 2022. [1](#)
- [4] Florian Bordes, Richard Yuanzhe Pang, Anurag Ajay, Alexander C. Li, Adrien Bardes, Suzanne Petryk, Oscar Mañas, Zhiqiu Lin, Anas Mahmoud, Bargav Jayaraman, Mark Ibrahim, Melissa Hall, Yunyang Xiong, Jonathan Lebensold, Candace Ross, Srihari Jayakumar, Chuan Guo, Diane Bouchacourt, Haider Al-Tahan, Karthik Padthe, Vasu Sharma, Hu Xu, Xiaoqing Ellen Tan, Megan Richards, Samuel Lavoie, Pietro Astolfi, Reyhane Askari Hemmat, Jun Chen, Kushal Tirumala, Rim Assouel, Mazda Moayeri, Arjang Talattof, Kamalika Chaudhuri, Zechun Liu, Xilun Chen, Quentin Garrido, Karen Ullrich, Aishwarya Agrawal, Kate Saenko, Asli Celikyilmaz, and Vikas Chandra. An introduction to vision-language modeling. *CoRR*, abs/2405.17247, 2024. [3](#)
- [5] Andrew P. Bradley. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognit.*, 30(7):1145–1159, 1997. [6](#)
- [6] Eyad Btoush, Xujuan Zhou, Rai Gururaian, Ka Ching Chan, and Xiaohui Tao. A survey on credit card fraud detection techniques in banking industry for cyber security. In *8th International Conference on Behavioral and Social Computing, BESC 2021, Doha, Qatar, October 29-31, 2021*, pages 1–7. IEEE, 2021. [1](#)
- [7] Jinyu Cai and Jicong Fan. Perturbation learning based anomaly detection. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [8] Chengwei Chen, Yuan Xie, Shaohui Lin, Ruizhi Qiao, Jian Zhou, Xin Tan, Yi Zhang, and Lizhuang Ma. Novelty detection via contrastive learning with negative data augmentation. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 606–614. ijcai.org, 2021. [1](#), [2](#)
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR, 2020. [1](#), [5](#)
- [10] Lucas Deecke, Robert A. Vandermeulen, Lukas Ruff, Stephan Mandt, and Marius Kloft. Image anomaly detection with generative adversarial networks. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2018, Dublin, Ireland, September 10-14, 2018, Proceedings, Part I*, volume 11051 of *Lecture Notes in Computer Science*, pages 3–17. Springer, 2018. [1](#), [2](#), [7](#)
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, pages 248–255. IEEE Computer Society, 2009. [6](#)
- [12] Tharindu Fernando, Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. Deep learning for medical anomaly detection - A survey. *ACM Comput. Surv.*, 54(7):141:1–141:37, 2022. [1](#)
- [13] Dazhi Fu, Zhao Zhang, and Jicong Fan. Dense projection for anomaly detection. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, pages 8398–8408. AAAI Press, 2024. [7](#)
- [14] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Un-supervised representation learning by predicting image rotations. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. [1](#), [2](#)
- [15] Izhak Golan and Ran El-Yaniv. Deep anomaly detection using geometric transformations. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 9781–9791, 2018. [1](#), [2](#), [3](#)
- [16] Sachin Goyal, Aditi Raghunathan, Moksh Jain, Harsha Vardhan Simhadri, and Prateek Jain. DROCC: deep robust one-class classification. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 3711–3721. PMLR, 2020. [1](#), [2](#), [6](#), [7](#)
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society, 2016. [6](#), [8](#)
- [18] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS*

- 2019, December 8-14, 2019, Vancouver, BC, Canada, pages 15637–15648, 2019. [1](#), [2](#), [3](#)
- [19] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. *CoRR*, abs/1503.02531, 2015. [5](#)
- [20] Wenpeng Hu, Mengyu Wang, Qi Qin, Jinwen Ma, and Bing Liu. HRN: A holistic approach to one class learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. [1](#), [3](#), [5](#), [6](#), [7](#)
- [21] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 4904–4916. PMLR, 2021. [3](#)
- [22] Ramneet Kaur, Susmit Jha, Anirban Roy, Sangdon Park, Edgar Dobriban, Oleg Sokolsky, and Insup Lee. idecode: Indistribution equivariance for conformal out-of-distribution detection. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 7104–7114. AAAI Press, 2022. [1](#), [2](#)
- [23] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. [2](#), [4](#)
- [24] Alex Krizhevsky. Learning multiple layers of features from tiny images. pages 32–33, 2009. [1](#), [2](#), [5](#), [6](#)
- [25] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86(11):2278–2324, 1998. [8](#)
- [26] Chun-Liang Li, Kihyuk Sohn, Jinsung Yoon, and Tomas Pfister. Cutpaste: Self-supervised learning for anomaly detection and localization. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 9664–9674. Computer Vision Foundation / IEEE, 2021. [1](#)
- [27] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, volume 8693 of *Lecture Notes in Computer Science*, pages 740–755. Springer, 2014. [2](#)
- [28] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 11966–11976. IEEE, 2022. [8](#)
- [29] Zhikang Liu, Yiming Zhou, Yuansheng Xu, and Zilei Wang. SimpNet: A simple network for image anomaly detection and localization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 20402–20411. IEEE, 2023. [1](#), [2](#), [6](#)
- [30] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. [6](#)
- [31] Nesryne Mejri, Laura Lopez-Fuentes, Kankana Roy, Pavel Chernakov, Enjie Ghorbel, and Djamilia Aouada. Unsupervised anomaly detection in time-series: An extensive evaluation and analysis of state-of-the-art methods. *Expert Syst. Appl.*, 256:124922, 2024. [1](#)
- [32] Hossein Mirzaei, Mojtaba Nafez, Mohammad Jafari, Mohammad Bagher Soltani, Mohammad Azizmalayeri, Jafar Habibi, Mohammad Sabokrou, and Mohammad Hossein Rohban. Universal novelty detection through adaptive contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22914–22923, June 2024. [2](#)
- [33] Mohamed Adel Musallam, Vincent Gaudilliere, Enjie Ghorbel, Kassem Al Ismaeil, Marcos Damian Perez, Michel Poucet, and Djamilia Aouada. Spacecraft recognition leveraging knowledge of space environment: Simulator, dataset, competition design and analysis. In *2021 IEEE International Conference on Image Processing Challenges (ICIPC)*, pages 11–15, 2021. [2](#), [5](#), [6](#)
- [34] Pramuditha Perera, Ramesh Nallapati, and Bing Xiang. OCGAN: one-class novelty detection using gans with constrained latent representations. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 2898–2906. Computer Vision Foundation / IEEE, 2019. [1](#), [2](#), [7](#)
- [35] Lukas Ruff, Nico Görnitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Robert A. Vandermeulen, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 4390–4399. PMLR, 2018. [1](#), [3](#), [5](#)
- [36] Sepehr Sameni, Kushal Kafle, Hao Tan, and Simon Jenni. Building vision-language models on solid foundations with masked distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14216–14226, June 2024. [3](#)
- [37] Thomas Schlegl, Philipp Seeböck, Sebastian M. Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *Information Processing in Medical Imaging - 25th International Conference, IPMI 2017, Boone, NC, USA, June 25-30, 2017, Proceedings*, volume 10265 of *Lecture Notes in Computer Science*, pages 146–157. Springer, 2017. [1](#), [2](#)
- [38] Bernhard Schölkopf, John C. Platt, John Shawe-Taylor, Alexander J. Smola, and Robert C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Comput.*, 13(7):1443–1471, 2001. [1](#), [3](#)

- [39] Vikash Sehwal, Mung Chiang, and Prateek Mittal. SSD: A unified framework for self-supervised outlier detection. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [1](#), [2](#), [3](#)
- [40] Oscar Serradilla, Ekhi Zugasti, Jon Rodriguez, and Urko Zurutuza. Deep learning models for predictive maintenance: a survey, comparison, challenges and prospects. *Appl. Intell.*, 52(10):10934–10964, 2022. [1](#)
- [41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [8](#)
- [42] Leslie N. Smith. Cyclical learning rates for training neural networks. In *2017 IEEE Winter Conference on Applications of Computer Vision, WACV 2017, Santa Rosa, CA, USA, March 24-31, 2017*, pages 464–472. IEEE Computer Society, 2017. [6](#)
- [43] Leslie N. Smith and Nicholay Topin. Super-convergence: Very fast training of residual networks using large learning rates. volume abs/1708.07120, 2017. [6](#)
- [44] Kihyuk Sohn, Chun-Liang Li, Jinsung Yoon, Minh Jin, and Tomas Pfister. Learning and evaluating representations for deep one-class classification. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [1](#), [2](#), [3](#)
- [45] Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin. CSI: novelty detection via contrastive learning on distributionally shifted instances. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. [1](#), [2](#), [3](#)
- [46] David M. J. Tax and Robert P. W. Duin. Support vector data description. *Mach. Learn.*, 54(1):45–66, 2004. [1](#), [3](#)
- [47] Guodong Wang, Yunhong Wang, Jie Qin, Dongming Zhang, Xiuguo Bao, and Di Huang. Unilaterally aggregated contrastive learning with hierarchical augmentation for anomaly detection. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 6865–6874. IEEE, 2023. [1](#), [2](#), [3](#)
- [48] Jingjing Wang, Sun Sun, and Yaoliang Yu. Multivariate triangular quantile maps for novelty detection. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5061–5072, 2019. [1](#), [2](#), [7](#)
- [49] Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 3733–3742. Computer Vision Foundation / IEEE Computer Society, 2018. [5](#)
- [50] Yunhe Zhang, Yan Sun, Jinyu Cai, and Jicong Fan. Deep orthogonal hypersphere compression for anomaly detection. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. [7](#)
- [51] Bo Zong, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Dae-ki Cho, and Haifeng Chen. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. [1](#), [2](#)