

# Out-of-Distribution Detection with Semantic Mismatch under Masking

Yijun Yang<sup>✉</sup>, Ruiyuan Gao<sup>✉</sup>, and Qiang Xu<sup>✉</sup>

CUHK REliable Computing Laboratory (CURE Lab.)

Department of Computer Science and Engineering

*The Chinese University of Hong Kong, Hong Kong S.A.R., China*

{yjyang, rygao, qxu}@cse.cuhk.edu.hk

**Abstract.** This paper proposes a novel out-of-distribution (OOD) detection framework named MOODCAT for image classifiers. MOODCAT masks a random portion of the input image and uses a generative model to synthesize the masked image to a new image conditioned on the classification result. It then calculates the semantic difference between the original image and the synthesized one for OOD detection. Compared to existing solutions, MOODCAT naturally learns the semantic information of the in-distribution data with the proposed mask and conditional synthesis strategy, which is critical to identify OODs. Experimental results demonstrate that MOODCAT outperforms state-of-the-art OOD detection solutions by a large margin. Our code is available at <https://github.com/cure-lab/MOODCat>.

**Keywords:** OOD detection, Robust AI, Generative model

## 1 Introduction

Deep neural networks (DNNs) are trained under a “close-world” assumption [12, 23], where all the samples fed to the model are assumed to follow a narrow semantic distribution. However, when deployed in the wild, the model is exposed to an “open-world” with all kinds of inputs not necessarily following this distribution [9]. Such out-of-distribution (OOD) samples with significantly different semantics may mislead DNN models and generate wrong prediction results with extremely high confidence, thereby hindering DNN’s deployment safety [1, 7, 14, 15, 33].

To distinguish OOD samples from the in-distribution (In-D) data, some propose to reuse the features extracted from the original DNN model to tell the difference [15, 26–28, 42, 43]. However, such a feature-sharing strategy inevitably results in the trade-off between the prediction accuracy for In-D samples and the OOD detection capabilities. There are also various density-based OOD detection methods [3, 35, 36], which try to model the In-D data with probabilistic measures such as energy and likelihood. However, the trustworthiness of these measures is not guaranteed [21]. Another popular OOD detection mechanism uses generative models (e.g., variational autoencoder (VAE)) to reconstruct the input [6, 38]. Based on the assumption that In-D data can be well reconstructed while OODs

cannot since they are not seen during training, one could measure the distance between the original input and the reconstructed one and detect OOD with a threshold. However, this assumption is not sound. There are cases where OODs are faithfully reconstructed with the generative models, causing misjudgements [21].

In this paper, we propose a novel distance-based OOD detection framework, named *Masked OOD Catcher* (MOODCAT), wherein we consider the semantic mismatch under masking as the distance metric. Specifically, for image classifiers, we first randomly mask a portion of the input image, use a generative model to synthesize the masked image to a new image conditioned on the classification result, and then calculate the semantic difference between the original image and the synthesized one for OOD detection.

Our insight is that, the classification result carries discriminative semantic information and it imposes strong constraints onto the synthesis procedure, especially when trying to recover the masked portions. With MOODCAT, for correctly classified In-D data, the generative model can use the unmasked region to make up the masked part with sufficient training. In contrast, for OOD samples that are semantically different, the synthesized image based on the classification result tends to be dramatically different, especially for the masked region.

MOODCAT is a standalone OOD detector, and it does not require fine-tuning the original classifier. Consequently, it can be combined with any classifier to equip it with OOD detection capability without affecting its accuracy. We perform comprehensive evaluations on standard OOD detection benchmarks [44] with six datasets and four detection settings. Results show that our method can outperform state-of-the-art (SOTA) solutions by a large margin. We summarize the contributions of this work in the following:

- We propose a novel OOD detection framework by identifying semantic mismatch under masking, MOODCAT. To the best of our knowledge, this is the first work that *explicitly* considers semantics information for OOD detection.
- We present a novel masking and conditional synthesis flow in MOODCAT, and investigate various masking strategies and conditional generator designs for OOD detection.
- To tell the semantic difference between the original image and the synthesized one, we employ an anomalous scoring model composed of various quality assessment metrics (e.g., DISTS [8] and LPIPS [52]) and a newly-proposed conditional binary classifier.

The rest of the paper is constructed as follows. Section 2 surveys related OOD detection methods. We detail our proposed MOODCAT framework in Section 3. Section 4 presents our experimental results and the corresponding ablation studies. Finally, Section 5 concludes this paper.

## 2 Related Work

### 2.1 Existing OOD Detection Methods

In general, OOD detection methods can be categorized into: classification-based, density-based and distance-based methods [45].

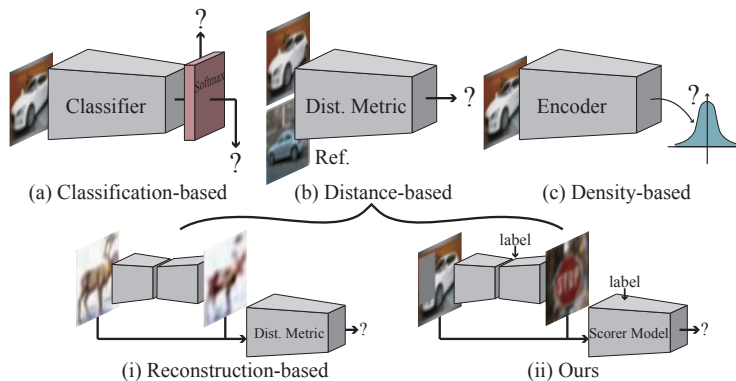


Fig. 1: Comparison of OOD Detection Methods. MOODCAT is a distance-based solution, and it relies on conditional image synthesis rather than reconstruction.

Classification-based methods derive OOD scores based on the output of DNNs, as shown in Fig. 1(a). Maximum Softmax Probability (MSP) [15] simply uses the maximum softmax probability as the indicator of In-D data. ODIN [26] applies a temperature scaling to the softmax value for OOD detection. Follow-up works include methods based on the output of DNNs [27, 28, 42, 43], the gradient of DNNs [19] and data generation or augmentation [40, 41]. Although simple to implement, most of them alter the training process of the original classifier, thereby reducing the classification accuracy for In-D samples.

Density-based methods usually apply some probabilistic models for the distribution of In-D samples and regard test data in low-density regions as OOD [3, 25, 36], as shown in Fig. 1(c). Some methods in this category also resort to generative models [35] to learn the distribution of data. However, recent research found that the learned density model may assign high likelihood value to some OODs, since the obtained likelihood could be dominated by low-level features such as location and variance instead of the high-level semantics, which is related to the specific network architecture and data used for learning [3, 30].

Distance-based methods consider that OODs should be relatively far away from In-Ds. They either calculate the centroids of In-D classes in the feature space [18, 49] (Fig. 1(b)) or reconstruct the input itself (Section. 2.2, Fig. 1(i)) for OOD detection. However, for high-level semantic features, their assumption for distance disparity may not hold, and high reconstruction quality cannot ensure In-Ds. In this paper, we use conditional synthesis on masked images to highlight the semantics difference in the image space.

## 2.2 Reconstruction-based OOD Detection

Reconstruction-based methods, which fall into the category of distance-based methods, are closely related to the proposed MOODCAT technique. These methods are based on the assumption that In-D data can be well-reconstructed from a

trained generative model, but OOD cannot as they are not seen during training (see Fig. 1(i)). Previous reconstruction-based detectors generally distinguish OOD samples by comparing pixel-level quality “degradation” of the reconstruction for given input [6, 38]. However, without prior-knowledge about OOD samples, there is no guarantee for such quality degradation. In contrast, MOODCAT tries to synthesize In-D images instead of reconstructing the inputs, which is in line with the objective of the generative model.

The framework of MOODCAT (Fig. 1(ii)) is inspired by [46], which detects adversarial examples (AE) by generating synthesized images conditioned on the output of the misled classifier. AE detection is quite different from OOD detection because adversarial examples are In-D samples with imperceptible perturbations. The classification label itself is sufficient to train the generative model to differentiate AEs and benign samples. This is not the case for OOD samples, which motivates the proposed MOODCAT solution for OOD detection, as detailed in Section 3.

### 2.3 OOD Detection with External OOD Data

Recently, some researchers propose to involve data from other datasets to simulate OOD samples for model training. Representative “OOD-aware” techniques include Outlier Exposure (OE) [16], Maximum Classifier Discrepancy (MCD) [48], and Unsupervised Dual Grouping (UDG) [44]. OE relies on large-scale purified OOD samples, whereas MCD and UDG only need extra unlabeled data, which contains both In-D and OOD data. However, all of them are classification-based methods, where external data are used to train a modified classifier model. Following the same unlabeled extra data setting of MCD and UDG, we present that including extra training data (In-D, OOD mixture) into the training process can further improve MOODCAT’s performance. Since our MOODCAT works independently with the original classifiers, MOODCAT will not degrade the accuracy of the original classifiers. We provide the detailed description in Section 3.6 and experimental results in Section 4.

### 2.4 Open Set Recognition

A similar problem to OOD detection is the so-called Open Set Recognition (OSR) problem, which aims to distinguish the known and unknown classes [34, 45]. Several existing works [10, 11, 31, 34] targeted on OSR also employ generative models, whereas differ from MOODCAT significantly. Specifically, OSRCI [31] uses Generative Adversarial Network (GAN) as data augmentation to train the classifier; C2AE [34] and CVAECapOSR [11] are conditional VAE/Autoencoder-based detectors. C2AE identifies outliers based on reconstruction errors, and requires K-time inference to give the final decision. In contrast, MOODCAT infers once and makes the decision based on semantic contradiction. CVAECapOSR use Conditional VAE (CVAE) to model the distribution of In-D samples and detect outliers at latent space (i.e., without using the generator), whereas MOODCAT detects at image space.




Fig. 2: Pipeline of MOODCAT. We first mask a portion of the input image. Next, a generative model synthesizes the masked image to a new image conditioned on  $y$ , and then an anomalous scoring model measures the semantic difference between the input image and the synthesized one for OOD detection.

### 3 Proposed Method

#### 3.1 Design Goals

This work considers the scenario where we have an In-D data trained classifier ( $\mathcal{C}$ ), which needs to be deployed in the wild. Consequently, the classifier will be threatened by OOD samples. We aim at building an OOD detection method, which can identify the OOD samples effectively without compromising the classification accuracy of the classifier. Our model is assumed to have access to the predicted label,  $y$ , but we do not modify any part of the classifier, including but may not limit to the architecture and the trained weights. As a result, our method can be a plug-and-play detector that easily cooperates with classifiers.

#### 3.2 Method Overview

As pointed out by [45], OOD samples ( $x_o \in \mathcal{O}$ ) are defined by label-shifted samples or samples with non-overlapping labels w.r.t the training data, or In-Ds ( $x_{in} \in \mathcal{I}$ ). Hence, the semantics of any OOD sample contradicts with any In-D sample. This is the observation that motivates us to design a framework for OOD detection by spotlighting their semantic discrepancy.

Fig. 2 depicts the overview of our method. The proposed Masked Out-of-Distribution Catcher (MOODCAT) contains three stages: *randomly masking*, *generative synthesis* and *scoring*. Specifically, we first randomly mask the input image  $x$  as  $x_m = \mathbf{M}(x)$ , where  $\mathbf{M}(\cdot)$  indicates the randomly masking operation. Then, we apply a generative model,  $\mathbf{G}$ , to synthesize a new image,  $x'$ , by taking  $x_m$  as the template and conditioning on the label  $y$ . Finally, we apply an anomalous scoring model to judge the discrepancy between the input and its synthesis.

Through masking,  $x_m$  will partially lose its original semantic meaning, and thus leave more space for  $\mathbf{G}$  to synthesize new content. With  $y$  as the condition,

the newly synthesized content should be consistent with the semantic meaning indicated by  $y$ . Note that, we use the ground truth label  $y$  to train the generative model and use the output from the classifier when inference, i.e.,  $y = \mathcal{C}(x)$ .

Here, we analyze different situations with In-D or OOD samples. On the one hand, if an In-D sample  $x_{in}$  comes, the predicted label  $y$  matches  $x_{in}$ 's intrinsic semantic meaning appropriately. Although the input image,  $x_{in-m}$ , is partially masked, there should be some visual clue related to its semantic meaning, e.g., wings of a bird or paws of a dog. As a result,  $\mathbf{G}$  can synthesize  $x'_{in}$  quite faithful to  $x_{in}$ . As exemplified in the upper half of Fig. 2, the synthesis of the "stop" sign can be very close to the original input. On the other hand, when it comes to an OOD sample  $x_o$ , the predicted label provided by the classifier is irreverent to  $x_o$ 's semantic meaning. Even with  $x_o$  as a template, the generative model will try its best to synthesize contents related to the semantic label. As a result, the mismatch of semantic meaning between input and label can be spotlighted by the discrepancy between input and its synthesis. As the example shows in the lower part in Fig. 2, if an OOD sample (*car*) is wrongly predicted as a "stop" sign. The synthesis will be highly related to the "stop" sign rather than the original image.

Through such conditional synthesis, we can spotlight the discrepancy caused by OOD samples. Thus OODs can be easily distinguished by comparing the pair of input with its synthesis,  $(x, x')$ .

### 3.3 Masking Mechanism

In MOODCAT, the generative model uses the input image as a template and synthesizes an image with the same semantic meaning as the given label. A high-quality synthesis can better highlight the contradiction. However, due to the intrinsic contradiction between the input image and the label for OOD, too much information from the input image can degrade the quality of generation. Therefore, we propose to apply masking on  $x$  to remove some redundant information while leaving more space for the generative model to synthesize.

The use of masking follows the key motivation of MOODCAT in OOD detection, which applies generation for synthesis rather than reconstruction. Previous reconstruction-based methods (e.g., [6, 38]) tend to reconstruct the image based on pixel-level dependency. In practice, the assumption that an OOD sample cannot be reconstructed well may not hold since they do not consider any semantics. However, our generative model aims at semantic synthesis. The masking mechanism can cooperate with the predicted label from the classifier to spotlight the contradiction caused by OOD.

The contribution of randomly masking is twofold: **1)** masking the input image can encourage the generative model to better depict the semantic meaning of the given label on the synthesis, especially to an OOD sample; **2)** masking, as a typical data augmentation method, can encourage the encoder to summarize the features of the input from a holistic perspective, thus improve the quality of synthesis, especially when synthesizing with In-Ds as templates. Obviously, the above two aspects both contribute to apart the behavior of In-D and OOD. As a result, a large discrepancy lies in the OOD sample and its synthesis.

### 3.4 Generative Model

The Generative model is responsible for generating a synthesis by taking both the masked input  $x_m$  and the pre-assigned semantic label  $y$  into consideration. As shown in Fig. 2, we select the Encoder (**E**) and Decoder (**D**) architecture as the generative model, i.e.  $\mathbf{G} = \mathbf{E} \cdot \mathbf{D}$ . This architecture is inspired by [46]. The encoder **E** acts as a feature extractor (as shown by the gray part in Fig. 2). By taking the masked image  $x_m$  as input, **E** is expected to capture necessary low-level features and encoder them as a latent vector,  $z = \mathbf{E}(x_m)$ . As done by VAE [20], we use the KL Divergence to regulate the latent vector  $z$ , which can be formulated as Eq. (1).

$$\mathcal{L}_{KLD} = D_{KL}[\mathcal{N}(\mu(x_m), \Sigma(x_m)) \parallel \mathcal{N}(0, 1)], \quad (1)$$

where  $\mathcal{N}(\mu, \Sigma)$  indicates the Gaussian distribution with respect to  $\mu$  and  $\Sigma$ . We use the reparameterization trick from VAE on the latent variable  $z$  during training,  $z = \mu(x_m) + \Sigma(x_m) \cdot \epsilon$ , where  $\epsilon \sim \mathcal{N}(0, 1)$ .

The decoder, **D**, is trained to generate a synthesis  $x' = \mathbf{D}(z, y)$ . The given semantic label  $y$  is used to control the semantic meaning of the synthesis, while  $z$  is used to provide low-level features from the template image  $x$ . This synthetic target is fulfilled through the class-conditional batch normalization layer [5]. This layer is usually used in conditional image generation [29, 51]. Since the normalization is determined by the given semantic label  $y$ , the semantic meaning of the synthesis can be highly dependent on it. As a result, if the semantic meaning of  $x$  is consistent with  $y$  (in the case of In-D samples), the synthesis can be highly close to  $x$ . However, if input an OOD sample, the semantic contradiction between the input image and the label will lead the synthesis to be far away from the input image, thus spotlighting the contradiction.

We implement **D** based on the generator architecture proposed in [2]. We apply the classic  $\ell_1$ ,  $\ell_2$  and  $\mathcal{SSIM}$  [37] as part of loss items to constrain that  $x'$  resembles  $x$ . Furthermore, we adopt the U-net based discriminator [39] to operate an adversarial loss on the training process to further improve the quality of the synthetic image. Compared with the vanilla discriminator, this U-net based one can additionally provide a per-pixel real/fake map to locate the fake parts in the image. Therefore, the generative model can be trained to focus on both local and global features with more realistic details. Due to space limitations, we detailed the training process and corresponding objective functions of **G** in Appendix.

### 3.5 Anomalous Scoring Model

As analyzed in the former sections, MOODCAT can generate high-quality syntheses in terms of similarity for In-D samples, but not for OOD samples. To distinguish OODs from In-Ds, we develop an anomalous scoring model. The proposed anomalous scoring model is built on two types of scorers: one is the *conditional binary classifier*, and the other is *Image Quality Assessment models* (IQA) [52]. Both can provide assessments of the syntheses by referring to input images.



**Conditional Binary Classifier.** Identifying the semantic mismatch lies between OOD and its synthesis can be seen as a binary classification task. With the trained generative model, we can train a binary classifier for this in a supervised way. This is feasible because the binary classifier can learn to identify OODs by the similarity between the given image and its conditional synthesis. We also provide the semantic label to the classifier for judgement, this can further ease the distinguishing procedure.

Note that, we do not rely on OOD data through training. To mitigate the lack of OOD samples during training, we leverage the In-D samples with the synthetic results under mismatched labels to simulate the behavior of OOD samples. As a result, the binary classifier can learn to identify the semantic mismatch, which is the spotlighted feature for OOD samples. Moreover, we condition the binary classifier on the semantic label via the projection layer [29]. With the prior knowledge of the class, the binary classifier can learn a fine-grained decision boundary for each class, leading to better performance.

Another good property of this training strategy is that there is no need for specific information from the DNN model, here the classifier ( $\mathcal{C}$ ), to be protected. Actually, the judgement of the binary classifier is only conditioned on given label, which is independent of the DNN model. Therefore, the trained conditional binary classifier can fit various DNN models in a plug-and-play manner.

Fig. 3 demonstrates the training process of the proposed conditional binary classifier,  $\mathcal{C}_b$ . *hinge loss* serves as the objective, as formulated in Eq. (2), where the  $x'_y$  indicates the synthesis generated under the groundtruth label  $y$  of  $x$ , and the  $(x, x'_y)$  is used as the positive pair.  $x'_{y'}$  depicts the synthesis generated under a randomly sampled mismatched semantic label  $y' \neq y$ , then  $(x, x'_{y'})$  is used as the negative pair during training. During inference,  $\mathcal{C}_b$  is used to score the input image pair. The score can be used to flag OOD samples by a given threshold.




Fig. 3: Training pipeline of the proposed Conditional Binary Classifier.

$$\mathcal{L}_{\mathcal{C}_b} = \text{ReLU}(1 - \mathcal{C}_b((x, x'_y), y)) + \text{ReLU}(1 + \mathcal{C}_b((x, x'_{y'}), y')) \quad (2)$$

**Image Quality Assessment Models.** IQA models [52] are widely adopted to evaluate the perceptual quality of a synthesis by referring to the source image in many computer vision tasks, such as denoising, super-resolution and compression. Here, we apply IQA models as the perceptual metric for the quality of synthesis, forming part of our anomalous scoring model. Since our generative model has already highlighted the contradiction caused by OOD through conditional synthesis, IQA models can be directly applied for detection.



In all, the *Conditional Binary Classifier* and *IQA* models work in a cascade way, where any scorer flags an OOD can lead to the final rejection. Different scoring mechanisms can evaluate the quality of generation from different perspectives, thus supporting each other for better performance than anyone alone.

### 3.6 Learning with Unlabeled Data

Recently, researchers have shown that including OOD data into the training process can improve the performance of OOD detection [16, 48]. However, they may rely on well labeled data that need to manually identify In-D and OOD elements<sup>1</sup>, e.g., [16]. On the contrary, unlabeled data can be easily collected from various sources with little cost. Actually, MOODCAT can be further improved with unlabeled data.

The Conditional Binary Classifier introduced in Section 3.5 is trained by treating In-D data with synthesis of mismatched semantic labels as negative pairs. This procedure can be directly combined with external data. Note that, our negative pair only require that the semantic label used for synthesis is different from that of input image. To eliminate possible In-D samples from unlabeled data, we can apply a classifier trained with In-D data to generate pseudo label for unlabeled data. Then, during training, when we randomly sample the mismatched semantic label, we uniformly sample from all possible labels other than the pseudo label. In this way, the sampled semantic label is ensured to be mismatched for both OOD and some possible In-D data from the unlabeled data. Therefore, all of them can be utilized correctly to improve the performance of MOODCAT.

## 4 Experiments

### 4.1 Evaluation Settings

**Benchmarks.** We evaluate MOODCAT on the most recent semantic OOD detection benchmarks, SC-OOD benchmarks [44]. SC-OOD benchmarks provide extensive semantic-level OOD detection settings for evaluation. Specifically, from SC-OOD, images from different datasets are filtered to ensure that only those containing different semantic meanings are considered as OOD samples. SC-OOD focus on the semantic difference between samples, thus being more practical for the real-world model deployment than other previous OOD benchmarks [15, 16, 26], which are built by setting one dataset as In-D and all others as OOD.

**Datasets.** Following the settings in [44], we employ CIFAR-10 [22], CIFAR-100 [22] as In-D samples, respectively, and others as OOD samples. When setting CIFAR-10 as In-D, we employ six datasets as OOD datasets, including SVHN [32], CIFAR-100 [22], TEXTURE [4], PLACES365 [53], LSUN [47] and TINY-IMAGENET [24]. For CIFAR-100 benchmarks, the OOD datasets are the same as that of CIFAR-10, except for swapping CIFAR-100 for CIFAR-10 as OOD.

<sup>1</sup> Note that, images from another dataset are not necessarily to be OOD w.r.t semantic meaning [44].

## 4.2 Evaluation Metrics

We employ FPR@TPR95%, AUROC, AUPR, and Classification Accuracy as the evaluation metrics following [28, 44]. In this paper, unless otherwise specified, we denote the In-D as Positive (P), and the OOD as Negative (N).

**FPR@TPR95%** presents the False Positive Rate (FPR) when the True Positive Rate (TPR) equals 95%. This metric reflects the ratio of falsely identified OOD when most of In-D samples are correctly recognized.

**AUROC.** The Area Under Receiver Operating Characteristic curve (AUROC) is an overall evaluation metric to reflect the detection capability of a detector.

**AUPR-In** AUPR calculates the Area Under the Precision-Recall curve. AUPR is a complementary metric that reflects the impact of imbalanced datasets. For AUPR-In metric, In-D samples are denoted as positive samples.

**AUPR-Out** indicates the same measure as AUPR-In mentioned above, whereas the OOD samples are deemed as positive during calculating AUPR-Out.

**Classification Accuracy** presents the classifier’s performance on the In-D samples. It indicates the impact on the original classifier caused by OOD detector.

## 4.3 Experimental Results

We evaluate MOODCAT with two settings: **1)** MOODCAT trained with In-D dataset only; **2)** MOODCAT trained with external unlabeled data (Section 3.6). We report the results in Table 1 and Table 2, respectively. Experiments are performed with ResNet18 [13] classifier<sup>2</sup> for fair comparison.

**Main Results.** As shown in Table 1 and Table 2, experimental results indicate that MOODCAT outperforms or at least on par with SOTA methods on CIFAR-10/CIFAR-100 benchmarks without/with external training data. Since our method detects OODs relying on their semantic-level mismatching instead of low-level distribution shift, the performance of MOODCAT is stable across various OODs. As a plug-and-play model, MOODCAT causes no classifier performance degradation.

**CIFAR-10 Benchmark.** Table 1 reports the detection performance of ODIN [26], EBO [28] and MOODCAT on CIFAR-10 benchmarks. ODIN and EBO are developed without external OOD data. When comparing these two methods, we implemented MOODCAT in the same setting. We additionally report results with external data for training to show the improvements. As shown in Table 1, MOODCAT outperforms ODIN and EBO in most cases. For example, for the AUROC, which reflects the overall performance of a detector, MOODCAT outperforms baselines on all six OOD benchmarks. Furthermore, MOODCAT presents a more stable performance across all OOD datasets. As for statistics, we report the standard deviation (Std) for each metric<sup>3</sup>. As in Table 1, the Std of MOODCAT can be much lower than baselines. To be more specific, EBO performs quite well when encounter OODs sourcing from SVHN (AUROC=92.08%), but when it comes to

<sup>2</sup> For more results under other classifier architectures (WRN28 [50], DenseNet [17]), please refer to Appendix.

<sup>3</sup> FPR@TPR95% is only a single point on the PR curve. It may not reflect the overall performance in terms of standard deviation.

Table 1: OOD Detection Performance on CIFAR-10 as In-D without using external OOD data for training. All the values are in percentages.  $\uparrow/\downarrow$  indicates higher/lower value is better. The best results are in **bold**. We also add our results with external data in gray.

Detection Methods	OOD	FPR@TPR95% $\downarrow$	AUROC $\uparrow$	AUPR In $\uparrow$	AUPR Out $\uparrow$	Classification Accuracy $\uparrow$
ODIN [26]	SVHN	52.27	83.26	63.76	92.60	<b>95.02</b>
	CIFAR-100	61.19	78.40	73.21	80.99	<b>95.02</b>
	TINY-IMAGENET	59.09	79.69	79.34	77.52	<b>92.54</b>
	TEXTURE	42.52	84.06	86.01	80.73	<b>95.02</b>
	LSUN	47.85	84.56	81.56	85.58	<b>95.02</b>
	PLACES365	53.94	82.01	54.92	93.30	<b>93.87</b>
	<b>Mean/Std</b>	52.00	82.00/2.48	73.13/11.79	85.12/6.59	<b>94.42/1.03</b>
EBO [28]	SVHN	<b>30.56</b>	92.08	80.95	96.28	<b>95.02</b>
	CIFAR-100	56.98	79.65	75.09	81.23	<b>95.02</b>
	TINY-IMAGENET	57.81	81.65	81.80	78.75	<b>92.54</b>
	TEXTURE	52.11	80.70	83.34	75.20	<b>95.02</b>
	LSUN	50.56	85.04	82.80	85.29	<b>95.02</b>
	PLACES365	52.16	83.86	58.96	93.90	<b>93.87</b>
	<b>Mean/Std</b>	50.03	83.83/4.51	77.16/9.40	85.11/8.44	<b>94.42/1.03</b>
Ours	SVHN	37.72/24.27	<b>92.99</b> /95.93	<b>87.43</b> /92.98	<b>96.70</b> /98.05	<b>95.02</b>
	CIFAR-100	<b>42.32</b> /39.92	<b>89.88</b> /91.45	<b>89.75</b> /91.54	<b>90.24</b> /91.73	<b>95.02</b>
	TINY-IMAGENET	<b>40.60</b> /32.41	<b>90.57</b> /93.34	<b>90.59</b> /93.63	<b>90.76</b> /93.41	<b>92.54</b>
	TEXTURE	<b>26.12</b> /6.86	<b>94.15</b> /98.69	<b>96.33</b> /99.29	<b>91.68</b> /97.71	<b>95.02</b>
	LSUN	<b>43.86</b> /33.31	<b>90.61</b> /93.40	<b>91.07</b> /93.85	<b>90.02</b> /93.22	<b>95.02</b>
	PLACES365	<b>42.34</b> /35.51	<b>90.16</b> /92.77	<b>75.28</b> /82.25	<b>96.83</b> /94.82	<b>93.87</b>
	<b>Mean</b>	<b>38.83</b> /28.71	<b>91.39</b> /94.27	<b>88.40</b> /92.26	<b>92.71</b> /94.82	<b>94.42</b>
	<b>Std</b>	-	<b>1.75</b> /2.61	<b>7.07</b> /5.57	<b>3.20</b> /2.56	<b>1.03</b>

CIFAR-100, the AUROC drops by 12% to 79.65%. On the contrary, the AUROC for MOODCAT are all around a high mean value. This stability may due to that MOODCAT utilize semantic information, which is exactly the definition of OOD. Comparing to the classification features or other low-level features, the semantic contradiction exists more generally.

As discussed in Section 3.6, MOODCAT can be equipped with external OOD data for better detection ability. In Table 1, we report the performance for MOODCAT using TINY-IMAGENET as external unlabeled training data in gray. These results evidence the effectiveness of the training strategy with external data for MOODCAT, where significant performance improvements can be found. Due to space limitation, for this setting, we only compare results with baselines on CIFAR-100 (see below) and leave that on CIFAR-10 benchmarks in Appendix. **CIFAR-100 Benchmark.** OE [16], MCD [48] and UDG [44] rely on TINY-IMAGENET as external OOD data for training. We use them as baselines to show the effectiveness of MOODCAT under this setting. Note that MOODCAT treats all external data as unlabeled during training.

Table 2 shows the results on CIFAR-100 benchmarks. As can be seen, though taking advantage of external data, MCD, OE and UDG suffer from the complicated CIFAR-100 dataset. By contrast, MOODCAT’s performance on CIFAR-100 is comparable with that on CIFAR-10 against same OOD shown in Table 1. This is due to MOODCAT’s detection mechanism, which relying on the semantic mis-

Table 2: OOD Detection Performance on CIFAR-100 as In-D with TINY-IMAGENET as external data for training. All the values are in percentages.  $\uparrow/\downarrow$  indicates higher/lower value is better. The best results are in **bold**. We also add our results without external data in gray.

Detection Methods	OOD	FPR@TPR95% $\downarrow$	AUROC $\uparrow$	AUPR In $\uparrow$	AUPR Out $\uparrow$	Classification Accuracy $\uparrow$
MCD [48]	SVHN	85.82	76.61	65.50	85.52	68.80
	CIFAR-10	87.74	73.15	76.51	67.24	68.80
	TINY-IMAGENET	84.46	75.32	85.11	59.49	62.21
	TEXTURE	83.97	73.46	83.11	56.79	68.80
	LSUN	86.08	74.05	84.21	58.62	67.51
	PLACES365	82.74	76.30	61.15	87.19	70.47
	<b>Mean/Std</b>	85.14	74.82/1.47	75.93/10.31	69.14/13.81	67.77/2.88
OE [16]	SVHN	68.87	84.23	75.11	91.41	70.49
	CIFAR-10	79.72	78.92	81.95	74.28	70.49
	TINY-IMAGENET	83.41	76.99	86.36	60.56	63.69
	TEXTURE	86.56	73.89	84.48	54.84	70.49
	LSUN	83.53	77.10	86.28	60.97	69.89
	PLACES365	78.24	79.62	67.13	88.89	72.02
	<b>Mean/Std</b>	80.06	78.46/3.46	80.22/ <b>7.66</b>	71.83/15.58	69.51/2.94
UDG [44]	SVHN	60.00	88.25	<b>81.46</b>	93.63	68.51
	CIFAR-10	83.35	76.18	78.92	71.15	68.51
	TINY-IMAGENET	81.73	77.18	86.00	61.67	61.80
	TEXTURE	75.04	79.53	87.63	65.49	68.51
	LSUN	78.70	76.79	84.74	63.05	67.10
	PLACES365	73.89	79.87	65.36	89.60	69.83
	<b>Mean/Std</b>	75.45	79.63/4.48	80.69/8.14	74.10/14.01	67.38/ <b>2.87</b>
Ours	SVHN	58.16/ <b>51.60</b>	87.38/ <b>88.99</b>	78.25/80.89	93.81/ <b>94.81</b>	<b>76.65</b>
	CIFAR-10	54.31/ <b>50.17</b>	85.91/ <b>87.76</b>	86.27/ <b>88.18</b>	85.91/ <b>87.79</b>	<b>76.65</b>
	TINY-IMAGENET	55.33/ <b>46.07</b>	86.95/ <b>89.42</b>	87.55/ <b>89.73</b>	86.67/ <b>89.28</b>	<b>69.56</b>
	TEXTURE	46.70/ <b>42.22</b>	89.20/ <b>90.56</b>	93.48/ <b>94.43</b>	83.28/ <b>85.13</b>	<b>76.65</b>
	LSUN	53.43/ <b>47.85</b>	87.98/ <b>89.96</b>	88.82/ <b>90.33</b>	87.32/ <b>89.23</b>	<b>76.10</b>
	PLACES365	54.20/ <b>47.72</b>	87.41/ <b>89.30</b>	71.68/ <b>74.83</b>	95.78/ <b>96.48</b>	<b>77.56</b>
	<b>Mean</b>	53.69/ <b>47.61</b>	87.47/ <b>89.33</b>	84.34/ <b>86.40</b>	88.80/ <b>90.45</b>	<b>75.53</b>
	<b>Std</b>	-	0.95/ <b>1.09</b>	7.19/7.94	4.33/ <b>4.89</b>	2.96

match. Even when the In-D dataset becomes complicated, the semantic mismatch in OODs is still obvious. For methods detect OODs based on DNN-extracted features, e.g. MCD, OE, UDG, they may suffer from poor decision boundaries as the number of classes increases. Similarly as in the case on CIFAR-10, to show the effectiveness of the proposed training strategy with external data, we also report the performance for MOODCAT without external data in gray. The improvement can be seen on all metrics across all OOD settings. For space limitation, we leave the comparison on CIFAR-100 benchmarks without external data in Appendix.

#### 4.4 Comparison with Open Set Recognition Methods

The experimental protocol in OSR is to randomly selecting  $K$  classes from a specific  $n$ -class dataset as “known” classes and the left  $n - K$  as “unknown” classes. To make a fair comparison, we retrain MOODCAT under the experimental settings in CVAECapOSR [11], where only 4 or 6 classes from CIFAR-10 are used for training. We report AUROC scores in Table 3, and the results for other methods

Table 3: Comparison with Open Set Recognition methods. AUROC scores on the detection of known and unknown classes. CIFAR indicates splitting CIFAR-10 to 6 known classes, and 4 unknown. CIFAR + $N$  samples known 4 classes form CIFAR-10,  $N$  unknown classes from CIFAR-100. The **bold** indicates the best. For more details about dataset splits, please refer to CVAECapOSR [11].

Method	CIFAR	CIFAR +10	CIFAR +50
OSRCI [31]	69.9 $\pm$ 3.8	83.8	82.7
C2AE [34]	71.1 $\pm$ 0.8	81.0 $\pm$ 0.5	80.3 $\pm$ 0.0
CVAECapOSR [11]	83.5 $\pm$ 2.3	88.8 $\pm$ 1.9	88.9 $\pm$ 1.7
<b>MOODCAT (ours)</b>	<b>89.48</b> $\pm$ 0.50	<b>89.36</b> $\pm$ 0.74	<b>89.23</b> $\pm$ 0.19

Table 4: OOD Detection Performance under different combinations of anomalous scorers. MOODCAT is trained on CIFAR-10 (In-D) without external data. OODs are from CIFAR-100. All the values are in percentages.  $\uparrow/\downarrow$  indicates the higher/lower value is better. The best results are in **bold**.

Anomalous Scorer	FPR@TPR95% $\downarrow$	AUROC $\uparrow$	AUPR-In $\uparrow$	AUPR-Out $\uparrow$
$\mathcal{C}_b$	42.80	89.13	88.58	89.85
$LPIPS$	76.62	73.93	72.68	73.23
$DISTS$	82.03	72.14	71.83	70.35
$\mathcal{C}_b + LPIPS$	42.14	89.49	89.18	90.11
$\mathcal{C}_b + DISTS$	42.31	89.35	89.09	89.98
$LPIPS + DISTS$	76.06	74.78	74.25	73.86
$\mathcal{C}_b + LPIPS + DISTS$	<b>41.95</b>	<b>89.57</b>	<b>89.30</b>	<b>90.16</b>

are from CVAECapOSR [11]. As can be seen, MOODCAT outperforms these methods in all three settings, especially in the original CIFAR setting, where MOODCAT outperforms the second best method about 6%.

#### 4.5 Ablation Study

In this section, we first analyze the effectiveness of every anomalous scorer and how scorers cooperate to achieve the final decent performance. Then we conduct experiments on masking, and give insights on masking’s effectiveness.

**Anomalous scoring model.** Table 4 summarizes the performances of every scorer, every two scorers and all three scorers working together. As can be seen in Table 4, the  $\mathcal{C}_b + LPIPS + DISTS$  combination wins the best detection performance in terms of all evaluation metrics, which means our proposed  $\mathcal{C}_b$  has a high flexibility to cooperate with IQA models (see Section 3.5). We can also observe that coupling scorers usually lead to a better detection capability than that of any single scorer within the coupling. However, adding extra scorers inevitably increases computational and memory overheads. The cost of basic version of MOODCAT, i.e.  $\mathcal{C}_b$ , **E**, **D**, is relatively small, i.e. Params 4.552M, MACs 0.408G, when compared to that of the widely adopted classifier architectures, e.g., ResNet50 (Params 23.251M, MACs 1.305G). Due to space limitations, we detail the computational and memory costs of MOODCAT and those of the baselines in

Table 5: Ablation study on Masking. We set the masking ratio as 0.3 for “Fixed High Ratio” and “Patched”, 0.1 for “Fixed Low Ratio”, and that of “Randomly” varies from 0.1 to 0.3. MOODCAT employs the **Randomly** masking style.

Training	Inference	FPR@95 ↓	AUROC ↑	AUPR In ↑	AUPR Out ↑
w/o mask	w/o mask	40.67	90.79	90.91	90.88
	Randomly	38.57 $\pm$ 0.75	90.99 $\pm$ 0.18	90.85 $\pm$ 0.21	91.31 $\pm$ 0.19
with mask	w/o mask	37.74	91.52	91.5	91.77
	Fix Low Ratio	37.5	91.64	91.6	91.9
	Fix High Ratio	36.16	91.44	91.03	91.9
	<b>Randomly</b>	<b>36.15<math>\pm</math>0.94</b>	<b>91.78<math>\pm</math>0.17</b>	<b>91.68<math>\pm</math>0.24</b>	<b>92.08<math>\pm</math>0.14</b>

Table 6: Ablation study on label conditioning

In-Data	Methods	FPR@95 ↓	AUROC ↑	AUPR In ↑	AUPR Out ↑
Cifar10 In-D	<b>ours</b>	<b>41.95</b>	<b>89.57</b>	<b>89.30</b>	<b>90.16</b>
Cifar100 OOD	uncond.	86.94	63.62	71.26	63.01
Cifar100 In-D	<b>ours</b>	<b>50.17</b>	<b>87.76</b>	<b>88.18</b>	<b>87.79</b>
Cifar10 OOD	uncond.	96.74	52.48	56.29	71.24

the Appendix. In practice, MOODCAT can achieve appropriate detection ability, by tailoring scorers in MOODCAT’s anomalous scoring model according to the application scenario, i.e. trade-off between the performance and costs, and we discussed this part in Appendix.

**Masking.** To evidence the effectiveness of masking holistically, we conduct another ablation study without masking in training. As shown in Table 5, the masking scheme indeed plays an important role in the performance of MOODCAT. Due to space limitation, we leave the ablation study on masking style in Appendix.

**Label Conditioning.** As for evaluating label conditioning, we degrade the cGAN to vanilla GAN without conditions. Table 6 reports this ablation study. We can see that our conditioning mechanism outperforms the unconditioned scheme by a large margin. As analyzed in Sec 3.2, the conditioning spotlights the semantic discrepancy between In-D and OOD to facilitate OOD detection.

## 5 Conclusion

In this paper, we propose a novel plug-and-play OOD detection method for image classifiers, MOODCAT, wherein we consider the semantic mismatch under masking as the distance metric. MOODCAT naturally learns the semantic information from the in-distribution data with the proposed mask and conditional synthesis framework. Experimental results demonstrate significantly better OOD detection capabilities of MOODCAT over SOTA solutions.

**Acknowledgements.** We appreciate the reviewers for their thoughtful comments and efforts towards improving this paper. This work was supported in part by General Research Fund of Hong Kong Research Grants Council (RGC) under Grant No. 14203521 and No. 14205420.

## References

1. Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., Mané, D.: Concrete problems in ai safety. arXiv preprint arXiv:1606.06565 (2016)
2. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. In: International Conference on Learning Representations (2018)
3. Choi, H., Jang, E., Alemi, A.A.: Waic, but why? generative ensembles for robust anomaly detection. arXiv preprint arXiv:1810.01392 (2018)
4. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2014)
5. De Vries, H., Strub, F., Mary, J., Larochelle, H., Pietquin, O., Courville, A.C.: Modulating early visual processing by language. *Advances in Neural Information Processing Systems* **30** (2017)
6. Denouden, T., Salay, R., Czarnecki, K., Abdelzad, V., Phan, B., Vernekar, S.: Improving reconstruction autoencoder out-of-distribution detection with mahalanobis distance. arXiv preprint arXiv:1812.02765 (2018)
7. Dietterich, T.G.: Steps toward robust artificial intelligence. *AI Magazine* **38**(3), 3–24 (2017)
8. Ding, K., Ma, K., Wang, S., Simoncelli, E.P.: Image quality assessment: Unifying structure and texture similarity. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020)
9. Drummond, N., Shearer, R.: The open world assumption. In: eSI Workshop: The Closed World of Databases meets the Open World of the Semantic Web. vol. 15 (2006)
10. Ge, Z., Demyanov, S., Chen, Z., Garnavi, R.: Generative openmax for multi-class open set classification. In: British Machine Vision Conference 2017. British Machine Vision Association and Society for Pattern Recognition (2017)
11. Guo, Y., Camporese, G., Yang, W., Sperduti, A., Ballan, L.: Conditional variational capsule network for open set recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 103–111 (2021)
12. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1026–1034 (2015)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
14. Hein, M., Andriushchenko, M., Bitterwolf, J.: Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 41–50 (2019)
15. Hendrycks, D., Gimpel, K.: A baseline for detecting misclassified and out-of-distribution examples in neural networks. In: 5th International Conference on Learning Representations (ICLR) (2017)
16. Hendrycks, D., Mazeika, M., Dietterich, T.: Deep anomaly detection with outlier exposure. In: International Conference on Learning Representations (2018)
17. Huang, G., Liu, Z., Pleiss, G., Van Der Maaten, L., Weinberger, K.: Convolutional networks with dense connectivity. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 1–1 (2019)



18. Huang, H., Li, Z., Wang, L., Chen, S., Zhou, X., Dong, B.: Feature space singularity for out-of-distribution detection. In: *Proceedings of the Workshop on Artificial Intelligence Safety (SafeAI)* (2021)
19. Huang, R., Geng, A., Li, Y.: On the importance of gradients for detecting distributional shifts in the wild. *Advances in Neural Information Processing Systems* **34** (2021)
20. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: *2nd International Conference on Learning Representations (ICLR)* (2014)
21. Kirichenko, P., Izmailov, P., Wilson, A.G.: Why normalizing flows fail to detect out-of-distribution data. *Advances in neural information processing systems* **33**, 20578–20589 (2020)
22. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
23. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*. vol. 25, pp. 1106–1114 (2012)
24. Le, Y., Yang, X.: Tiny imagenet visual recognition challenge. *CS 231N* **7**(7), 3 (2015)
25. Lee, K., Lee, K., Lee, H., Shin, J.: A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems* **31** (2018)
26. Liang, S., Li, Y., Srikant, R.: Enhancing the reliability of out-of-distribution image detection in neural networks. In: *6th International Conference on Learning Representations, ICLR 2018* (2018)
27. Lin, Z., Roy, S.D., Li, Y.: Mood: Multi-level out-of-distribution detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15313–15323 (2021)
28. Liu, W., Wang, X., Owens, J., Li, Y.: Energy-based out-of-distribution detection. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 21464–21475. Curran Associates, Inc. (2020)
29. Miyato, T., Koyama, M.: cgans with projection discriminator. In: *International Conference on Learning Representations* (2018)
30. Nalisnick, E., Matsukawa, A., Teh, Y.W., Gorur, D., Lakshminarayanan, B.: Do deep generative models know what they don’t know? In: *International Conference on Learning Representations* (2019)
31. Neal, L., Olson, M., Fern, X., Wong, W.K., Li, F.: Open set learning with counterfactual images. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 613–628 (2018)
32. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning. In: *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011* (2011)
33. Nguyen, A., Yosinski, J., Clune, J.: Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 427–436 (2015)
34. Oza, P., Patel, V.M.: C2ae: Class conditioned auto-encoder for open-set recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2307–2316 (2019)
35. Pidhorskyi, S., Almohsen, R., Doretto, G.: Generative probabilistic novelty detection with adversarial autoencoders. *Advances in neural information processing systems* **31** (2018)

36. Ren, J., Liu, P.J., Fertig, E., Snoek, J., Poplin, R., Deprieto, M., Dillon, J., Lakshminarayanan, B.: Likelihood ratios for out-of-distribution detection. *Advances in Neural Information Processing Systems* **32** (2019)
37. Sara, U., Akter, M., Uddin, M.S.: Image quality assessment through fsim, ssim, mse and psnr—a comparative study. *Journal of Computer and Communications* **7**(3), 8–18 (2019)
38. Schlegl, T., Seeböck, P., Waldstein, S.M., Schmidt-Erfurth, U., Langs, G.: Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In: *International conference on information processing in medical imaging*. pp. 146–157. Springer (2017)
39. Schonfeld, E., Schiele, B., Khoreva, A.: A u-net based discriminator for generative adversarial networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8207–8216 (2020)
40. Sricharan, K., Srivastava, A.: Building robust classifiers through generation of confident out of distribution examples. *arXiv preprint arXiv:1812.00239* (2018)
41. Vernekar, S., Gaurav, A., Abdelzad, V., Denouden, T., Salay, R., Czarnecki, K.: Out-of-distribution detection in classifiers via generation. *arXiv preprint arXiv:1910.04241* (2019)
42. Wang, H., Liu, W., Bocchieri, A., Li, Y.: Can multi-label classification networks know what they don’t know? *Advances in Neural Information Processing Systems* **34** (2021)
43. Wang, Y., Li, B., Che, T., Zhou, K., Liu, Z., Li, D.: Energy-based open-world uncertainty modeling for confidence calibration. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9302–9311 (2021)
44. Yang, J., Wang, H., Feng, L., Yan, X., Zheng, H., Zhang, W., Liu, Z.: Semantically coherent out-of-distribution detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8301–8309 (2021)
45. Yang, J., Zhou, K., Li, Y., Liu, Z.: Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334* (2021)
46. Yang, Y., Gao, R., Li, Y., Lai, Q., Xu, Q.: What you see is not what the network infers: Detecting adversarial examples based on semantic contradiction. In: *Network and Distributed System Security Symposium (NDSS)* (2022)
47. Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., Xiao, J.: Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop (2016)
48. Yu, Q., Aizawa, K.: Unsupervised out-of-distribution detection by maximum classifier discrepancy. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2019)
49. Zaeemzadeh, A., Bisagno, N., Sambugaro, Z., Conci, N., Rahnavard, N., Shah, M.: Out-of-distribution detection using union of 1-dimensional subspaces. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9452–9461 (2021)
50. Zagoruyko, S., Komodakis, N.: Wide residual networks. In: *British Machine Vision Conference 2016*. British Machine Vision Association (2016)
51. Zhang, H., Goodfellow, I., Metaxas, D., Odena, A.: Self-attention generative adversarial networks. In: *International Conference on Machine Learning* (2019)
52. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 586–595 (2018)
53. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2017)