

# FewSOME: One-Class Few Shot Anomaly Detection with Siamese Networks

Niamh Belton<sup>1,2</sup> Misgina Tsighe Hagos<sup>1,3</sup> Aonghus Lawlor<sup>3,4</sup> Kathleen M. Curran<sup>1,2</sup>

<sup>1</sup>Science Foundation Ireland Centre for Research Training in Machine Learning

<sup>2</sup>School of Medicine, University College Dublin

<sup>3</sup>School of Computer Science, University College Dublin

<sup>4</sup>Insight Centre for Data Analytics, University College Dublin, Dublin, Ireland

{niamh.belton, misgina.hagos}@ucdconnect.ie {aonghus.lawlor, kathleen.curran}@ucd.ie

## Abstract

*Recent Anomaly Detection techniques have progressed the field considerably but at the cost of increasingly complex training pipelines. Such techniques require large amounts of training data, resulting in computationally expensive algorithms that are unsuitable for settings where only a small amount of normal samples are available for training. We propose ‘Few Shot anOMaly dEtECTION’ (FewSOME), a deep One-Class Anomaly Detection algorithm with the ability to accurately detect anomalies having trained on ‘few’ examples of the normal class and no examples of the anomalous class. We describe FewSOME to be of low complexity given its low data requirement and short training time. FewSOME is aided by pretrained weights with an architecture based on Siamese Networks. By means of an ablation study, we demonstrate how our proposed loss, ‘Stop Loss’, improves the robustness of FewSOME. Our experiments demonstrate that FewSOME performs at state-of-the-art level on benchmark datasets MNIST, CIFAR-10, F-MNIST and MVTec AD while training on only 30 normal samples, a minute fraction of the data that existing methods are trained on. Moreover, our experiments show FewSOME to be robust to contaminated datasets. We also report F1 score and balanced accuracy in addition to AUC as a benchmark for future techniques to be compared against. Code available; <https://github.com/niamhbelton/FewSOME>.*

## 1. Introduction

Anomaly Detection (AD) refers to any technique that attempts to detect samples that are substantially distinct from the majority of other samples in a dataset. In recent years, Machine Learning algorithms for AD have become increasingly accurate, but at the cost of increased complexity. Self-

supervised techniques, in particular, have made significant breakthroughs in terms of accuracy. However, most of these methods have complicated training pipelines that involve performing extensive transformations and augmentations to the data during training. Such complex algorithms are not suitable for the specific few-shot setting where only a few examples (i.e. shots) of the normal class exist and no examples of the anomalous class exist. In this paper, we propose ‘Few Shot anOMaly dEtECTION’ (FewSOME), a deep one-class AD algorithm that performs at State-of-the-Art (SOTA) level with a fraction of the complexity of existing methods in terms of training data size and training time. FewSOME has a Siamese-like architecture that consists of multiple branches of neural networks where each neural network has shared weights. Data samples are input in tandem into each neural network during training with the objective of transforming the data samples to representations that are within close proximity of each other. This architecture is suitable for AD given that Siamese Networks are known for making predictions about unknown class distributions [17]. As is common in recent AD techniques, FewSOME benefits from weight initialisation on the ImageNet dataset [7, 8, 25]. FewSOME is trained using our proposed loss, ‘Stop Loss’ denoted as  $L_{stop}$  to prevent ‘Representational Collapse’. Representational Collapse is a known issue in one-class AD tasks where the model learns to map all inputs to a constant output. Our proposed loss, ‘Stop Loss’, can prevent this and improve the robustness of the model.

In this paper, we present FewSOME, a new SOTA in the field of Few Shot Anomaly Detection (FSAD) on normal data samples. In sections 5.1 and 5.2, we show that FewSOME can perform competitively with SOTA techniques that were trained in the classical AD setting (i.e. trained on large datasets) whilst having significantly lower complexity in terms of training data size and training time. Our experiments show that FewSOME is robust to contaminated

datasets (section 5.3) and by means of an ablation study, we show how our proposed loss ‘Stop Loss’ boosts model performance and prevents representational collapse (section 6.1). Additionally, we report F1 score and balanced accuracy in addition to Area Under the Curve (AUC) as a benchmark for future techniques to be compared against.

## 2. Related Work

Traditionally, AD was performed using simple Machine Learning approaches such as Isolation Forest [20] or statistical methods such as Kernel Density Estimation (KDE) [15] and One-Class Support Vector Machine (OC-SVM) [29]. However, the requirement for manual feature engineering and the poor computational scalability of these approaches led to the exploration of more advanced techniques for AD.

**Deep One-Class.** Deep One-Class methods aim to learn the features of normality in a dataset and classify data samples that diverge substantially from the extracted normal features as anomalies. DeepSVDD [26] is one of the most commonly known one-class AD techniques. It trains a network that transforms the training samples to a representation space where normal samples are contained inside a hypersphere. Data samples whose representations fall outside the radius of the hypersphere are considered to be anomalies. The Deep Robust One-Class Classification (DROCC) [12] significantly outperformed DeepSVDD by margins up to 20% on some AD tasks by generating synthetic anomalies on which the model could train on. A more recent method, PANDA [25], shows that SOTA AD performance can be achieved by training large neural network architectures of ResNet-152 and WideResNet-50 with pretrained ImageNet weights [8]. The newly proposed, Interpolated Gaussian Descriptor (IGD) [7] trains a Gaussian classifier, initialised with pretrained weights on the ImageNet dataset, to minimise the distance between representations of normal images and the centre of the normal image distribution whilst also adversarially interpolating training samples. They implement a critic module to constrain the training of IGD to be based on normal samples that are representative of the majority of normal samples, rather than anomalous samples that may be present in the training data.

**Generative.** Deep Autoencoders and Generative Adversarial Networks (GANs) have dominated the field of AD in previous years [1, 11, 23, 28, 32, 37, 38]. Autoencoders learn a mapping from the input to a latent space of lower dimensionality. The images are then reconstructed from the latent space and anomalies are identified by large reconstruction errors. GANs for AD learn to generate normal samples. During testing, the model attempts to match the test sample to a point in the generator’s latent space. It then reconstructs the test sample based on this point and anomalies are identified by large reconstruction errors. Although Autoencoders

and GANs can accurately detect anomalies, they have various limitations; it can be difficult to estimate the intrinsic dimensionality of the data i.e. it requires manual selection of the latent space dimensionality, they do not specifically target anomaly detection, they often have a significant number of network parameters, they require large amounts of training data and it is computationally expensive to reconstruct the entire image.

**Self-Supervised.** Self-Supervised techniques have progressed the AD field considerably in terms of accuracy. GEOM [10], GOAD [4] and RotNet [9] are examples of transformation based methods that apply transforms to the training data such as flipping and rotating. They then train a classifier to predict the transform applied. The training sample is classified as normal if the model can predict the transform and an anomaly otherwise. Another method, Contrasting Shifted Instances (CSI) [34], contrasts each training sample with an augmentation of itself, where the augmentation is considered to be a different class to the original training sample. CSI can therefore, contrastively learn meaningful representations of the normal class. Although these techniques significantly outperform existing methods, recent literature under-reports their limitations such as the requirement for domain specific transformations. For example, augmentations such as flipping an image of class ‘6’ will convert this to class ‘9’ in the handwritten digit dataset, MNIST. Therefore, many of the recent approaches have not tested their technique on the standard AD benchmark MNIST dataset, highlighting their lack of robustness. It was also previously shown that the performance of GEOM significantly decreases if applied to data that has been augmented, whilst the performance of the one-class method DROCC was unaffected [12]. These transformations/augmentations also increase the complexity of the training and testing pipeline with CSI requiring up to 40 augmentations at test time and thus, limiting their application in the real-world.

**Few-Shot Anomaly Detection (FSAD).** DeepSAD [27], an extension of DeepSVDD, was one of the first AD techniques to adapt their loss function and training strategy to consider labelled normal and anomalous data. Later studies, focused on the FSAD setting where there is a large amount of unlabelled samples and few shots of the anomalous class available for training [21, 22, 35, 36]. Liznerski et al. (2022) [21] studied the zero-shot setting where the model trained on all available unlabelled data but no anomalies. They use pretrained weights from the Contrastive Language-Image Pre-training (CLIP) [24] model. CLIP was trained on a large dataset of image and text pairs with the objective of related pairs being within close proximity to each other in representation space and unrelated pairs being distant from each other in representation space. Other works have studied the separate FSAD setting where multiple normal classes

exist in the data and there is only a few shots of each normal class available for training [2, 14]. These techniques aim to not only identify anomalies that are distant from the normal data but they also aim to identify anomalies that exist in-between the normal classes. FewSOME tackles the less studied FSAD setting, that requires no level of supervision, of training on only few examples of the normal class. Recent works have studied this setting for the specific use case of detecting anomalies in industrial settings [3, 31]. They also study the zero-shot setting, where the model has not seen any examples of the normal or anomalous class but detect anomalies by identifying image regions that break the homogeneity of the input data sample. The first work to train on few shots of the normal class for more general AD was the Hierarchical Transformation-Discriminating Generative (HTDG) model [33]. They employ a generative model and use self-supervision to perform AD based on only one, five and ten shots of the normal data. Unlike FewSOME, they compare their performance to competing methods trained only in the few shot setting, whilst we show in sections 5.1 and 5.2 that FewSOME’s performance surpasses that of HTDG and FewSOME also performs at the same level of existing SOTA that were trained in the classical AD setting (i.e. trained on large datasets).

**Representational Collapse.** Techniques with examples of the anomalous class or techniques that synthetically generate anomalous samples can learn meaningful representations of the normal class through contrastive learning. However, one-class AD algorithms can be susceptible to Representational Collapse as their objective to minimise the representation space between training samples can be easily satisfied by learning to map all inputs to a constant output. One-Class method, DeepSVDD prevents Representational Collapse by removing bias terms and freezing the value of its hypersphere centre during training. PANDA prevents collapse by using early stopping and another technique known as ‘elastic regularisation’. Relatedly in the field of visual representation learning, a recent technique known as ‘Stop Gradient’ has been empirically proven to avoid Representational Collapse by preventing the backflow of gradients through one of the branches of a Siamese network [6].

### 3. FewSOME

This section outlines FewSOME, visualised in Figure 1. Given a dataset  $X \subseteq \mathbb{R}^d$ , a sample of size  $N$  is randomly sampled from  $X$  so that the training data is  $R_N = \{r_1, \dots, r_N\} \subseteq X$ . This is referred to as the ‘Reference Set’. This is different to the standard training set as the size of the Reference Set is typically a small fraction of the size of the training set. A neural network,  $f$ , is then trained to transform the input space  $R_N \subseteq \mathbb{R}^d$  to the representation space  $f(R_N) \subseteq \mathbb{R}^l$ . The representation space of  $R_N$  is de-

noted as  $f(R_N) = \{f(r_1), \dots, f(r_N)\}$ , where  $f(r_i)$  is the representation of  $r_i$  in the form of a 1D feature embedding with dimensions  $1 \times l$ . The objective of the network,  $f$ , is to learn weights  $W$  that minimise the Euclidean Distance (ED) between feature embeddings of the Reference Set. This is achieved by employing a distance-based loss function. The distance component of the loss function,  $D(r)$  consists of  $L_{dist}$  and Stop Loss, denoted as  $L_{stop}$  (equation 1). In each epoch of the training process,  $D(r_i)$  is computed for each  $r_i \in R_N$ .

$$D(r_i) = \underbrace{\sum_{k=1}^K (||f(r_i) - f(r_k)||)}_{L_{dist}} + \underbrace{\alpha ||f(r_i) - stop(f^*(r_a))||}_{L_{stop}} \quad (1)$$

#### 3.1. $L_{dist}$

The  $L_{dist}$  component represents the Siamese architecture of FewSOME. FewSOME differs from the typical Siamese Network as it trains on only one class and the number of branches in the network is a hyper-parameter of the model. The hyper-parameter  $K$  is the number of data samples in  $R_N$  that are input into  $f$  in tandem with  $r_i$ . The term  $||f(\cdot) - f(\cdot)||$  is the ED between two feature embeddings divided by  $\sqrt{l}$  where  $l$  is the dimension of the 1D feature embeddings. As the final layer of  $f$  is a Sigmoid layer, it can be proven that dividing the ED by  $\sqrt{l}$  results in the term  $||f(\cdot) - f(\cdot)||$  being bounded between zero and one.

For each  $r_i$  during training,  $r_k$  is selected randomly. A variation of FewSOME, named Smart FewSOME (S-FewSOME) selects  $r_k$  such that  $r_k = \operatorname{argmax}_{r_k} ||f(r_i) - f(r_k)||$  where  $r_i$  and  $r_k \in R_N$ . This is commonly done in distance based loss functions such as triplet loss [30]. This can speed up the convergence of the model, however, it increases the complexity of the model.

#### 3.2. $L_{stop}$

The  $L_{stop}$  component of  $D(r_i)$  computes the distance between  $f(r_i)$  and  $f(r_a)$  where  $r_a \in R_N$ . The data sample,  $r_a$  is used to compute the ‘anchor’,  $f(r_a)$ . It is arbitrarily selected from  $R_N$  prior to training. The values of  $f(r_a)$  are computed as the first initial pass through the model. Stop Gradient [6], denoted as  $stop$  in equation 1 is then applied to  $f(r_a)$ . This prevents the back-flow of gradients i.e. the weights are updated so that  $f(r_i)$  is moved closer to  $f(r_a)$ , rather than the weights updating to move both  $f(r_i)$  and  $f(r_a)$  closer to each other. This results in the value of  $f(r_a)$  being frozen for the duration of model training. Frozen values are denoted by  $*$  in  $f^*(r_a)$ . We found that the selection of  $r_a$  had an insignificant impact on model performance and model convergence (see also section 6.2). Hyper-parameter  $\alpha$ ,  $0 \leq \alpha \leq 1$ , controls the trade-off be-

tween  $L_{dist}$  and  $L_{stop}$ . The anchor in FewSOME behaves differently to DeepSVDD’s frozen centre. For example, if  $\alpha < 1$ , minimising the distance from  $f(r_i)$  to the anchor  $f(r_a)$  is not the primary objective when minimising  $D(r_i)$  and therefore, the anchor does not necessarily become the centre. Our proposed loss is also less complex as the anchor is computed based on a randomly selected training sample and it does not require the representations for all training samples to be calculated in order to assign its value, as is required when assigning the value of DeepSVDD’s centre,  $c$ .

Equation 2 shows the objective function. The first term is equal to Contrastive loss<sup>1</sup> when  $y = 0$ . During training,  $y$  will always equal zero in FewSOME as all samples in  $R_N$  are of the same class. The last term shows the weight decay regularization where  $0 \leq \lambda \leq 1$ .

$$\min_{D(r_i), W} \frac{1}{N} \sum_{i=1}^N \frac{1}{2} D(r_i)^2 + \frac{\lambda}{2} \|W\|^2 \quad (2)$$

During testing, samples from a test set  $T \subseteq \mathbb{R}^d$  are input into  $f$  to transform them to the representation space. The transformed test samples are then ranked by an anomaly scoring function,  $s$ . Existing anomaly scoring functions score a test sample based on its distance to the centre of the normal class representation space. We found that assigning an anomaly score to test point,  $t_i \in T$ , based on the distance between  $f(t_i)$  and the nearest representation in the Reference Set,  $f(r) \in f(R_N)$ , and the distance from  $f(t_i)$  to the anchor,  $f^*(r_a)$  gave optimal results (equation 3).

$$s(t_i) = \operatorname{argmin}_{f(r)} \{ \|f(t_i) - f(r)\| + \alpha \|f(t_i) - f^*(r_a)\| \} \quad (3)$$

### 3.3. Properties to Prevent Representational Collapse

We define Representational Collapse as the case in which all feature embeddings of the Reference Set are equal to the mean feature embedding of the Reference Set i.e.  $\forall r \in R_N : \frac{1}{N} \sum_{i=1}^N f(r_i) = f(r)$ . In simpler terms, collapse results in the model mapping all inputs to a constant output. We outline two scenarios in which collapse could occur and the measures that were taken to prevent it;

1. **In the absence of  $L_{stop}$ :** if the weights of any layer in  $f$  converge to all zeros, this would result in  $\forall r \in R_N : \frac{1}{N} \sum_{i=1}^N f(r_i) = f(r)$ . Although the model has not learned the features of normality, it has in theory

<sup>1</sup>Contrastive Loss =  $\frac{1}{N} \sum_{i=1}^N \frac{1}{2} (1 - y_i) (\|x_i - x_j\|)^2 + \frac{1}{2} (y_i) \operatorname{MAX}(0, \operatorname{margin} - (\|x_i - x_j\|)^2)$ , where  $N$  is the number of samples,  $y$  is the label ( $y = 0$  when both samples are the same,  $y = 1$  when two samples are different),  $\|x_i - x_j\|$  is the ED between two data samples  $x_i$  and  $x_j$  and the margin is an arbitrary small number.

achieved its objective of minimising the representation space with  $L_{dist} = 0$ . However, we prevent this scenario by including  $L_{stop}$  which minimises the distance between  $f(R_N)$  and non-zero value  $f^*(r_a)$ . This ensures that the model has not achieved its objective by outputting constant representations.

2. **In the presence of  $L_{stop}$ :** if the weights of any layer in  $f$  converge to all zeros and bias terms exist,  $f$  can achieve its objective by updating the bias terms so that each feature embedding is equal to the feature embedding of the anchor i.e.  $\frac{1}{N} \sum_{i=1}^N f(r_i) = f^*(r_a)$ . Similar to DeepSVDD, we prevent this type of representational collapse by removing bias terms.

## 4. Experiments

FewSOME was evaluated on four benchmark datasets, MNIST [19], Fashion-MNIST (F-MNIST) [39], CIFAR-10 [18] and MVTEC AD [5]. The former three datasets consist of ten classes. Following standard AD protocol, each class was set to normal and all other classes were set to anomalies. MVTEC AD has 15 classes consisting of high-resolution images of different industry objects and textures. Within each class, there is a training set of normal images and a test set of normal and anomalous samples. Ten models were trained for each normal class. Each of the ten models were trained on different randomly sampled Reference sets from the provided training data.

### 4.1. FewSOME Implementation Details

The model backbone for MNIST, CIFAR-10 and F-MNIST is a simple architecture we name ‘VGG-3’ based on the first three blocks of VGG-16. Following the training protocol of previous SOTA DROCC, 20% of the official training data was used as a validation set to tune the model parameters. The hyper-parameters of FewSOME include the Reference Set size,  $N \in \{30, 60\}$ ,  $K \in \{1, 2, 3\}$ , feature embedding dimension,  $l \in \{1024, 2048\}$ , batch size  $\in \{1, 8, 16, 30\}$ ,  $\alpha$  in the range  $[0.01, 1]$  and learning rate in the range  $[1e - 2, 1e - 8]$ . The early stopping criteria for FewSOME is to stop training once the rate at which the training loss is decreasing is less than 0.5% for a patience of two. The optimiser was Adam [16]. The image dimensions in MVTEC AD range from  $700 \times 700$  to  $1024 \times 1024$ . These images were rescaled to  $128 \times 128$ . A larger architecture, ResNet-18 [13] was then employed as the model backbone and the feature embedding dimension,  $l$  was set to 1000. For MVTEC, the average best result (averaged over ten seeds) on the test set was reported as was done for competing methods for fair comparison. Weights were initialised using Kaiming initialisation or with weights pretrained on ImageNet.

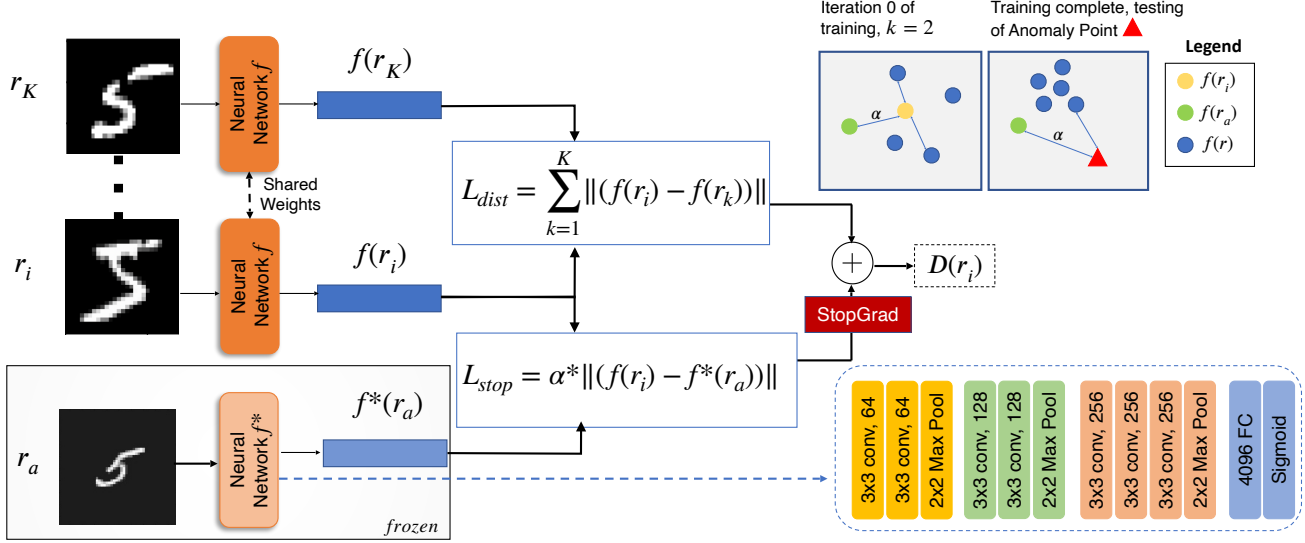


Figure 1. Training process for the MNIST dataset when the normal class is equal to five. In each iteration,  $r_i \in R_N$  and  $\{r_1, \dots, r_K\} \in R_N$  are input into the network  $f$  to obtain the feature embeddings,  $f(r_i)$  and  $\{f(r_1), \dots, f(r_K)\}$ . The values of  $f^*(r_a)$  are obtained in the first initial pass through the model and they are unchanged throughout training. The  $L_{dist}$  is calculated in parallel with  $L_{stop}$  to obtain  $D(r_i)$ . The model is trained with the objective of minimising  $D(r_i)$ . The VGG-3 architecture is shown in the figure where ‘conv’ refers to a convolution, ‘Max Pool’ refers to a Max Pooling operation and ‘FC’ refers to a Fully Connected layer. The figure also depicts a hypothetical example of the feature embeddings projected into 2D space. The blue dots are the feature embeddings of the Reference Set. The green dot is  $r_a \in R_N$ , the yellow dot is  $r_i \in R_N$  and the red triangle is an anomaly during testing. Before training all feature embeddings are dispersed in the 2D space. After training, the feature embeddings of the Reference Set are closer together, while the anomalous sample is distant from the Reference Set.

## 4.2. Competing Methods Implementation Details

We compare FewSOME against typical baseline methods. We primarily compare FewSOME against DeepSVDD and DROCC as they are proven SOTA also belonging to the ‘Deep One-Class’ class of AD techniques. As the original papers do not report on all benchmark datasets, we implement both models for F-MNIST and MVTEC and we additionally implement DROCC for MNIST. We tuned DROCC’s hyper-parameters using grid-search on a validation set of 20% of the provided training data for MNIST, CIFAR-10 and F-MNIST. For MVTEC, the average best result (averaged over ten seeds) on the test set was reported. For MVTEC, DROCC’s model backbone was scaled up from a LeNet architecture to a ResNet-18. As DeepSVDD requires pretraining of an autoencoder, it was too computationally expensive to increase the size of its architecture to a ResNet-18 for training on MVTEC. Therefore, the images were further downsized to  $64 \times 64$  and we increased the number of model parameters by two million by adding convolutional layers and kernels. Although newer techniques such as PANDA and IGD also belong to the ‘Deep One-Class’ category, it would not be a fair comparison to compare them to FewSOME given FewSOME’s simple training pipeline and smaller architecture. The recent technique,

PANDA employs ResNet-152 and WideResNet-50 architectures. IGD additionally has a complex training pipeline that involves reconstructing the images and adversarially generating training samples.

## 5. Results

### 5.1. Comparative Analysis

Table 1 shows the Area Under the Curve (AUC) averaged over all classes (ten seeds per class) compared with other SOTA on the provided test sets. Figures with an asterisk are reported based on our implementation. FewSOME and the Smart FewSOME (S-FewSOME) variation show substantial performance improvements over baseline and generative models with training data sizes of  $N = 30$  for MNIST, CIFAR-10 and F-MNIST and  $N = 60$  for MVTEC, a fraction of the training data that competing methods use (competing methods train on 5,000 images for CIFAR-10 and 6,000 images for MNIST and F-MNIST). It can be noted that S-FewSOME results in only marginal increases in performance in some cases compared to standard FewSOME, indicating that the lower complexity standard FewSOME model is sufficient. As it is known that 90% of the test data are anomalies, test samples with anomaly scores in the 10th percentile and above are classified as anomalies for all

Table 1. AUCs in % averaged over each normal class (ten seeds per class) on MNIST, CIFAR-10, F-MNIST and MVTec AD. \* Results based on our implementation. † Few-Shot learning approaches (trained on few shots of normal data).

Model	MNIST	CIFAR-10	F-MNIST	MVTec
<b>Baselines</b>				
OC-SVM [29]	91.3	64.8	×	×
KDE [15]	86.9	64.9	×	×
<b>Generative</b>				
DCAE [32]	89.7	59.4	×	×
MinLGAN [38]	×	68.7	×	×
LSA [1]	×	64.1	×	×
OCGAN [23]	97.5	65.7	×	×
AnoGAN [28]	91.3	61.8	-	50.3
MEMAE [11]	97.5	60.9	×	×
CAVGA [37]	98.6	73.7	88.5	×
HTDG† [33]	87.2	70.2	91.2	78.0
<b>One-Class</b>				
DeepSVDD [26]	92.4	64.8	90.7*	72.2*
DROCC [12]	87.8*	<b>76.9</b>	90.5*	74.5*
FewSOME†	98.0	76.6	<b>93.1</b>	<b>84.7</b>
S-FewSOME†	<b>98.1</b>	<b>76.9</b>	<b>93.1</b>	82.3

models. The F1 score and balanced accuracy were then calculated, averaged across all classes and reported in Table 2. Although FewSOME and DROCC perform similarly for CIFAR-10 in terms of AUC, Table 2 demonstrates that FewSOME outperforms it in terms of F1 score and balanced accuracy. This highlights the necessity for metrics in addition to AUC to be reported with all AD algorithms.

## 5.2. Training on Few Normal Samples

Five models were trained for each normal class in MNIST and CIFAR-10 with varying number of shots (i.e. training dataset sizes) ranging from two to 50. We repeat the experiments with competing methods DeepSVDD and DROCC. We also report the results of the previous SOTA in FSAD on normal data samples, HTDG. HTDG, as outlined in section 2, is a generative model that uses self-supervision to detect anomalies. In this section, we introduce an additional competing method, IGD [7]. IGD, as outlined previously, is a deep one-class method that uses ImageNet weights and adversarially generates samples during training. They then classify anomalies by measuring their distance to a Gaussian centre of the normal class representation space. We include this as a competing method as they specifically demonstrate the performance of their method on small datasets.

For this experiment, FewSOME and DeepSVDD were trained until the rate at which the training loss was decreasing was less than 0.5% for a patience of two. As DROCC does not benefit from pretrained weights, we used the same

Table 2. F1 score and Balanced Accuracy in % averaged over each normal class (ten seeds per class) on MNIST, CIFAR-10, F-MNIST and MVTec AD. All results are based on our implementation. The model ‘S-FewSOME’ is the Smart FewSOME variation.

	MNIST	CIFAR-10	FMNIST	MVTec
<b>F1</b>				
DeepSVDD	96.6	91.0	95.4	83.1
DROCC	95.0	92.1	95.3	83.7
FewSOME	<b>98.2</b>	92.1	<b>96.0</b>	<b>85.7</b>
S-FewSOME	<b>98.2</b>	<b>92.9</b>	<b>96.0</b>	84.9
<b>Balanced Accuracy</b>				
DeepSVDD	82.7	55.1	76.8	57.1
DROCC	73.7	58.0	76.2	58.0
FewSOME	90.9	60.6	<b>80.2</b>	<b>63.2</b>
S-FewSOME	<b>91.0</b>	<b>64.5</b>	<b>80.2</b>	61.6

stopping criteria but with increased patience from two to ten. Due to IGD’s architecture of a generator and critic module that behaves similar to a GAN discriminator, monitoring the training loss was an ineffective method for early stopping. Therefore, we evaluate it on the test set every 20 epochs and report the best result. Table 3 displays the AUC averaged over all normal classes for the varying training data sizes. There are three primary findings from Table 3. Firstly, FewSOME achieves almost optimal performance on as little as  $N = 2$ . Secondly, FewSOME outperforms the current SOTA in FSAD on normal examples, HTDG. Thirdly, FewSOME outperforms IGD, the SOTA for AD on small datasets. This is a notable result as IGD has a complex training pipeline that consists of a classifier, generator and a critic module that performs similar to a GAN discriminator, whilst FewSOME is a single network with a small number of model parameters ranging from four million to 11 million depending on whether VGG-3 or ResNet-18 is employed.

Figure 2 visualises how FewSOME reaches its peak performance when  $N = 30$ , whilst both DROCC and DeepSVDD require all available training data to achieve their peak performance. The figure also shows FewSOME requires less training time than competing methods.

## 5.3. Contaminated Datasets

The following experiment was conducted by contaminating the training data with anomalies at rates of 1%, 5%, 10% and 20%. The results were averaged over normal classes. The results reported in Table 4 show that FewSOME is robust to contamination in the data, achieving 96% of its peak performance on both CIFAR-10 and MNIST despite 20% of the Reference Set being contaminated with anomalies. At all levels of contamination, FewSOME is outperforming DeepSVDD and DROCC.




Figure 2. Top row: Percentage of peak performance achieved for each model at different training dataset sizes (averaged over all classes, 5 seeds per class) for MNIST and CIFAR-10. Bottom row: The average training time for each model at different training dataset sizes (averaged over all classes, 5 seeds per class) for MNIST and CIFAR-10. Stars represent the training data size where the model achieves its peak AUC.

Table 3. AUC in % averaged over each normal class on MNIST and CIFAR-10 for shots (i.e. training set sizes) ranging from two to 50. All results are based on our implementation with the exception of HTDG. DeepSVDD has been abbreviated to DSVDD and FewSOME has been abbreviated to FSOME.

Model	2	5	10	20	30	40	50
<b>MNIST</b>							
HTDG	×	85.9	87.2	×	×	×	×
DSVDD	75.9	78.8	80.0	80.9	81.1	81.4	81.6
DROCC	64.3	70.3	66.1	74.2	72.1	69.7	70.9
IGD	80.1	83.4	88.5	90.2	92.8	93.9	94.7
FSOME	<b>90.1</b>	<b>95.5</b>	<b>97.0</b>	<b>97.8</b>	<b>98.1</b>	<b>98.2</b>	<b>98.2</b>
<b>CIFAR-10</b>							
HTDG	×	67.5	70.2	×	×	×	×
DSVDD	57.4	58.6	59.7	60.0	59.4	59.2	59.8
DROCC	54.2	55.3	55.6	55.5	56.2	55.1	55.5
IGD	54.2	58.2	65.4	73.3	74.6	75.8	74.6
FSOME	<b>64.4</b>	<b>69.6</b>	<b>72.5</b>	<b>75.1</b>	<b>76.6</b>	<b>76.2</b>	<b>75.8</b>

## 6. Further Analysis

### 6.1. Ablation Study on $L_{stop}$

A barplot in Figure 3 shows the percentage performance increase in average AUC over ten seeds for each normal

Table 4. AUC in % averaged over each class (five seeds per class) at varying levels of contamination in the training data. Results based on our implementation.

Model	1%	5%	10%	20%
<b>MNIST</b>				
DeepSVDD	92.9	90.4	87.4	82.9
DROCC	82.8	82.7	77.9	76.9
FewSOME	<b>97.2</b>	<b>96.7</b>	<b>95.8</b>	<b>93.7</b>
<b>CIFAR-10</b>				
DeepSVDD	63.1	62.3	61.5	60.3
DROCC	70.1	69.6	68.9	68.4
FewSOME	<b>76.0</b>	<b>75.1</b>	<b>74.8</b>	<b>73.7</b>

class in the MVTEC dataset as a result of including  $L_{stop}$  in the loss function. The plot shows significant increases of up to 60% in one case, demonstrating the importance of including  $L_{stop}$  in the loss function. Representational collapse can be identified when the model immediately reaches an almost-zero training loss [6]. The training plots in Figure 3 show that the train loss immediately achieves almost zero values without  $L_{stop}$ . However, the inclusion of  $L_{stop}$  regularises the model to prevent representational collapse and




Figure 3. The barplot shows the percentage performance increase in average AUC over ten seeds as a result of including  $L_{stop}$  in the loss function. The figure also depicts two examples of the training loss at each epoch of training with and without  $L_{stop}$ .

boosts the model performance.

## 6.2. Sensitivity Analysis on Reference Set, N and choice of Anchor

A sensitivity analysis was conducted to assess how the model performance is impacted depending on (a) the choice of Reference set, (b) choice of N and (c) choice of the training sample,  $r_a$  that is used to compute the anchor. For each value of N, ten models were trained based on different randomly sampled Reference sets and randomly sampled  $r_a$ . Figure 4 shows the model’s performance on MNIST with the normal class set to zero and CIFAR-10 with the normal class set to ‘airplane’. The results show that the performance increases as N is increased from two to five before plateauing at  $N = 5$  for both datasets. From this, we outline the following guidelines for choosing N.

1. Increase N until the performance begins to plateau and there is little variability in model performance between choice of Reference Set.
2. In the absence of labelled data, our experiments have shown  $N = 30$  to achieve good performance across all datasets.

As there is little variability in model performance once N is large enough (i.e. greater than  $N = 5$ ), we can also conclude that the selection of the Reference set and  $r_a$  does not impact performance. Therefore, random sampling is an effective method for selecting the Reference set and  $r_a$ .

## 7. Conclusion

In this paper, we presented FewSOME, a new SOTA in the field of FSAD on normal data samples. By reporting on AUC, F1 scores, balanced accuracy and training times, we have demonstrated FewSOME’s ability to perform at SOTA level having trained on a fraction of the data that existing approaches require. Our extensive experiments have

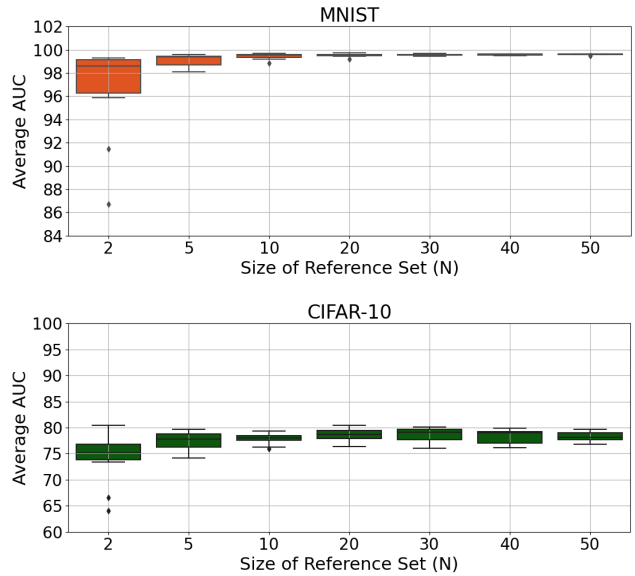


Figure 4. AUC in % averaged over ten different randomly sampled Reference Sets for different values of N and different selections of the anchor.

shown that our proposed loss, ‘Stop Loss’ prevents representational collapse and boosts performance. We have also demonstrated that FewSOME is robust to contaminated datasets and it is insensitive to choice of Reference Set and anchor.

## Acknowledgements

This work was funded by Science Foundation Ireland through the SFI Centre for Research Training in Machine Learning (Grant No. 18/CRT/6183). This work is supported by the Insight Centre for Data Analytics under Grant Number SFI/12/RC/2289 P2.

## References

- [1] Davide Abati, Angelo Porrello, Simone Calderara, and Rita Cucchiara. Latent space autoregression for novelty detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 481–490, 2019. 2, 6
- [2] Shin Ando and Ayaka Yamamoto. Anomaly detection via few-shot learning on normality. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, 2022. 3
- [3] Toshimichi Aota, Lloyd Teh Tzer Tong, and Takayuki Okatani. Zero-shot versus many-shot: Unsupervised texture anomaly detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5564–5572, 2023. 3



- [4] Liron Bergman and Yedid Hoshen. Classification-based anomaly detection for general data. *International Conference on Learning Representations*, 2020. 2
- [5] Paul Bergmann, Kilian Batzner, Michael Fauser, David Sattlegger, and Carsten Steger. The mvtec anomaly detection dataset: a comprehensive real-world dataset for unsupervised anomaly detection. *International Journal of Computer Vision*, 129(4):1038–1059, 2021. 4
- [6] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15750–15758, 2021. 3, 7
- [7] Yuanhong Chen, Yu Tian, Guansong Pang, and Gustavo Carneiro. Deep one-class classification via interpolated gaussian descriptor. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 383–392, 2022. 1, 2, 6
- [8] 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 conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 1, 2
- [9] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *International Conference on Learning Representations (ICLR)*, 2018. 2
- [10] Izhak Golan and Ran El-Yaniv. Deep anomaly detection using geometric transformations. *Advances in neural information processing systems*, 31, 2018. 2
- [11] Dong Gong, Lingqiao Liu, Vuong Le, Budhaditya Saha, Moussa Reda Mansour, Svetha Venkatesh, and Anton van den Hengel. Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1705–1714, 2019. 2, 6
- [12] Sachin Goyal, Aditi Raghunathan, Moksh Jain, Harsha Vardhan Simhadri, and Prateek Jain. Drocc: Deep robust one-class classification. In *International conference on machine learning*, pages 3711–3721. PMLR, 2020. 2, 6
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [14] Chaoqin Huang, Haoyan Guan, Aofan Jiang, Ya Zhang, Michael Spratling, and Yan-Feng Wang. Registration based few-shot anomaly detection. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIV*, pages 303–319. Springer, 2022. 3
- [15] JooSeuk Kim and Clayton D Scott. Robust kernel density estimation. *The Journal of Machine Learning Research*, 13(1):2529–2565, 2012. 2, 6
- [16] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR) (Poster)*, 2015. 4
- [17] Gregory Koch, Richard Zemel, Ruslan Salakhutdinov, et al. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2. Lille, 2015. 1
- [18] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 4
- [19] Yann LeCun, Corinna Cortes, and Chris Burges. Mnist handwritten digit database, 2010. 4
- [20] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 eighth IEEE international conference on data mining*, pages 413–422. IEEE, 2008. 2
- [21] Philipp Liznerski, Lukas Ruff, Robert A Vandermeulen, Billy Joe Franks, Klaus-Robert Müller, and Marius Kloft. Exposing outlier exposure: What can be learned from few, one, and zero outlier images. *Transactions on Machine Learning Research*, 2022. 2
- [22] Guansong Pang, Chunhua Shen, and Anton van den Hengel. Deep anomaly detection with deviation networks. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 353–362, 2019. 2
- [23] Pramuditha Perera, Ramesh Nallapati, and Bing Xiang. Ocgan: One-class novelty detection using gans with constrained latent representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2898–2906, 2019. 2, 6
- [24] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [25] Tal Reiss, Niv Cohen, Liron Bergman, and Yedid Hoshen. Panda: Adapting pretrained features for anomaly detection and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2806–2814, 2021. 1, 2
- [26] Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *ICML*, pages 4393–4402. PMLR, 2018. 2, 6
- [27] Lukas Ruff, Robert A Vandermeulen, Nico Görnitz, Alexander Binder, Emmanuel Müller, Klaus-Robert Müller, and Marius Kloft. Deep semi-supervised anomaly detection. *International Conference on Learning Representations (ICLR)*, 2019. 2
- [28] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Georg Langs, and Ursula Schmidt-Erfurth. f-anogan: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis*, 54:30–44, 2019. 2, 6
- [29] Bernhard Schölkopf, John C Platt, John Shawe-Taylor, Alex J Smola, and Robert C Williamson. Estimating the support of a high-dimensional distribution. *Neural computation*, 13(7):1443–1471, 2001. 2, 6
- [30] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. 3
- [31] Eli Schwartz, Assaf Arbel, Leonid Karlinsky, Sivan Harary, Florian Scheidegger, Sivan Doherty, and Raja Giryes. Maeday: Mae for few and zero shot anomaly-detection. 2022. 3

- [32] Philipp Seeböck, Sebastian Waldstein, Sophie Riedl, Bianca Gerendas, René Donner, Thomas Schlegl, Ursula Schmidt-Erfurth, and Georg Langs. Identifying and categorizing anomalies in retinal imaging data. 12 2016. [2](#), [6](#)
- [33] Shelly Sheynin, Sagie Benaim, and Lior Wolf. A hierarchical transformation-discriminating generative model for few shot anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8495–8504, 2021. [3](#), [6](#)
- [34] Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin. Csi: Novelty detection via contrastive learning on distributionally shifted instances. *Advances in neural information processing systems*, 33:11839–11852, 2020. [2](#)
- [35] Hironori Takimoto, Junya Seki, Sulhayanti F. Situju, and Akihiro Kanagawa. Anomaly detection using siamese network with attention mechanism for few-shot learning. *Applied Artificial Intelligence*, 36(1):2094885, 2022. [2](#)
- [36] Yu Tian, Gabriel Maicas, Leonardo Zorron Cheng Tao Pu, Rajvinder Singh, Johan W Verjans, and Gustavo Carneiro. Few-shot anomaly detection for polyp frames from colonoscopy. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VI 23*, pages 274–284. Springer, 2020. [2](#)
- [37] Shashanka Venkataramanan, Kuan-Chuan Peng, Rajat Vikram Singh, and Abhijit Mahalanobis. Attention guided anomaly localization in images. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII*, pages 485–503. Springer, 2020. [2](#), [6](#)
- [38] Chu Wang, Yan-Ming Zhang, and Cheng-Lin Liu. Anomaly detection via minimum likelihood generative adversarial networks. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 1121–1126. IEEE, 2018. [2](#), [6](#)
- [39] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. 08 2017. [4](#)