THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Few-Shot Adaptive Detection of Objects of Concern Using Generative Models with Negative Retraining

Citation for published version:

Dionelis, N, Yaghoobi Vaighan, M & Tsaftaris, SA 2021, Few-Shot Adaptive Detection of Objects of Concern Using Generative Models with Negative Retraining. in *2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI)*. IEEE, pp. 8, The 33rd IEEE International Conference on Tools with Artificial Intelligence , 1/11/21. <https://doi.org/10.1109/ICTAI52525.2021.00086>

Digital Object Identifier (DOI):

[10.1109/ICTAI52525.2021.00086](https://doi.org/10.1109/ICTAI52525.2021.00086)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Few-Shot Adaptive Detection of Objects of Concern Using Generative Models with Negative Retraining

Nikolaos Dionelis

Electronics and Electrical Engineering

The University of Edinburgh, UK

Contact email: nikolaos.dionelis@ed.ac.uk

Mehrdad Yaghoobi

Electronics and Electrical Engineering

The University of Edinburgh, UK

myvaigha@ed.ac.uk

Sotirios A. Tsafaris

*Electronics and Electrical Engineering**Digital Communications Research Institute*

The University of Edinburgh, UK

s.tsafaris@ed.ac.uk

Abstract—Detecting objects which we are interested in, Objects of Concern (OoC), is nowadays attracting attention. In aviation and transport, it is important to robustly detect OoC in security images. OoC are rare, differ from typical samples, and may be unknown during training. Most OoC detection methods need to be trained on large datasets to achieve good performance and have limited real-world generalization ability. A large variety of samples is needed, and it is expensive to collect and label large datasets due to the rarity of OoC. To address such limitations, we propose the negative REtraining with Few-shots Generative Adversarial Network (REFGAN) for detecting OoC. REFGAN aims at automatically identifying OoC by learning from Objects of No Concern (OoNC) and OoC. Our methodology comprises learning a prior using OoNC, and few-shot adaptation using the Few-Shot OoC (FSOoC). We propose a methodology to robustly perform few-shot adaptive detection of OoC using GANs and learned distribution boundaries. The evaluation of REFGAN on the Baggage SIXray dataset shows that when FSOoC are used, our model outperforms the prior improving OoC detection, and outperforms recent benchmarks by approximately 6.3% in mean values. REFGAN using few-shots of 80 samples shows a robust comparable performance compared to REFGAN using all the samples for retraining and model adaptation. REFGAN can detect unknown OoC and its evaluation on SIXray and CIFAR-10 shows robustness against the number of few-shot samples of OoC. REFGAN on CIFAR-10 outperforms benchmarks by approximately 16% using few-shots of 80 and of 10 samples.

Index Terms—Generative Adversarial Network, Anomaly Detection, Detection of Objects of Concern, Baggage X-Ray

I. INTRODUCTION

In security in aviation and transport, accurately detecting Objects of Concern (OoC), for which we only have few samples, is important. Robust OoC detection using few-shots is attracting attention. OoC are rare and may not be known before inference. It is expensive to label large datasets. This is a challenge for existing methods which have limited generalization ability and need to be trained on big datasets to yield good performance. To address such limitations, we propose the combination of Out-of-Distribution (OoD) detection with few-shot OoC detection [1]–[4]. OoD detection is the detection of samples that are different from typical data, outside the normal data distribution. OoC detection refers to the detection of unknown items. Few-shot learning is learning new classes for which we have few-shots without training from scratch [1]. This is done by learning a prior which is required by necessity since learning directly from few samples is not useful due to noisy parameter estimates.

We propose the negative REtraining with Few-shots Generative Adversarial Network (REFGAN) for detecting OoC and adaptively learning novel OoC classes. REFGAN detects OoC that do not conform to normality, learning from Objects of No Concern (OoNC) and OoC. REFGAN addresses the rarity of OoC and comprises (a) learning a prior using OoNC, and (b) model adaptation using the new OoC and the Few-Shot OoC (FSOoC). REFGAN generates OoD samples close to the boundary of the data distribution of OoNC in a self-supervised data augmentation manner, using only normal class samples. Our algorithm is adaptive as FSOoC are incrementally learned. REFGAN, which is a methodology that performs few-shot adaptive detection of OoC using GANs, adapts our model to the introduction of few-shots of new classes of OoC [1].

Our contribution is the development of REFGAN for few-shot adaptive OoC detection combining self-supervised OoD detection with iterative few-shot adaptation to improve the OoC detection performance and detect unknown OoC. Our methodology performs retraining by including negative samples from the generated distribution boundary. REFGAN performs negative data augmentation, iterative adaptation by including the generated boundary of OoNC, OoC and FSOoC, and OoD detection with active negative sampling [3], [4], [10], [11].

REFGAN combines OoD detection with few-shot learning and iterative retraining and adaptation to improve robust OoC detection. REFGAN strengthens our model for OoD detection by retraining including the learned distribution boundary and the OoC few-shots. REFGAN is architecture-agnostic and can be integrated with f-divergence-based GANs [9], [8]. It avoids assigning non-zero probability to outliers which leads to a high number of false negatives. The evaluation of REFGAN both on Baggage SIXray and on CIFAR-10 shows that when FSOoC adaptation is used for Outlier Exposure, then the OoC detection performance improves and REFGAN outperforms recent benchmarks. The robustness sensitivity analysis of the performance of REFGAN to the variation of the number of few-shot samples on SIXray and CIFAR-10, shows that REFGAN is on average robust to using few-shot samples compared to using all the available samples, and outperforms the benchmarks.

II. RELATED WORK

OoC detection can be cast as anomaly detection, which is challenging in high-dimensional spaces. In spite of the progress

Fig. 1: REFGAN where the blue points are OoNC, the red points are FSOoC, and the green points are $B(z)$ samples.

Fig. 2: Training of the proposed REFGAN using the FSOoC samples, together with active negative sampling and training.

ushered in by the use of generative models, robustness remains a challenge that is unlikely to be overcome solely by the existing methods. A challenge is the rarity of OoC. To improve OoD detection in the few-shot setting, novel methodologies should be developed [2]. Generative models learn the bulk of the data distribution, minimizing a distribution divergence.

OoD detection can be done by performing sample generation on the distribution boundary and then inference. Estimating the distribution along with its boundary is hard due to sampling complexity and sampling from low probability density regions. In [3], we have proposed the Boundary of Distribution Support Generator (BDSG) for OoD detection. In [4], we have devised the Tail of Distribution GAN (TailGAN) for sample generation on the boundary. The contributions of BDSG and TailGAN are: (a) Training generative models for OoD detection; (b) Develop algorithms for boundary generation obviating rarity, sampling complexity, mode collapse; and (c) Generation of strong and adversarial anomalies on the boundary of the distribution.

Few-shot learning is used in security when we are interested in rare occurrences with few samples. Few-shot learning uses prior knowledge to augment the learned classes [12]–[15]. By updating the prior, few-shot learning generalizes to unseen classes containing only few-shots. Deep learning models, such as the ones presented in [5]–[7], perform classification and are trained and evaluated on baggage X-ray data, but do not generalize to unseen classes, do not perform OoD detection, and do not perform few-shot learning. Discriminative models do not include OoD as a concept; they lack OoD detection ability. Also, the performance on SIXray is lacking. SIXray is

a baggage X-ray dataset with OoNC and the OoC classes of gun, knife, wrench, plier, scissors, and hammer [17], [18].

The GANomaly and Skip-GANomaly auto-encoders [22], [23] perform OoD detection and are tested on the University Baggage Anomaly (UBA) and Full Firearm Operational Benign (FFOB) datasets. The Meta-Transfer Learning model [24] decomposes X-ray images into energy tensors and then performs few-shot classification [25], [26]. The Transfer Learning CNN model [27] evaluates whether inter-scanner generalisation exists over a multiple class detection problem. It investigates the transferability of several architectures between varying X-ray scanners with differing imaging geometries, image resolutions, and material colour profiles. Using a Faster R-CNN architecture with ResNet classification, an AUPRC of 0.86 is achieved on Baggage SIXray data for two class objects (guns, knives).

The Outlier Exposure model trains detectors against outliers to improve the OoD detection performance and detect unseen anomalies while the Hierarchical Transformation Discriminating (HTD) generative model considers the setting of few-shot OoD detection, captures the multi-scale patch distributions of images, and uses transformations to optimize discriminators to improve few-shot OoD detection [28], [25]. The robustness sensitivity analysis for the mean performance of GEOM on CIFAR-10, to the variation of the number of few-shots, shows that when using 80 and 10 few-shots, the mean AUC is reduced by 15% and 28% compared to without few-shots [28], [1].

III. OUR PROPOSED ALGORITHM

We propose the REFGAN model for the detection of OoC which include prohibited items. REFGAN first learns the prior which is a distribution on the OoNC, as shown in Fig. 1. We train the prior, and we learn the data distribution of OoNC. To improve this learned distribution, we generate samples on the distribution boundary and perform retraining by including the boundary samples. The REFGAN model generates data on the boundary of the OoNC distribution. REFGAN performs self-supervised learning using retraining by including the boundary of OoNC, to improve both the learning of the data distribution of OoNC and OoD detection. During inference, REFGAN computes a discriminator score to identify abnormal OoC and detect deviations from normality, as presented in Fig. 2.

Next, during the adaptation phase, new few-shot samples are introduced. To improve the learned distribution of OoC, we create data on the support boundary of the OoC distribution. Then, we perform retraining by including the boundary samples. These few-shot OoC (FSOoC), including their corresponding boundaries, are used to negatively retrain the prior, to more accurately learn the data distribution. Our detection algorithm performs improved anomaly and out-of-data-distribution detection with respect to the FSOoC. In this way, we update the prior of OoNC by performing negative retraining using the few-shot samples, and we perform OoD detection using the updated prior. REFGAN performs few-shot adaptation using the provided OoC and also combines robust anomaly and OoD detection with FSOoC detection and Outlier Exposure [28]. We have devised a mechanism for incorporating low-probability

samples and OoC few-shot samples into our training algorithm. Our anomaly score, which is discriminator-based, is used on the updated prior, after few-shot adaptation, to detect OoC.

IV. OPTIMIZATION TASKS OF REFGAN

The REFGAN algorithm performs the following optimizations. *Step 1* learns the OoNC distribution, i.e. PRIOR. *Step 1* also finds the boundary of the support of this learned distribution, automating and mechanizing negative data augmentation. Our model performs Automatic NEgative Data Augmentation (ANEDA). By performing ANEDA on the OoNC, we generate negative samples on the boundary. Next, *Step 1* more accurately learns the distribution of the OoNC using the learned support boundary of the OoNC distribution, and this is the REFGAN BOUNDARY model. REFGAN is data-independent, in contrast to state-of-the-art models that are dataset-dependent [28].

Step 2 learns the distribution of each class of OoC. *Step 2* includes optimizations for OoC, in line with *Step 1*, using generated data on their distribution boundaries. By performing ANEDA, we create OoC samples. *Step 3* refines the REFGAN BOUNDARY from *Step 1* for OoNC to avoid and move away from the few-shots of OoC from *Step 2*. REFGAN uses few-shot samples of OoC [25], together with the generated strong anomalies, to more accurately learn the OoNC distribution. In *Step 4*, the Anomaly Detector is trained to robustly detect OoC. OoD detection is performed using the distribution boundary, and distribution metrics from this boundary [3], [4], [29].

A. Learning the Prior

We first learn a prior distribution on OoNC from the Baggage SIXray and CIFAR-10 sets. We set both a distribution and a f-divergence on OoNC using a f-divergence-based GAN [9]. We estimate the distribution boundary of OoNC [3], [4], and we retrain using the learned boundary for OoD detection.

In the adaptation stage, we update the prior using OoC and FSOoC samples [1], [2]. In the REFGAN few-shot training phase, we train using the given few-shot dataset, we generate samples on its boundary, and we perform negative retraining by setting the samples from the novel OoC class, together with the corresponding learned boundary of the FSOoC, as OoD negative data. These few-shots constitute Outlier Exposure. We incorporate prior knowledge into the design of our model by updating the prior and improving the learned distribution of the OoNC data to more effectively learn and detect OoC.

For the prior learning algorithm, REFGAN has the following steps. In *Step 1*, we train a f-divergence-based GAN [9]. The inputs are the training OoNC samples, $\mathbf{x}$. The outputs are the generator-discriminator pair, (G, D) . We use $\mathbf{z} \sim p_{\mathbf{z}} = N(\mathbf{0}, \mathbf{I})$. We use $\mathbf{x} \sim p_{\mathbf{x}}$ and $G(\mathbf{z}) \sim p_g$ [3], [4]. The optimization is

$$\arg \min_G \max_D \mathbb{E}_{\mathbf{x}} \log(D(\mathbf{x})) + \mathbb{E}_{\mathbf{z}} \log(1 - D(G(\mathbf{z}))). \quad (1)$$

In *Step 1*, we train the boundary generator, $B(\mathbf{z})$. For our model B , it is the boundary of G . Using $B(\mathbf{z})$, we create strong and adversarial anomalies on the distribution boundary. Using the f-divergence from *Step 1*, m [3], [4], the optimization is

$$\arg \min_B -m(B(\mathbf{z}), G) + \mu d(B(\mathbf{z}), G) + \nu s(B(\mathbf{z}), \mathbf{z}). \quad (2)$$

The scattering measure, $s(B(\mathbf{z}), \mathbf{z})$, preserves distance proportionality in the $\mathbf{z}$ and $\mathbf{x}$ spaces, eliminating mode collapse. In (2), the l_p -norm distance, $d(B(\mathbf{z}), G)$, and $s(B(\mathbf{z}), \mathbf{z})$ are

$$d(B(\mathbf{z}), G) = \min_{j=1, \dots, Q} \|B(\mathbf{z}) - G(\mathbf{z}_j)\|_2 \quad (3)$$

$$s(B(\mathbf{z}_i), \mathbf{z}_i) = \frac{1}{N-1} \sum_{j=1, j \neq i}^N \frac{\|\mathbf{z}_i - \mathbf{z}_j\|_2}{\|B(\mathbf{z}_i) - B(\mathbf{z}_j)\|_2} \quad (4)$$

where N is the training batch size. Q is the inference size.

In *Step 1*, we negatively retrain our GAN on OoNC using the learned boundary. The inputs are the GAN samples, $G(\mathbf{z})$, and the boundary, $B(\mathbf{z})$. The output is the new generator G' which more accurately learns and captures the OoNC distribution. Using $G'(\mathbf{z}) \sim p_{g'}$ and the discriminator C , the optimization is

$$\arg \min_{G'} \max_C \mathbb{E}_{\mathbf{x}} \log(C(\mathbf{x})) + \mathbb{E}_{\mathbf{z}} \log(1 - C(G'(\mathbf{z}))) + \mathbb{E}_{\mathbf{z}} \log(1 - C(B(\mathbf{z}))) + \mathbb{E}_{\mathbf{z}} \log(C(G(\mathbf{z}))). \quad (5)$$

Regarding the weighted sum of f-divergences in (5), when the weight of the term involving the learned distribution boundary is far from zero, then we perform a degree of exploration. We use the boundary for OoD detection along with distribution metrics which include any f-divergence and are more robust than Euclidean distances as they model probability, rarity, and expected log-likelihood. Euclidean distances are also not robust for multimodal distributions [30]. For boundary-based OoD detection, the low-probability boundary is defined as high f-divergence from the distribution and zero point-set distance.

B. Few-Shot Adaptation

Our few-shot adaptation algorithm has the following tasks. In *Step 2*, we use the few-shot OoC performing the optimisation task (1), after we have performed positive data augmentation using geometrical affine transformations according to [28].

In *Step 2*, we train our boundary generator, $B'(\mathbf{z})$, for the FSOoC using (2). In *Step 2*, we negatively retrain our GAN model on the FSOoC data with the generated distribution boundary, $B'(\mathbf{z})$, using the optimization task given by (5).

In *Step 3*, we negatively retrain our GAN model on the OoNC data samples obtained from *Step 1* to move away from the estimated data distributions of the OoC from *Step 2*.

In *Step 3*, as in *Step 1*, we negatively retrain our GAN on the OoNC samples from *Step 1* to move away from the estimated distributions of the OoC from *Step 2*, to more accurately learn the distribution of OoNC. With retraining and model refinement, we correct our predictions in low-probability OoNC regions and in regions of OoC. The REFGAN algorithm adapts our model to the introduction of few-shot samples of novel OoC to improve the performance of OoC detection using Outlier Exposure. These few-shots are outliers, anomalies, and OoD [1], [25]. REFGAN is able to detect accurately and robustly these few-shots because we have pooled apart the learned distributions and their generated boundaries from these few-shots, and this is the idea behind the REFGAN methodology. In this way, using *Step 3*, REFGAN separates, pools apart, and splits the learned data distributions from the few-shots.

The inputs are $G'(\mathbf{z})$ and the few-shot OoC, $\mathbf{x}^*$. The output is the generator $G''(\mathbf{z}) \sim p_{g''}$. Using the discriminator C' , as in (5), to avoid and move away from $\mathbf{x}^*$, the optimization is

$$\begin{aligned} \arg \min_{G''} \max_{C'} & \mathbb{E}_{\mathbf{x}} \log(C'(\mathbf{x})) + \mathbb{E}_{\mathbf{z}} \log(1 - C'(G''(\mathbf{z}))) + \\ & \mathbb{E}_{\mathbf{x}} \log(1 - C'(\mathbf{x}^*)) + \mathbb{E}_{\mathbf{z}} \log(1 - C'(B(\mathbf{z}))) \\ & + \mathbb{E}_{\mathbf{z}} \log(C'(G'(\mathbf{z}))) + \mathbb{E}_{\mathbf{z}} \log(1 - C'(B'(\mathbf{z}))). \end{aligned} \quad (6)$$

The REFGAN methodology performs OoD detection with Outlier Exposure with Few-Shots. We perform iterative adaptation to outliers for OoD detection, as well as few-shot adaptation to known outlier classes. REFGAN performs refinement using few-shot outliers and the introduction of novel classes [1]. REFGAN is adaptive and takes into account OoC as they are introduced. We adapt the prior to detect OoC by moving away from a priori known OoC using two poles, one positive and one negative, in line with Contrastive Learning [15]. REFGAN uses low-shot adaptation by performing retraining to match the target domain from the source domain with few-shots.

REFGAN uses distribution metrics to solve the few-shots issue. GAN discriminators find distribution metrics, including any f-divergence. GANs operate on samples from distributions. Distribution metrics can be computed between a model of a distribution and samples from a distribution, between models of distributions, or between samples. Few-shots from a new class arise [1], [25], and the REFGAN algorithm uses these samples and their generated boundaries as OoD to negatively retrain our model. The incremental retraining term, $\mathbb{E}_{\mathbf{x}} \log(1 - C'(\mathbf{x}^*))$, in (6) is a discrete distribution using Dirac functions. Using (5) and (6), REFGAN performs active negative sampling and training and learned negative data augmentation. By sampling from low-probability regions, using the distribution boundary, we improve few-shot robustness and model performance.

Our REFGAN methodology, compared to others for OoD detection with Outlier Exposure, Few-Shot Learning, Domain Adaptation, and Prior Learning, performs iteratively adaptive OoD Detection to the introduction of new outliers and OoC. REFGAN improves anomaly and OoD detection after new known OoC are introduced with outlier exposure. Our REFGAN algorithm addresses the OoD detection plus the Few-Shot Learning problem [13]. It intersects with other methods such as [28], [15]. REFGAN proposes a solution for the Generalized Few-Shot Learning setting [13], [14], which subsumes the Few-Shot Learning problem as a special case and is more realistic [31], [32]. In the Generalized Few-Shot Learning setting, the REFGAN prior and its various boundaries are computed on many samples where GANs are known to perform well, and the few-shots are OoD that REFGAN can robustly detect.

C. Inference Discriminator

Algorithm for Step 4. The inputs are FSOoC, $\mathbf{x}^*$, and $G''(\mathbf{z})$. The output is the discriminator J , and the optimization is

$$\begin{aligned} \arg \max_J & \mathbb{E}_{\mathbf{x}} \log(1 - J(\mathbf{x})) + \mathbb{E}_{\mathbf{z}} \log(1 - J(G''(\mathbf{z}))) + \\ & \mathbb{E}_{\mathbf{x}} \log(J(\mathbf{x}^*)) + \mathbb{E}_{\mathbf{z}} \log(J(B(\mathbf{z}))) + \mathbb{E}_{\mathbf{z}} \log(J(B'(\mathbf{z}))). \end{aligned} \quad (7)$$

For any queried test sample, $\tilde{\mathbf{x}}$, the Inference Score (IS) is

$$IS(\tilde{\mathbf{x}}) = J(\tilde{\mathbf{x}}) + \lambda fD(G'', \tilde{\mathbf{x}}) \quad (8)$$

where $fD(G'', \tilde{\mathbf{x}})$ is the f-divergence between G'' and $\tilde{\mathbf{x}}$.

The inference mechanism for OoC detection is given by (8). The classification decision is given by: $\tilde{\mathbf{x}}$ is OoNC if $IS(\tilde{\mathbf{x}}) < \tau$, where τ is a predefined threshold [3,4]. $\tilde{\mathbf{x}}$ is OoC otherwise.

D. Multi-Class Classification

Our REFGAN methodology can perform multi-class classification and recognize the class of OoC using Nearest Neighbors with f-divergences [15]. The output of our algorithm is class of OoC, class of OoNC, or OoD; this is a multi-class comparison with few-shots with $M+2$ classes for M OoC classes and one OoNC class. We develop a model to discriminate $M+1$ classes plus OoD, in contrast with discriminative models which do not include OoD as a concept and do not perform OoD detection. ResNet50 in [17] cannot detect OoD to OoNC and the 6 OoC classes. OoC are not confined in finite sets and REFGAN finds OoNC, known OoC, or unseen OoC. For class label assignment, using the class-independent threshold $\tilde{\tau}$, REFGAN finds

$$\begin{aligned} c = \arg \min & (fD(G''_1, \tilde{\mathbf{x}}), fD(G''_2, \tilde{\mathbf{x}}), \dots \\ & fD(G''_{M+1}, \tilde{\mathbf{x}})), \text{ and OoD when } fD(G''_c, \tilde{\mathbf{x}}) > \tilde{\tau}. \end{aligned} \quad (9)$$

V. EVALUATION OF REFGAN

Evaluation Metrics. We evaluate REFGAN using several metrics, including the Area Under the Receiver Operating Characteristics curve (AUC), the Area Under the Precision Recall Curve (AUPRC), and the Accuracy (ACC) [1], [25]. The REFGAN methodology performs multi-class comparison with few-shot samples. Our generator is trained to keep away from samples from each class of OoC. We improve the distribution of OoNC by moving away from each class of OoC, and we report the improved REFGAN performance in AUC-type scores.

Models. We train a f-divergence GAN, the KL-Wasserstein GAN (KLWGAN) [9] for (1) in *Steps* 1 and 2. We train Convolutional Neural Networks (CNN) with batch normalization for our models [3]. We train our prior using (1)-(5). This prior includes learned negative data augmentation and the generation of distribution boundary samples. These generated samples are on the boundary of the data distribution of the OoNC.

Active negative sampling and training, together with self-supervised learning, are used to sample strong anomalies close to the boundary of the OoNC distribution. The model in [33] generates OoD examples in the Mahalanobis distance; in contrast, REFGAN generates OoD samples in the f-divergence distribution metric. We perform positive data augmentation on the OoC, and using (1)-(5) the model includes learned negative data augmentation and the generation of the distribution boundary for the OoC. In the few-shot setting [25], in the few-shot adaptation stage, we update this prior using the optimization in (6) to more accurately learn the OoNC distribution. During inference, REFGAN detects prohibited OoC samples in a robust manner, by calculating our proposed inference metrics.

Data. Focusing on real-world security, we use OoNC and OoC from SIXray [17] that comprises 3131 guns, 1943 knives,

TABLE I: Ablation study of REFGAN for OoC detection on SIXray data, in AUC|AUPRC|ACC, compared to PRIOR.

	REFGAN	REFGAN BOUNDARY	PRIOR
MEAN	0.81 0.84 0.74	0.77 0.78 0.66	0.66 0.69 0.61
GUN	0.94 0.93 0.87	0.93 0.92 0.86	0.77 0.74 0.69
KNIFE	0.91 0.87 0.82	0.86 0.84 0.77	0.71 0.70 0.62
WRENCH	0.67 0.76 0.67	0.72 0.71 0.61	0.63 0.66 0.56
PLIERS	0.82 0.87 0.76	0.79 0.85 0.70	0.68 0.69 0.59
SCISSORS	0.68 0.70 0.52	0.62 0.65 0.38	0.56 0.63 0.60
HAMMER	0.86 0.91 0.79	0.77 0.76 0.74	0.63 0.69 0.59

2199 wrenches, 3961 pliers, 983 scissors, and 60 hammers. We train the REFGAN prior on the OoNC class of SIXray. OoNC follow a multimodal distribution with a disjoint-components support because glasses, hangers, and clothes are OoNC.

For few-shot adaptation, we train REFGAN by updating the prior on the OoC classes of SIXray. During inference, we test the prior and REFGAN on samples from OoNC and OoC.

We also train and test REFGAN on CIFAR-10. In the few-shot adaptation stage, we train REFGAN by refining the prior on the OoC class of CIFAR-10 [28]. During inference, we test REFGAN on samples from both the OoNC class and the OoC class of CIFAR-10. We perform One-Class Classification (OCC) evaluation on CIFAR-10 [28], [3], using retraining with the learned boundary and positive data augmentation. We perform retraining using the boundary to better learn the distribution of OoC and OoNC, and then inference for OoD detection.

Benchmarks. For SIXray, we use OoD detection benchmarks, [17], [24], [27]. For CIFAR-10 data, we use [28].

A. Evaluation of REFGAN on Baggage SIXray

Prior: We train the prior on data from the OoNC class of SIXray. We evaluate the performance of this prior [9]. Table I presents the performance of PRIOR. The prior model achieves an average AUC of 0.66, AUPRC 0.69, and ACC 0.61.

Improved Prior Learning: We update the prior on baggage data from the OoNC class of SIXray. We test the performance of this prior. Table I examines the performance of REFGAN BOUNDARY corresponding to *Step 1*. REFGAN BOUNDARY achieves mean AUC 0.77, AUPRC 0.78, and ACC 0.66.

Few-Shot Adaptation: In the few-shot adaptation stage, we refine the prior using OoNC and OoC, as well as their learned boundaries. We use few-shots of outliers for retraining [1].

Ablation Study. Table I presents the performance of REFGAN in AUC, AUPRC, and ACC on SIXray. It shows that REFGAN achieves mean AUC 0.81, AUPRC 0.84, and ACC 0.74 and outperforms the corresponding PRIOR model [9] and the improved prior, i.e. the REFGAN BOUNDARY model.

Table I shows the performance of REFGAN on SIXray on average and per class. It shows the ablation study of REFGAN for OoC detection in comparison to REFGAN BOUNDARY,

TABLE II: Performance of REFGAN for OoC detection on SIXray data in AUC|AUPRC|ACC, compared to benchmarks.

	REFGAN	SIXRAY [17]	TR. LEARN- ING [27]
MEAN	0.81 0.84 0.74	— 0.79 0.65	
GUN	0.94 0.93 0.87	— 0.87 0.79	— 0.91 —
KNIFE	0.91 0.87 0.82	— 0.86 0.77	— 0.81 —
WRENCH	0.67 0.76 0.67	— 0.70 0.59	
PLIERS	0.82 0.87 0.76	— 0.88 0.72	
SCISSORS	0.68 0.70 0.52	— 0.66 0.40	
HAMMER	0.86 0.91 0.79		

TABLE III: Mean OoC detection performance of REFGAN on SIXray in AUC|AUPRC|ACC, compared to benchmarks.

	REFGAN [17]	SIXRAY [17]	META- TR. [24]	TRANSFER LE. [27]
(ALL)	0.81 0.84 0.74	— 0.79 0.65	— 0.65 0.89	
(GUN, KNIFE)	0.93 0.90 0.85	— 0.87 0.78		— 0.86 —

corresponding to *Step 1*, and to PRIOR. In this ablation study, for the OoC class of hammers, REFGAN uses only the 60 few-shot available samples [17]. The ablation study of REFGAN presented in Table I indicates that REFGAN improves upon the prior for OoC detection. When comparing with the PRIOR and REFGAN BOUNDARY models, Table I shows that we improve the performance on average. The mean performance of REFGAN in AUC, AUPRC, and ACC is higher compared to that of PRIOR and REFGAN BOUNDARY. REFGAN differs from our REFGAN BOUNDARY model by performing few-shot learning using the outlier exposure of samples of OoC. By pooling apart and separating the learned data distribution from the outlier few-shot samples in *Step 3*, REFGAN outperforms both the PRIOR and the REFGAN BOUNDARY models.

REFGAN improves the mean performance by 22.7%, 21.7% and 21.3%, in the AUC, AUPRC and ACC metrics respectively, compared to the PRIOR model [9]. Compared to REFGAN BOUNDARY, the REFGAN model achieves an improvement of approximately 5.2% 7.7%, and 12.1%, respectively.

Performance Compared to Benchmarks. Table II presents the OoC detection performance of REFGAN which performs few-shot anomaly and OoD detection. The algorithm in [17] at inference time assigns a class label to the samples, in addition to OoC detection. Table II shows the performance of REFGAN compared to benchmarks [17], [27]. Table III presents the mean performance of REFGAN compared to [17], [24], [27].

Robustness to Few-Shots. Table IV shows the robustness sensitivity analysis of REFGAN in AUC, AUPRC, and ACC to the number of few-shots. Fig. 3 shows the performance of

TABLE IV: Robustness sensitivity analysis of REFGAN to the number of few-shots on SIXray, in AUC|AUPRC|ACC.

	REFGAN	REFGAN FS 80	REFGAN FS 10
MEAN	0.81 0.84 0.74	0.79 0.77 0.76	0.71 0.73 0.65
GUN	0.94 0.93 0.87	0.89 0.81 0.83	0.76 0.74 0.65
KNIFE	0.91 0.87 0.82	0.80 0.85 0.74	0.75 0.73 0.66
WRENCH	0.67 0.76 0.67	0.75 0.64 0.80	0.69 0.70 0.62
PLIERS	0.82 0.87 0.76	0.88 0.78 0.82	0.77 0.74 0.69
SCISSORS	0.68 0.70 0.52	0.63 0.72 0.57	0.57 0.66 0.61
HAMMER	0.86 0.91 0.79	0.86 0.91 0.79	0.74 0.78 0.67

Fig. 3: REFGAN performance on SIXray using few-shots of 80 and 10 and all available samples, per class and in mean.

REFGAN, in AUPRC. We test REFGAN with respect to its few-shot sensitivity using datasets of variable length of OoC, using all data and few-shots of 80 and 10 per OoC class.

Fig. 3 presents the robustness sensitivity of REFGAN on SIXray in AUPRC per class compared to variable few-shots. Table IV and Fig. 3 show that REFGAN using few-shots of 80 samples is robust and is also comparable to REFGAN using all the available samples. REFGAN using 80 few-shots achieves comparable AUC and ACC to our REFGAN model using all the available image samples. Table IV shows the slight percentage decrease of 1.5% of the average performance of REFGAN in AUC, when 80 few-shot samples are used. It shows a mean decrease of 11.4% when 10 few-shots are used, compared to when all the available data samples are used in REFGAN.

Ablation Study for Few-Shots. REFGAN achieves a mean ACC of 0.76 for 80 few-shots, as shown in Table IV. According to Tables I and IV, the mean ACC improves from 0.66 using the REFGAN BOUNDARY, to 0.76 using REFGAN for 80 few-shots: 15% improvement. For 10 few-shots, REFGAN achieves a similar mean ACC compared to REFGAN BOUNDARY.

B. Evaluation of REFGAN on CIFAR-10 Data

We evaluate the proposed REFGAN model on the CIFAR-10 dataset using the OCC evaluation methodology, as in [28].

TABLE V: Ablation study of REFGAN for OoC detection on CIFAR-10 data in AUC|AUPRC|ACC, compared to PRIOR.

	REFGAN	REFGAN BOUNDARY	REFGAN PRIOR
MEAN	0.89 0.88 0.82	0.86 0.88 0.82	0.63 0.64 0.60
PLANE	0.92 0.89 0.85	0.90 0.99 0.88	0.67 0.61 0.68
CAR	0.90 0.88 0.82	0.88 0.95 0.85	0.59 0.68 0.67
BIRD	0.83 0.79 0.76	0.79 0.91 0.78	0.61 0.62 0.60
CAT	0.89 0.93 0.87	0.87 0.95 0.83	0.72 0.78 0.68
DEER	0.84 0.79 0.77	0.82 0.74 0.73	0.58 0.57 0.58
DOG	0.85 0.81 0.78	0.81 0.77 0.79	0.64 0.58 0.67
FROG	0.88 0.92 0.86	0.85 0.81 0.82	0.57 0.59 0.48
HORSE	0.90 0.87 0.81	0.83 0.74 0.73	0.55 0.52 0.53
SHIP	0.95 0.96 0.88	0.94 0.97 0.92	0.62 0.63 0.55
TRUCK	0.90 0.93 0.82	0.88 0.93 0.88	0.71 0.72 0.55

Ablation Study. In Table V, we present the ablation study of REFGAN for OoC detection on CIFAR-10 in AUC, AUPRC, and ACC, compared to PRIOR and REFGAN BOUNDARY. Table V shows that using few-shots, our REFGAN methodology improves the OoC detection performance. REFGAN improves the mean OoC detection performance by 41.3% and 37.5%, in AUC and AUPRC, respectively, compared to the PRIOR. Compared to REFGAN BOUNDARY, REFGAN achieves an improvement of about 3.5% on average, in AUC. REFGAN outperforms REFGAN BOUNDARY in AUC, per class.

Ablation Study for Few-Shots. Table V and Fig. 4 present that REFGAN shows robustness to the variation of the number of few-shot samples introduced, per class and on average.

Performance compared to Benchmarks. Table VI presents REFGAN’s OoC detection performance on CIFAR-10 [28], [25]. We use the OCC evaluation methodology, and we report the AUC, AUPRC, and ACC. We compare REFGAN to GEOM, GOAD, and DROCC [28]. REFGAN is on average comparable to GEOM and GOAD, and it also outperforms DROCC.

Robustness to Few-Shots. Table VII shows the robustness sensitivity analysis of REFGAN on CIFAR-10 to the number of few-shots in the training and testing datasets. Both Table VII and Fig. 4 show that REFGAN using few-shot samples of 80 images is robust and is comparable to REFGAN using all the available data samples for retraining and model adaptation.

Comparing the performance of REFGAN on SIXray and CIFAR-10 in Figs. 3 and 4, we observe that REFGAN for few-shots of 10 and 80 samples shows smaller deviations compared to REFGAN using all available data on CIFAR-10, compared to SIXray. According to Fig. 4, REFGAN using 10 few-shot samples, as well as 80 samples, on CIFAR-10 is more robust than REFGAN on SIXray, and shows smaller deviations.

Table VIII shows the mean OoC detection performance for few-shots of 80 and 10 in AUC, AUPRC, and ACC. REFGAN

TABLE VI: OoC detection performance of REFGAN in AUC|AUPRC on CIFAR-10, compared to benchmarks.

	REFGAN	GEOM	GOAD	DROCC
MEAN	0.89 0.88	0.86 0.53	0.88 —	0.74 —
PLANE	0.92 0.89	0.75 0.29	0.73 —	0.82 —
CAR	0.90 0.88	0.96 0.79	0.96 —	0.77 —
BIRD	0.83 0.79	0.78 0.33	0.83 —	0.46 —
CAT	0.89 0.93	0.72 0.33	0.78 —	0.67 —
DEER	0.81 0.75	0.88 0.43	0.88 —	0.74 —
DOG	0.85 0.81	0.88 0.53	0.87 —	0.74 —
FROG	0.88 0.92	0.83 0.49	0.89 —	0.74 —
HORSE	0.90 0.87	0.96 0.81	0.96 —	0.71 —
SHIP	0.95 0.96	0.93 0.68	0.94 —	0.80 —
TRUCK	0.90 0.93	0.91 0.59	0.93 —	0.73 —

on CIFAR-10 achieves good performance and outperforms benchmarks for few-shots of 80 and 10 samples. In Table VIII, we compare REFGAN with the recent few-shot benchmarks of GEOM, GOAD, DROCC, and Hierarchical Transformation Discrimination (HTD) [28]. Table VIII shows that REFGAN outperforms the benchmarks by approximately 16% using few-shots of 80 and 10 samples. Table IX shows the performance of REFGAN on both the SIXray and the CIFAR-10 datasets in average AUC-type metrics, in relation to the impact of the number of few-shot image examples present in the dataset.

VI. CONCLUSION

We have proposed REFGAN to perform negative retraining with few-shots for detecting OoC. REFGAN uses the learned boundary of OoNC in a self-supervised learning manner. We first learn the prior and in our few-shot adaptation step for OoC, we retrain by setting the samples from the novel class of OoC as OoD. REFGAN uses few-shot outlier exposure [28], [34]. It adapts our model to the introduction of few-shots of new OoC. These few-shots are outliers that REFGAN is able to identify robustly. This holds because we have moved away the learned data distributions, together with their learned boundaries, from these few-shot samples. Our REFGAN algorithm is based on retraining using the self-supervised generated boundary [3].

This is the first work that develops a methodology combining negative training using the learned distribution boundary with a GAN to robustly detect few-shots of OoC. The evaluation shows that anomaly detection with Outlier Exposure combined with few-shot OoD detection [33], [35], [36], within the REFGAN methodology, is beneficial. In the few-shot setting [25], [1], first moving away from and then detecting FSOoC is effective. The mean AUPRC improves on SIXray from 0.69 using PRIOR to 0.84 using REFGAN: improvement of 21.7%. The AUPRC improves in mean values from 0.78 using REFGAN BOUNDARY in Table I to 0.84 using the proposed REFGAN model: improvement of 7.7%. REFGAN outperforms

TABLE VII: Robustness sensitivity analysis of REFGAN to the number of few-shots on CIFAR, in AUC|AUPRC|ACC.

	REFGAN	REFGAN FS 80	REFGAN FS 10
MEAN	0.89 0.88 0.82	0.88 0.88 0.80	0.82 0.87 0.76
PLANE	0.92 0.89 0.85	0.86 0.89 0.77	0.73 0.96 0.70
CAR	0.90 0.88 0.82	0.88 0.87 0.84	0.86 0.85 0.75
BIRD	0.83 0.79 0.76	0.82 0.85 0.74	0.79 0.80 0.71
CAT	0.89 0.93 0.87	0.90 0.92 0.75	0.86 0.90 0.81
DEER	0.84 0.79 0.77	0.85 0.87 0.76	0.81 0.85 0.72
DOG	0.85 0.81 0.78	0.88 0.81 0.80	0.86 0.86 0.77
FROG	0.88 0.92 0.86	0.91 0.89 0.82	0.87 0.90 0.80
HORSE	0.90 0.87 0.81	0.89 0.84 0.84	0.88 0.81 0.80
SHIP	0.95 0.96 0.88	0.94 0.91 0.86	0.74 0.91 0.73
TRUCK	0.90 0.93 0.82	0.89 0.92 0.81	0.84 0.89 0.79

Fig. 4: REFGAN performance on CIFAR-10 using few-shots of 80 and 10 and all available samples, per class and in mean.

TABLE VIII: OoC detection performance of REFGAN on CIFAR-10 in AUC and AUPRC, for FSOoC of 80 and 10.

	REFGAN	GEOM	GOAD	DROCC	HTD [28]
MEAN	0.89 0.88	0.86 0.54	0.88 —	0.74 —	—
FS 80	0.88 0.88	0.74 —	0.56 —	0.58 —	0.76 —
FS 10	0.82 0.87	0.62 —	0.61 —	0.57 —	0.70 —

benchmark models by approximately 6.3% in mean values. The robustness sensitivity analysis of REFGAN to the variation of the number of few-shots on SIXray data shows that it achieves on average robustness. It achieves comparable results when using 80 few-shots compared to using all available samples.

On CIFAR-10, REFGAN shows a mean AUC improvement of approximately 40% compared to PRIOR. REFGAN yields an improvement of about 3.5% compared to REFGAN BOUND-

TABLE IX: REFGAN performance on SIXray and CIFAR-10 in mean AUC|AUPRC|ACC against the number of few-shots.

DATASET: OoC DATA	AUC	AUPRC	ACC
SIXRAY: AVAILABLE	0.81	0.82	0.73
SIXRAY: FS 80	0.79	0.77	0.76
SIXRAY: FS 10	0.71	0.73	0.65
CIFAR-10: AVAILABLE	0.89	0.88	0.82
CIFAR-10: FS 80	0.88	0.88	0.80
CIFAR-10: FS 10	0.82	0.87	0.76

ARY. REFGAN outperforms benchmarks by approximately 16% using few-shots of 80 and 10. It shows robustness to the number of few-shots of OoC. REFGAN using few-shots of 80 and 10 shows a robust comparable performance compared to REFGAN using all samples for retraining and adaptation.

In this work, we show that REFGAN, using FSOoC samples for OoC detection, increases the OoC detection performance, outperforms benchmarks using few-shots, and is robust to few-shots of OoC for 80 and 10 samples [26], [37], [38]. Our generated boundary-based retraining methodology significantly improves the few-shot performance per class and in mean values. For specific OoC per class in SIXray and CIFAR-10, we observe that when using REFGAN and REFGAN BOUNDARY, classes which were misdetected using the PRIOR, now have improved performance. In the future, we will evaluate REFGAN using a large number of classes using (9) to avoid softmax [39] and compare it to [15] and meta-learning methods that address the few-shot learning setting [25], [40], [41]. In addition to few-shot OoD detection, our algorithm will also output confidence and bounding boxes around anomalies/objects within a robust integrated classification-with-OoD-detection framework.

ACKNOWLEDGMENT

This work was supported by the UK EPSRC Grant Number EP/S000631/1 and the UK MOD UDRC in Signal Processing.

REFERENCES

- [1] T. Robinson, "Few-Shot Learning for Defence and Security," in Proceedings Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications II, SPIE, vol. 11413, pp. 93-106, 2020.
- [2] Defence Science and Technology Laboratory (Dstl), "Machine Learning with Limited Data," 2020. Online: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/941511/Machine_Learning_with_Limited_Data_original_version_.pdf
- [3] N. Dionelis, M. Yaghoobi, and S. A. Tsafaris, "Boundary of Distribution Support Generator (BDSG): Sample Generation on the Boundary," in Proc. International Conference on Image Processing (ICIP), 2020.
- [4] N. Dionelis, M. Yaghoobi, and S. A. Tsafaris, "Tail of Distribution GAN (TailGAN): Generative-Adversarial-Network-Based Boundary Formation," in Proc. Sensor Signal Processing for Defence (SSPD), 2020.
- [5] Taimur Hassan, et al., "Detecting Prohibited Items in X-Ray Images: A Contour Proposal Learning Approach", in Proceedings ICIP, 2020.
- [6] Domingo Mery, et al., "X-Ray Baggage Inspection With Computer Vision: A Survey", in Proceeding IEEEAccess, 2020.
- [7] S. Akcay, et al., "Towards Automatic Threat Detection: A Survey of Advances within X-ray Security Imaging", arXiv:2001.01293, 2020.
- [8] S. Nowozin, et al., "f-GAN: Training Generative Neural Samplers using Variational Divergence Minimization," in Proc. NeurIPS, 2016.
- [9] J. Song and S. Ermon, "Bridging the Gap Between f-GANs and Wasserstein GANs," in Proc. ICML, pp. 9078-9087, vol. 119, 2020.
- [10] J. Sipple, "Interpretable, Multidimensional, Multimodal Anomaly Detection with Negative Sampling," in Proc. Conference ICML, 2020.
- [11] A. Sinha, et al., "Negative Data Augmentation," in Proceedings International Conference on Learning Representations (ICLR), 2021.
- [12] Yaqing Wang, et al., "Generalizing from a Few Examples: A Survey on Few-Shot Learning", arXiv:1904.05046, 2020.
- [13] Bharath Hariharan, et al., "Low-Shot Visual Recognition by Shrinking and Hallucinating Features," in Proc. Conference ICCV, 2017.
- [14] Yu-Xiong Wang, et al., "Low-shot learning from imaginary data," in Proceedings CVPR, 2018.
- [15] Priya Goyal, et al., "Self-supervised Pretraining of Visual Features in the Wild," arXiv:2103.01988, March 2021.
- [16] Mathilde Caron, et al., "Unsupervised Learning of Visual Features by Contrasting Cluster Assignments," in Proc. NeurIPS, 2020.
- [17] Caijing Miao, et al., "SIXray : A Large-scale Security Inspection X-ray Benchmark for Prohibited Item Discovery," in Proc. CVPR, 2019.
- [18] D. Mery, V. Rizzo, et al., "GDxray: The database of X-ray images for nondestructive testing," Journal of Nondestructive Evaluation, 2015.
- [19] Yanlu Wei, et al., "Occluded Prohibited Items Detection: An X-ray Security Inspection Benchmark and De-occlusion Attention Module," in Proc. ACM International Conference on Multimedia, p. 138-146, 2020.
- [20] M. Z. Zaheer, et al., "Old is Gold: Redefining the Adversarially Learned One-Class Classifier Training Paradigm," arXiv:2004.07657, 2020.
- [21] L. Ruff, R. Vandermeulen, N. Gornitz, et al., "Deep Semi-Supervised Anomaly Detection," in Proceedings ICLR, 2020.
- [22] S. Aksay, A. Atapour-Abarghouei, and T. Breckon, "GANomaly: Semi-Supervised Anomaly Detection," arXiv:1805.06725 [cs.CV], 2018.
- [23] S. Aksay, et al., "Skip-GANomaly: Skip Connected and Adversarially Trained Encoder-Decoder Anomaly Detection," arXiv:1901.08954, 2019.
- [24] T. Hassan, M. Shafay, S. Akcay, and S. Khan, "Meta-Transfer Learning Driven Tensor-Shot Detector for the Autonomous Localization and Recognition of Concealed Baggage Threats," 2020.
- [25] T. Jeong and H. Kim, "OoD-MAML: Meta-Learning for Few-Shot Out-of-Distribution Detection and Classification," in Proc. NeurIPS, 2020.
- [26] E. Robb, et al., "Few-Shot Adaptation of Generative Adversarial Networks," arXiv preprint, arXiv:2010.11943, 2020.
- [27] Y. Gaus, et al., "Evaluating the Transferability and Adversarial Discrimination of CNNs for Threat Object Detection and Classification within X-Ray Security Imagery," arXiv:1911.08966, 2019.
- [28] S. Sheynin, "A Hierarchical Transformation-Discriminating Generative Model for Few Shot Anomaly Detection," arXiv:2104.14535, 2021.
- [29] R. Devon Hjelm, et al., "Boundary-Seeking Generative Adversarial Networks," arXiv:1702.08431, 2018.
- [30] Heiga Zen, "Statistical Parametric Speech Synthesis: From HMM to LSTM-RNN," Google Research, Slide 77 of 104, "Limitations of DNN-based acoustic modeling," 2015. Online: <https://static.googleusercontent.com/media/research.google.com/en/pubs/archive/44312.pdf>
- [31] V. Verma, et al., "Meta-Learning for Generalized Zero-Shot Learning," in Proc. AAAI Conference Artificial Intelligence, 34(4), 6062-6069, 2020.
- [32] J. Kim, et al., "Model-agnostic boundary-adversarial sampling for test-time generalization in few-shot learning," in Conference ECCV, 2020.
- [33] V. Sehwag, et al., "SSD: A unified framework for self-supervised outlier detection," in Proc. International Conference ICLR, 2021.
- [34] S. Fort, J. Ren, and B. Lakshminarayanan, "Exploring the Limits of Out-of-Distribution Detection," arXiv:2106.03004 preprint, 2021.
- [35] D. Kimura, et al., "Adversarial Discriminative Attention for Robust Anomaly Detection," in Proc. Winter Conference on Applications of Computer Vision (WACV), p. 2172-2181, 2020.
- [36] A. Frikha, et al., "Few-Shot One-Class Classification via Meta-Learning," in Proc. AAAI Conference on Artificial Intelligence, 2021.
- [37] K. Wang, et al., "Few-shot Out-of-Distribution Detection," ICML Workshop on Uncertainty and Robustness in Deep Learning, 2020.
- [38] J. Tack, et al., "CSI: Novelty Detection via Contrastive Learning on Distributionally Shifted Instances," in Proc. NeurIPS, 2020.
- [39] L. He, et al., "Softmax dissection: Towards understanding intra- and inter-class objective for embedding learning," in Proc. AAAI, 2020.
- [40] X. Chen, et al., "A Boundary Based Out-Of-Distribution Classifier for Generalized Zero-Shot Learning," in Proc. Conference ECCV, 2020.
- [41] H. Kvinge, et al., "One Representation to Rule Them All: Identifying Out-of-Support Examples in Few-Shot Learning," arXiv:2106.01423.