



# Rethinking Out-of-Distribution Detection From a Human-Centric Perspective

Yao Zhu<sup>1</sup> · Yuefeng Chen<sup>2</sup> · Xiaodan Li<sup>2</sup> · Rong Zhang<sup>2</sup> · Hui Xue<sup>2</sup> · Xiang Tian<sup>1,4</sup> · Rongxin Jiang<sup>1,4</sup> · Bolun Zheng<sup>3,4</sup> · Yaowu Chen<sup>1,5</sup>

Received: 30 November 2022 / Accepted: 22 April 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

## Abstract

**Out-Of-Distribution (OOD)** detection has received broad attention over the years, aiming to ensure the reliability and safety of deep neural networks (DNNs) in real-world scenarios by rejecting incorrect predictions. However, we notice a discrepancy between the conventional evaluation vs. the essential purpose of OOD detection. On the one hand, the conventional evaluation exclusively considers risks caused by label-space distribution shifts while ignoring the risks from input-space distribution shifts. On the other hand, the conventional evaluation reward detection methods for not rejecting the misclassified image in the validation dataset. However, the misclassified image can also cause risks and should be rejected. We appeal to rethink OOD detection from a human-centric perspective, that a proper detection method should reject the case that the deep model's prediction mismatches the human expectations and adopt the case that the deep model's prediction meets the human expectations. We propose a human-centric evaluation and conduct extensive experiments on 45 classifiers and 8 test datasets. We find that the simple baseline OOD detection method can achieve comparable and even better performance than the recently proposed methods, which means that the development in OOD detection in the past years may be overestimated. Additionally, our experiments demonstrate that model selection is non-trivial for OOD detection and should be considered as an integral of the proposed method, which differs from the claim in existing works that proposed methods are universal across different models.

**Keywords** Out-of-Distribution detection · AI reliability · Image classification

---

Communicated by Oliver Zendel.

✉ Rongxin Jiang  
tianx@zju.edu.cn

✉ Bolun Zheng  
blzheng@hdu.edu.cn

Yao Zhu  
ee\_zhuy@zju.edu.cn

Yuefeng Chen  
yuefeng.chenyf@alibaba-inc.com

Xiaodan Li  
fiona.lxd@alibaba-inc.com

Rong Zhang  
stone.zhangr@alibaba-inc.com

Hui Xue  
hui.xueh@alibaba-inc.com

Xiang Tian  
rongxinj@zju.edu.cn

## 1 Introduction

A fundamental component of application safety is modeling the expected operational domain, which provides boundaries for when it is sensible to trust the application and when it is not. However, it is challenging to define such an operational

Yaowu Chen  
cyw@mail.bme.zju.edu.cn

<sup>1</sup> Zhejiang University, Hangzhou 310027, Zhejiang, China

<sup>2</sup> The Security Department, Alibaba Group, Hangzhou 310023, Zhejiang, China

<sup>3</sup> Hangzhou Dianzi University, Hangzhou 310018, Zhejiang, China

<sup>4</sup> Zhejiang Provincial Key Laboratory for Network Multimedia Technologies, Hangzhou 310027, Zhejiang, China

<sup>5</sup> Zhejiang University Embedded System Engineering Research Center, Ministry of Education of China, Hangzhou 310027, Zhejiang, China

domain for machine learning programs, especially for visual classifiers based on Deep Neural Networks (DNNs), which could lead to potentially catastrophic consequences in real-world applications.

Out-of-distribution (OOD) detection methods (Hendrycks & Gimpel, 2017; Liu et al., 2020; Yang et al., 2021, 2022) are developed to determine the application scope of deep models in real-world scenarios, working on the same eventual goal of detecting risk cases that deep models can't give reliable predictions on. There have been plenty of works on designing a criterion to construct the application scope of deep models, including deriving the criterion from the features extracted by the model (Lee et al., 2018; Sun et al., 2022), the outputs of the model (Hendrycks & Gimpel, 2017; Liu et al., 2020), the combination of features and outputs (Wang et al., 2022), and the gradient of the model (Huang et al., 2021).

Existing OOD detection methods have achieved satisfactory performance on conventional evaluations (Huang & Li, 2021; Sun et al., 2021; Huang et al., 2021; Sun & Li, 2022), which roughly consider the whole validation dataset of the model as the in-distribution dataset and the dataset that has disjoint labels with respect to the in-distribution dataset as the out-of-distribution dataset. We notice that there exists a discrepancy between the conventional evaluation and the essential goal of OOD detection that enhances the reliability and safety of deep models.

First, the "distribution" in the "in-distribution" in the conventional evaluation is defined over the whole validation dataset or the training dataset (Hendrycks & Gimpel, 2017; Sun et al., 2021) of the deep model, and the conventional evaluations reward OOD detection methods for not detecting the misclassified images in the validation dataset. This does not meet humans' safety needs because the misclassified input can also threaten the security of the model and should be rejected rather than kept. Second, the conventional evaluations exclusively focus on the risk caused by the label-space-shifted inputs (Huang & Li, 2021; Yang et al., 2021) that belong to new categories. However, the input-space shifts (Huang et al., 2021; Hendrycks & Dietterich, 2018; Hendrycks et al., 2021a) where inputs can be corruption-shifted or domain-shifted but retain the label information of the original dataset may also cause risks for classification. It is challenging for conventional evaluation to determine to what extent input space data variation should be regarded as out of distribution and rejected. Roughly rejecting or adopting all the input-space-shifted images is unreasonable because deep models can reliably handle part of the input-space-shifted images, and the risk comes from the misclassified part. Figure 1 shows that the conventional evaluation may not be suitable to determine which cases humans wish to reject or keep.

In this paper, we appeal to rethink OOD detection from a human-centric perspective, which can provide proper solu-

tions for the above two dilemmas of conventional evaluation. First, according to human security needs, examples that the OOD detector regards as "in-distribution" should be reliable examples of the deep model. Thus, the "distribution" in "in-distribution" should be defined over the images on which the deep model's prediction meets the human expectation in the training dataset rather than over the whole training dataset or validation dataset like conventional evaluations. Second, we extend the evaluation of OOD detection to consider both label-space and input-space shifts. Specifically, we propose the OOD detection methods to reject the input on which the model's prediction mismatches human expectations and keep the input on which the model's prediction meets human expectations.

In practice, we use the ground-truth labels annotated by humans of the dataset as alternatives to human expectations. As shown in Fig. 1, from the human-centric perspective, the OOD detector should keep making predictions on images that can be correctly classified by the model even though the input space of the image is shifted but should reject images that models can not make correct predictions on. The human-centric evaluation method is more in line with the essential purpose of OOD detection and the safety requirements of humans.

We have conducted extensive empirical evaluations of OOD detection methods with various model architectures, including MLP, CNN, and the recent transformer-based model ViT. Our results lead to two major conclusions. First, the baseline method maximum softmax probability (MSP (Hendrycks & Gimpel, 2017)) achieves comparable and even better performance on different distribution-shifted cases when compared to recent popular OOD detection alternatives. This finding suggests that the efforts over the years may be overestimated and not have intrinsically addressed human safety concerns, and demonstrates that the need for rethinking in this field may outweigh the need to propose new detection methods. Second, model architecture and training regime selection matter in OOD detection. Some OOD detection methods almost fail on a certain model but may achieve the best performance on another model. This finding reflects that ignoring the selection of models and claiming the universality across different models, like the existing OOD detection methods, may not be rigorous.

The main contributions of our paper are summarized as follows:

- We provide a novel OOD detection evaluation method from a human-centric perspective that takes whether the model's prediction meets human expectations into account.
- We found that current progress in OOD detection might have been overestimated. Most advanced methods tailor-made for OOD detection are still on par with or even

	Validation Images		Input-Space-Shifted ( retain the original label information )						Label-Space-Shifted (new classes)		
Correctness											
Conventional Evaluation	Keep	Keep	Ignore	Ignore	Ignore	Ignore	Ignore	Ignore	Reject	Reject	Reject
Our Evaluation	Reject	Keep	Keep	Keep	Reject	Keep	Keep	Reject	Reject	Reject	Reject
Correctness											
Conventional Evaluation	Keep	Keep	Ignore	Ignore	Ignore	Ignore	Ignore	Ignore	Reject	Reject	Reject
Our Evaluation	Reject	Keep	Keep	Keep	Reject	Keep	Keep	Reject	Reject	Reject	Reject

**Fig. 1** Images in the blue area are validation images of the ResNet-50 pretrained on ImageNet (Deng et al., 2009). Images in the yellow area are input-space-shifted images that retain the label information of the original dataset (including corruptions and stylized versions). Images in the green area are label-space-shifted images that belong to new categories of the model. “Correctness” means whether the model’s prediction matches the human expectation (the ground truth

label). Conventional evaluations roughly regard all validation images as in-distribution images and exclusively consider the risk caused by label-space shifts while ignoring the risk caused by input-space shifts. Our human-centric evaluation proposes to reject the risk from both the input space shift and the label space shift and keep making predictions on the images that can be correctly classified by the model

worse than the baseline method maximum softmax probability (MSP) under the human-centric evaluation.

- We found that the training regime and model architecture selection matter in OOD detection, which should be regarded as an integral part of the proposed method rather than claiming the proposed method is universal across different models.

The rest of the paper is organized as follows. Section II summarizes the literature related to OOD detection methods and evaluation methods. In Section III, we first provide the formulation of the OOD detection task. Then, we introduce the conventional evaluation method and analyze its drawbacks. After that, we appeal for attention to the discrepancy between the existing evaluation method and the essential goal of the OOD detection task. Moreover, we propose to take the model’s performance into account and rethink the OOD detection task from a human-centric perspective. In Section IV, we first conduct OOD detection experiments with existing methods on different model architectures, which shows that the simple baseline method maximum softmax probability (MSP) beats the other recent methods from the literature. Then, we investigate the influence of model capacity, training strategy, and model calibration on detection performance. Section V gives some conclusive results.

## 2 Related Work

It is well known that a deep neural network can yield unreliable predictions on anomalous inputs from an unknown distribution. When adopted in safety-critical systems such as medical diagnosis or autonomous driving, it is essential to detect examples that the classifier likely fails to make proper predictions. The task of OOD detection is to distinguish the out-of-distribution input (on which the prediction is unreliable) from the in-distribution (ID) input (on which the prediction is reliable). OOD detection has received wide attention because it is critical to ensuring the safety of deep neural networks.

### 2.1 OOD Detection Methods

Directly estimating the density of the examples can be a natural approach to quantify the uncertainty of inputs, which explicitly models the distribution of ID examples with probabilistic models and distinguishes the OOD vs. ID examples through the likelihood (Nalisnick et al., 2018; Kobyzev et al., 2020; Zisselman et al., 2020; Serrà et al., 2019; Ren et al., 2019; Xiao et al., 2020; Kirichenko et al., 2020). However, these methods require costly training of the probabilistic

models, and performance lags behind the classification-based approaches (Huang & Li, 2021).

The classification-based approaches, including confidence enhancement methods and post-hoc methods, use the classifier or a part of the classifier (e.g., feature extractor) to construct a criterion for OOD detection.

Confidence enhancement methods propose retraining the classifier to enhance the sensitivity to OOD examples. Taking advantage of the adversarial training optimization technique, Hein et al. (2019) endow low confidence to the examples far away from the training data and high confidence to the training data. Moreover, some methods introduce a set of collected OOD examples into the training process to enhance the uncertainty estimation and enforce low confidence around the OOD examples (Hendrycks et al., 2018; Papadopoulos et al., 2021; Chen et al., 2021). The correlations between the collected and real OOD examples have a great influence on the performance of these outliers exposure methods (Wu et al., 2021).

Post-hoc methods focus on improving the OOD detection performance with the pre-trained classifiers rather than retraining a model, which is beneficial for adopting OOD detection in real-world scenarios and large-scale settings. These methods derive OOD criterion from different spaces of the deep classifier, e.g., output space (Hendrycks & Gimpel, 2017; Liu et al., 2020; Liang et al., 2018; Sun et al., 2021; Zhu et al., 2022), feature space (Lee et al., 2018; Sun et al., 2022), and gradient space (Huang et al., 2021). The methods based on the output space start from a simple baseline MSP (Hendrycks & Gimpel, 2017), which hypothesizes that the classifier outputs a higher maximum softmax probability on the ID example than the OOD example. ODIN (Liang et al., 2018) introduces a large sufficiently temperature factor and adversarial perturbation to amplify the difference between the softmax probability between the ID and OOD examples. Liu et al. (2020) analyze the limitations of softmax probability in OOD detection and propose to use the negative energy as a criterion. The negative energy is termed as the energy score, and the examples with low energy scores are regarded as OOD examples, and vice versa. Sun et al. (2021) and Zhu et al. (2022) propose to rectify the features of the classifier to improve the detection performance. Virtual-logit Matching (ViM) (Wang et al., 2022) proposes a softmax score which is jointly determined by the feature and the existing logits. The methods based on the feature space suppose the features of OOD examples should be relatively far away from that of in-distribution classes. Lee et al. (2018) model the distribution of feature representations with a mixture of Gaussians and propose using the feature-level Mahalanobis distance as an OOD criterion. GradNorm (Huang et al., 2021) investigates the gradient space of the classifier and shows that the gradients of the categorical cross-entropy loss can be used as an uncertainty measurement. KNN-OOD (Sun et al., 2022)

calculates the distance to the  $k$ -th nearest neighbor between the features of the test input and the features of the model's training dataset, demonstrating the effectiveness of this non-parametric detection metric. Galesso et al. (2022) investigate the efficacy of non-parametric nearest-neighbor distance for complex domains, such as semantic segmentation, and find that model selection has quite a substantial impact on the detection algorithm's performance.

## 2.2 OOD Evaluation Datasets

OOD examples are generally one of two types: i) **label-space-shifted** examples (Van et al., 2018; Xiao et al., 2010; Cimpoi et al., 2014; Zhou et al., 2017; Hendrycks et al., 2021b) and ii) **input-space-shifted** examples (Hendrycks & Dietterich, 2018; Hendrycks et al., 2021a; Geirhos et al., 2019). The label-space-shifted examples belong to a new category that is different from the training dataset, therefore, these examples would not be correctly predicted by deep models. The label-space-shifted examples are conventionally used in OOD evaluation. iNaturalist (Van et al., 2018) is a natural fine-grained dataset that contains images whose labels are disjoint from ImageNet-1k. The Scene UNderstanding (SUN) (Xiao et al., 2010) dataset and Places (Zhou et al., 2017) dataset are scene recognition datasets that can be used as label-space-shifted datasets against ImageNet. The Describable Textures Dataset (Cimpoi et al., 2014) is a texture dataset, which can be divided into 47 categories according to human perception. ImageNet-O (Hendrycks et al., 2021b) contains anomalies of unforeseen classes, which should result in low-confidence predictions and enable us to evaluate the out-of-distribution detection method when the label distribution shifts. The shifts in input space (Hendrycks & Dietterich, 2018; Hendrycks et al., 2021a, b), where images can be corruption-shifted and domain-shifted while remaining in the same label space, are commonly used to evaluate model robustness and domain generalization performance. ImageNet-C (Hendrycks & Dietterich, 2018) applies different corruptions to the ImageNet validation dataset, including noise, blur, compression, etc. ImageNet-R (Hendrycks et al., 2021a) contains various renditions of ImageNet classes, including art, cartoons, graffiti, sketches, etc. The conventional OOD detection evaluation generally focuses on the risk caused by the label-space shifts while ignoring the risk from the input-space shifts.

## 3 OOD Evaluation

### 3.1 Task Formulation

Given a deep classifier to solve the classification problem with  $K$  classes whose labels are denoted as  $\mathcal{Y} =$

$\{1, 2, \dots, K\}$ . Let  $\mathcal{X}$  be the input space. Let  $f : \mathcal{X} \mapsto \mathbb{R}^K$  represent the pre-trained classifier. Suppose that the deep classifier can give reliable predictions on the in-distribution data and the distribution of the in-distribution data  $\mathcal{D}_{in}$  is denoted as  $P_0$ .

OOD detection methods aim to improve the reliability of safety-critical deep models by filtering out the samples for which the model cannot make proper predictions. Thus, OOD detection tasks commonly determine a reject region  $\mathcal{R}$  based on a criterion and a threshold. For any test input  $\mathbf{x} \in \mathcal{X}$ , the classifier rejects to give a prediction on the test input if  $\mathbf{x} \in \mathcal{R}$ . One can use the classifier  $f$  or a part of  $f$  (e.g., feature extractor) to construct a criterion  $T(\mathbf{x}; f)$ , where  $\mathbf{x}$  is the test input. Then, the reject region can be written as  $\mathcal{R} = \{\mathbf{x} : T(\mathbf{x}; f) \leq \gamma\}$ , where  $\gamma$  is the threshold.

The OOD detection criterion is required to give higher scores for the in-distribution examples that classifiers can make reliable predictions and give lower scores for the out-of-distribution examples that classifiers can not handle. When considering the out-of-distribution detection in the classification task to ensure the safety of deep neural networks, the output of the classifier can be formulated as follows:

$$(f, T)(\mathbf{x}; \gamma) := \begin{cases} \text{Rejection,} & \text{if } T(\mathbf{x}; f) \leq \gamma \\ f(\mathbf{x}), & \text{otherwise.} \end{cases} \quad (1)$$

The alarm is triggered when  $T(\mathbf{x}; f)$  falls below the threshold  $\gamma$ .

### 3.2 Conventional OOD Evaluation

Recently proposed OOD detection methods are commonly assessed on conventional OOD evaluation, which uses the FPR95 and AUROC to measure the detection performance. The whole validation dataset of the deep model is typically considered the in-distribution dataset. The label space of the out-of-distribution dataset is disjointed from the in-distribution dataset.

**FPR95:** the false positive rate of OOD (negative) examples when the true positive rate of in-distribution (positive) examples is as high as 95%. Given a threshold  $\gamma$ , the confusion matrix for OOD detection can be expressed as follows:

$$TP(\gamma) = \sum_{i=1}^n (1 - y_{ood}) \cdot \mathbb{I}(T(\mathbf{x}_i; f) > \gamma), \quad (2)$$

$$FN(\gamma) = \sum_{i=1}^n (1 - y_{ood}) \cdot \mathbb{I}(T(\mathbf{x}_i; f) \leq \gamma), \quad (3)$$

$$FP(\gamma) = \sum_{i=1}^n (y_{ood}) \cdot \mathbb{I}(T(\mathbf{x}_i; f) > \gamma), \quad (4)$$

$$TN(\gamma) = \sum_{i=1}^n (y_{ood}) \cdot \mathbb{I}(T(\mathbf{x}_i; f) \leq \gamma), \quad (5)$$

where  $y_{ood}$  is the OOD label which is 1 for the OOD example and 0 for the ID example. In the conventional OOD evaluation, the OOD example is of the unseen category that the model hasn't been trained on.  $\mathbb{I}(Event)$  represent the indicator function:

$$\mathbb{I}(Event) := \begin{cases} 1, & \text{if the Event is True} \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

The true positive rate (TPR) can be computed as:

$$TPR(\gamma) = \frac{TP(\gamma)}{TP(\gamma) + FN(\gamma)}, \quad (7)$$

The false positive rate (FPR) can be computed as:

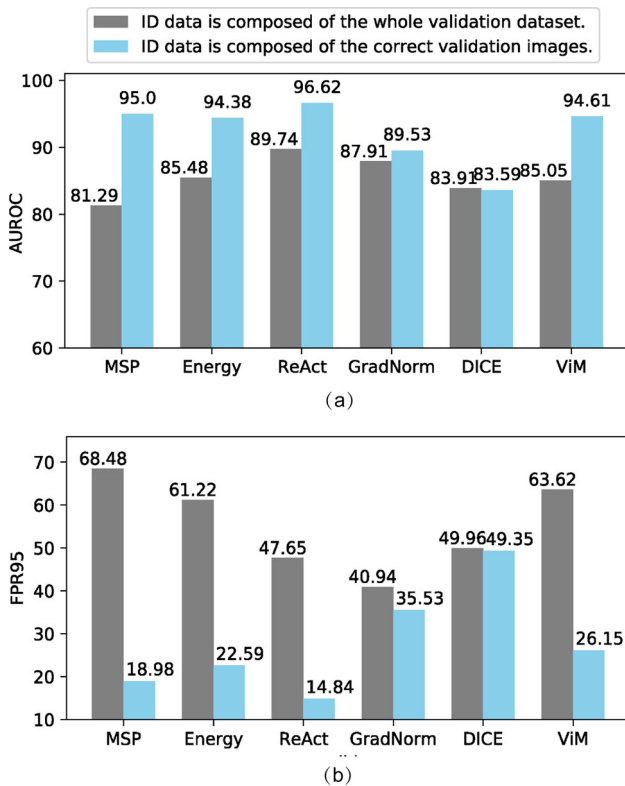
$$FPR(\gamma) = \frac{FP(\gamma)}{FP(\gamma) + TN(\gamma)}. \quad (8)$$

**AUROC:** the area under the receiver operating characteristic curve (ROC), which is the plot of TPR vs. FPR.

### 3.3 Drawbacks of Conventional OOD Evaluation

Existing evaluations roughly regard all the validation images that share the same label space with the training dataset as the in-distribution dataset and regard images of new classes as out-of-distribution images. However, the essential goal of the OOD task is to improve the model's safety by detecting examples in which models can not give reliable predictions. Thus, there exists a discrepancy between the OOD task and the conventional OOD evaluation.

First, regarding all the validation images (including misclassified and correctly classified examples) as the in-distribution dataset is contrary to the human security needs that in-distribution images should be reliably classified. According to the goal of OOD detection, the detection method should assign lower scores to the unrecognizable examples and reject to make predictions on these examples. However, existing evaluation penalizes the OOD detection method for assigning lower scores for misclassified examples in the validation dataset and rewards the OOD detection method for assigning higher scores for these misclassified examples. The discrepancy between the commonly stated purpose of OOD detection and the evaluation method may lead to the underestimation of some OOD detection methods. As shown in Fig. 2, the range of in-distribution data has a significant impact on OOD detection performance. When considering all the validation images of ImageNet (Deng et al., 2009) as in-distribution images, the baseline method MSP



**Fig. 2** **a** The AUROC for different methods (Higher is better) on ResNet-50. **b** The FPR95 for different methods (lower is better) on ResNet-50. The Resnet-50 is pre-trained on the ImageNet dataset. The performance is averaged on four commonly used OOD datasets (iNaturalist, Places, SUN, and Textures). The gray bar indicates that we regard all Imagenet validation examples as ID examples. The blue bar indicates that we regard only the correctly classified examples in the Imagenet validation as ID examples

(Hendrycks & Gimpel, 2017) lags behind other methods. However, when eliminating the influence of the misclassified images, MSP shows comparable performance with other recent methods. Specifically, the performance of the MSP in FPR95 is 27.54% worse than that of GradNorm when regarding all validation images as in-distribution images but surpasses GradNorm by 16.55% when only regarding the correctly classified images as in-distribution images.

Second, the conventional OOD evaluation method roughly considers the risk caused by label-space-shifted examples whose label does not belong to the categories of the training dataset and ignores the risk caused by input-space-shifted cases. It is unreasonable to regard input-space-shifted examples as in-distribution examples because these examples may lead to failure of the classification and cause potential risks. It is also not advisable to arbitrarily treat all the input-space-shifted examples as OOD examples. Consider the following scenario: the input-space-shifted data is constructed by adding small Gaussian noises. Deep models can

provide reliable prediction on this input-space-shifted data, and the OOD detector should not reject this data.

### 3.4 Human-Centric OOD Evaluation

Aware of drawbacks of conventional OOD detection evaluation, we propose to rethink the OOD detection methods from a human-centric perspective. That is, the “distribution” in “in-distribution” and “out-of-distribution” should indicate the distribution of the images on which the model’s prediction meets the humans’ expectations in the training dataset rather than the distribution of the whole training dataset. We denote this distribution as  $P_0$ . “In-distribution” means the test example is drawn from the distribution  $P_0$  where the deep models can provide reliable predictions on. Test examples that deviate away from the distribution  $P_0$  should be rejected. From a human-centric perspective, the examples on which the model’s prediction mismatches the human’s expectation should be rejected, and examples on which the model’s prediction meets the human’s expectation should be kept.

The ground-truth labels of images in the test datasets are commonly annotated by humans. In practice, we use the ground truth label of the image as an alternative to human expectation. The confusion matrix can be reformulated as follows with a given threshold  $\gamma$ :

$$TP_{cor}(\gamma) = \sum_{i=1}^n (y_{cor}) \cdot \mathbb{I}(T(\mathbf{x}_i; f) > \gamma), \quad (9)$$

$$FN_{cor}(\gamma) = \sum_{i=1}^n (y_{cor}) \cdot \mathbb{I}(T(\mathbf{x}_i; f) \leq \gamma), \quad (10)$$

$$FP_{cor}(\gamma) = \sum_{i=1}^n (1 - y_{cor}) \cdot \mathbb{I}(T(\mathbf{x}_i; f) > \gamma), \quad (11)$$

$$TN_{cor}(\gamma) = \sum_{i=1}^n (1 - y_{cor}) \cdot \mathbb{I}(T(\mathbf{x}_i; f) \leq \gamma), \quad (12)$$

where  $y_{cor}$  indicates whether the image can be correctly classified.  $y_{cor}$  is 1 for the example on which the model’s prediction matches the ground-truth label.

We propose to measure the error rate at different thresholds. The threshold is estimated on the correctly classified images in the training dataset (in-distribution images). The error comes from two aspects: one is the assignment of recognizable examples to the reject region ( $FN_{cor}$ ), and the other is the failure to reject unrecognizable examples ( $FP_{cor}$ ).

The detection error rate (DER) can be formulated as:

$$DER(\gamma) = \frac{FN_{cor}(\gamma) + FP_{cor}(\gamma)}{TP_{cor}(\gamma) + FN_{cor}(\gamma) + FP_{cor}(\gamma) + TN_{cor}(\gamma)}. \quad (13)$$

**Algorithm 1** Human-Centric Evaluation.

Require: Classifier  $f$ , in-distribution training dataset  $\mathcal{D}_{in} = \{\mathbf{x}_i, y_i\}_{i=1}^N$ , a test dataset  $\mathcal{D}_{test}$ , an OOD criterion  $T(\mathbf{x}; f)$  and percentage  $p$ .

► **Estimate the reject region :**

ID\_Scores = [ ]

**for**  $i = 1, 2, \dots, N$  **do**

score =  $T(\mathbf{x}_i; f)$

Add score to ID\_Scores

**end for**

$\gamma = \text{quantile}(\text{ID\_Scores}, p)$

► **Measure the DER on  $\mathcal{D}_{test}$  :**

$$DER(\gamma) = \frac{FN_{cor}(\gamma) + FP_{cor}(\gamma)}{TP_{cor}(\gamma) + FN_{cor}(\gamma) + FP_{cor}(\gamma) + TN_{cor}(\gamma)}$$

Lower DER indicates better OOD detection performance.

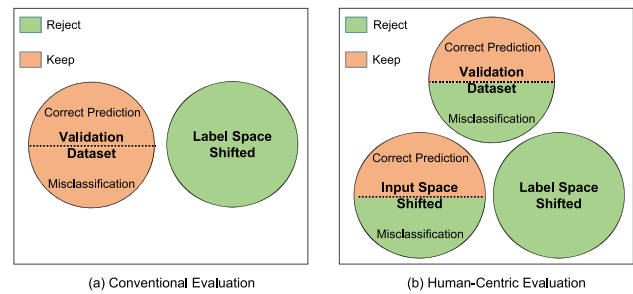
Our evaluation method is outlined in Alg. 1. It is worth noting that our human-centric OOD evaluation method differs from the conventional evaluation method in two aspects. First, the data used to estimate the reject region is different. OOD detection is a single-sample testing task, and the reject region should be determined by known in-distribution examples. We regard correctly classified examples in the training dataset as in-distribution examples and choose the threshold according to the scores of these examples. The conventional evaluation estimates the rejection threshold based on the model's validation dataset in which the misclassified images may lead to the underestimation of the reject region.

Second, the criterion for judging whether the detection result is correct is different. The conventional evaluation requires the OOD detection methods to assign lower scores for the images with unseen categories of the model, while our human-centric evaluation requires OOD detection methods to assign lower scores for the images that the model can not provide a reliable prediction. Figure 3 illustrates the difference between our human-centric evaluation and the conventional evaluation. Our human-centric OOD evaluation broadens the scope of evaluating the OOD detection performance from exclusively detecting risks caused by label-space-shifted images to detecting risks caused by multi-type distribution-shifted cases.

## 4 Experiments

### 4.1 Implementation

We evaluate different OOD detection methods with human-centric evaluation in Alg. 1. We name the DER (in Eq. (13)) whose threshold  $\gamma$  is chosen at the 1st and 5th percentile of scores of the correctly classified images in the training dataset (in-distribution data) of the deep model as DER99 and DER95. The threshold of DER95 means that 95% of in-distribution examples' scores are higher than this threshold.



**Fig. 3** **a** The conventional evaluation aims to evaluate the performance of the OOD detection method on distinguishing the label-space-shifted images from the validation dataset (which is regarded as an in-distribution dataset). **b** Our human-centric evaluation takes the model's performance into account and aims to evaluate the performance of the OOD detection method in rejecting the misclassified images and keeping the correctly predicted images

An OOD detection method should assign higher scores than the threshold to examples for which the model can make reliable predictions and lower scores than the threshold to examples that the model can not properly handle. All experiments in this paper are run on Tesla V100.

**Algorithms.** We evaluate nine OOD detection methods in our experiments:

- **MSP** score (Hendrycks & Gimpel, 2017) uses the maximum softmax probabilities as the criterion score and supposes deep models assign higher probabilities to in-distribution examples.
- **Mahalanobis** score (Lee et al., 2018) computes the minimum Mahalanobis distance between the feature of the test example and the class-wise centroids.
- **KL-Matching** (Hendrycks et al., 2022) uses the minimum KL-divergence between the softmax and the mean class-conditional distributions as the OOD indicator.
- **Energy** score (Liu et al., 2020) proposes an energy function that maps the logit outputs to a scalar through a convenient logsumexp operator. Examples with higher energy are considered OOD examples.
- **ReAct** (Sun et al., 2021) proposes to rectify the activation before calculating the Energy score.
- **GradNorm** (Huang et al., 2021) is an OOD detection method utilizing information extracted from the gradient space.
- **KNN** (Sun et al., 2022) demonstrates the efficacy of the non-parametric nearest-neighbor distance for OOD detection.
- **ViM** (Wang et al., 2022) considers both the class-agnostic score from feature space and the In-Distribution (ID) class-dependent logits in OOD detection.
- **DICE** (Sun & Li, 2022) ranks weights of the model based on a measure of contribution and selectively uses the most salient weights to derive the output for OOD detection.

Table 1 OOD detection performance (DER99 || DER95) on different models that were pre-trained on the ImageNet training dataset

	Method	Textures	iNaturalist	Places	SUN	ImageNet-Val	ImageNet-C	ImageNetV2	ImageNet-R	Average
VGG19	MSP	39.50  22.22	20.14  9.50	<b>46.52  24.29</b>	43.35  22.55	<b>19.99  20.72</b>	<b>25.77  23.34</b>	<b>25.70  23.98</b>	<b>33.70  28.10</b>	<b>31.83  21.84</b>
	Maha	98.53  95.25	99.99  99.91	99.97  99.87	99.98  99.87	27.88  29.99	50.66  50.96	40.48  41.40	66.43  66.47	72.99  72.97
	Energy	34.10  16.37	20.50  5.87	51.88  24.76	41.72  17.49	24.09  23.54	32.23  27.91	32.59  28.20	35.41   <b>26.20</b>	34.07  21.29
	KLM	41.83  20.16	23.07  8.87	57.11  30.10	55.72  27.81	21.77  21.28	28.24  24.03	28.29  24.53	36.08  27.82	36.51  23.07
	ReAct	<b>32.29  15.05</b>	<b>20.04  5.79</b>	51.74  24.98	<b>41.61  18.04</b>	24.07  23.48	32.12  27.71	32.58  28.17	35.15  26.29	33.70   <b>21.19</b>
	Grad	51.28  28.94	50.83  27.96	78.29  56.77	65.50  42.77	26.72  27.04	48.37  45.95	38.52  36.55	60.94  52.69	52.56  39.83
	KNN	98.56  95.35	99.97  99.60	99.95  99.45	99.98  99.76	27.78  29.88	50.56  49.46	40.35  41.22	66.12  64.18	72.95  72.66
	ViM	81.31  53.88	98.86  90.08	98.13  87.62	98.99  91.66	26.44  26.02	46.49  36.52	38.63  33.92	57.89  41.21	68.34  57.61
	DICE	74.10  44.17	85.83  56.15	91.63  72.74	84.65  59.68	27.26  28.28	50.35  49.97	39.78  39.39	65.88  63.43	64.93  51.73
	RN50	MSP	41.45  24.59	23.70  10.49	<b>45.11  24.01</b>	42.43  22.03	<b>18.78  20.42</b>	<b>25.34  23.50</b>	<b>24.35  23.27</b>	<b>31.58  28.79</b>
Maha	66.86  36.15	99.11  93.76	99.66  96.79	99.69  96.72	25.85  27.86	51.14  45.77	39.08  39.96	63.46  60.11	68.11  62.14	
Energy	48.39  24.06	35.83  11.33	59.32  31.67	52.86  25.72	22.79  23.18	31.74  26.85	31.61  27.42	34.27  29.45	39.60  24.96	
KLM	40.64  20.48	23.32  8.12	53.52  25.62	50.58  22.84	19.98  21.00	26.52  23.84	26.59  24.01	34.32  28.56	34.43  21.81	
ReAct	43.24  17.71	<b>20.14  5.50</b>	49.73   <b>22.19</b>	<b>41.80  16.71</b>	22.93  22.93	33.05  26.55	31.97  27.20	35.22  28.55	34.76   <b>20.92</b>	
Grad	51.95  34.31	39.25  20.41	72.07  51.04	59.09  38.47	25.55  26.87	48.81  45.26	37.92  37.55	52.53  44.75	48.40  37.33	
KNN	89.61  71.15	99.53  96.19	99.16  94.36	99.52  95.86	25.76  27.27	48.11  38.73	38.67  38.58	62.36  56.28	69.22  63.47	
ViM	<b>14.29  5.59</b>	71.16  25.73	81.62  49.37	81.69  45.79	24.23  23.50	28.02  26.82	35.32  29.37	36.91   <b>27.58</b>	46.66  29.22	
DICE	63.90  45.82	59.80  37.18	83.52  65.86	71.47  52.23	25.79  27.55	51.41  51.20	38.93  39.76	59.60  55.57	56.80  46.90	
NASNetLarge	MSP	46.99  32.16	35.62  22.81	48.43  31.58	45.14  29.51	<b>15.44  19.30</b>	22.12  23.87	21.53  23.42	35.34  33.20	33.83  26.98
	Maha	<b>22.85  11.86</b>	<b>10.18  4.11</b>	42.12  22.79	38.43   <b>18.79</b>	17.05  21.82	21.87  26.92	22.47  25.06	25.15  28.52	<b>25.02  19.98</b>
	Energy	87.41  72.80	92.00  78.44	94.51  82.66	92.80  77.50	18.84  21.03	32.29  31.01	29.47  30.00	51.41  47.14	62.34  55.07
	KLM	32.77  21.10	21.46  13.20	<b>38.21  22.81</b>	<b>35.17  20.38</b>	15.81  19.82	<b>21.12  23.37</b>	<b>20.80  22.88</b>	30.70  28.81	27.00  21.55
	ReAct	88.60  64.38	89.99  61.33	93.29  76.38	88.05  66.25	18.77  20.89	33.63  30.70	30.07  29.46	53.70  45.72	62.01  49.39
	Grad	98.69  96.65	99.80  99.37	99.95  99.40	99.55  98.31	18.79  20.71	34.38  35.32	30.06  31.25	54.79  54.81	67.00  66.98
	KNN	32.32  16.88	16.78  6.80	48.12  26.54	44.73  22.46	17.33  21.93	21.95  26.04	23.10  25.24	<b>24.10  26.50</b>	26.78  20.77
	ViM	29.27  14.29	14.33  6.15	50.06  29.48	45.98  24.83	17.03  21.02	21.92  25.29	22.94  24.43	24.73  26.52	28.28  21.50
	DICE	99.40  97.62	99.76  99.51	99.90  99.17	99.38  98.30	18.66  20.66	34.31  35.34	30.01  31.10	54.80  54.82	67.04  67.07



Table 1 continued

	Method	Textures	iNaturalist	Places	SUN	ImageNet-Val	ImageNet-C	ImageNetV2	ImageNet-R	Average
Inception-v4	MSP	41.03  25.98	28.25  15.60	42.99   <b>24.33</b>	39.12   <b>21.59</b>	<b>16.64</b>    <b>20.35</b>	22.29  24.20	21.67  23.27	31.31  30.16	30.41  23.18
	Maha	34.54  14.45	27.84   <b>6.97</b>	73.86  41.41	71.97  37.29	20.03  21.63	27.63  25.93	28.95  26.70	37.34  27.36	40.27  25.22
	Energy	63.81  46.24	63.22  41.98	72.05  51.59	68.62  47.79	19.81  21.91	30.45  28.53	28.12  27.58	40.54  35.79	48.33  37.68
	KLIM	33.62  21.84	20.50  12.38	<b>41.18</b>   24.84	<b>37.68</b>   21.79	16.89  20.36	<b>22.21</b>    <b>23.68</b>	<b>21.61</b>    <b>22.75</b>	<b>29.70</b>   28.53	<b>27.92</b>    <b>22.02</b>
	ReAct	49.41  30.37	37.02  20.24	63.26  42.31	56.42  36.11	19.40  21.38	28.26  26.71	27.02  26.12	36.70  31.75	39.69  29.37
	Grad	96.65  93.14	98.73  96.73	99.70  98.84	99.12  97.19	21.37  23.55	40.59  41.45	32.68  33.94	57.84  57.71	68.33  67.82
	KNN	44.08  16.93	34.38  8.12	67.84  36.67	67.18  32.06	19.65  21.99	26.02  25.79	26.94  27.05	31.74  25.62	40.00  24.75
	ViM	<b>26.13</b>    <b>12.70</b>	<b>22.64</b>   9.39	66.25  44.38	60.21  37.34	19.44  21.21	25.88  25.83	27.48  26.00	29.77   <b>25.21</b>	34.73  25.26
	DICE	99.86  99.11	99.98  99.93	99.98  99.80	99.94  99.66	21.37  23.36	40.62  41.50	32.68  33.82	57.93  57.97	69.04  69.39

A lower detection error rate (DER99 and DER95) indicates better performance. All methods are post hoc and can be directly used for pre-trained models. We consider 8 test datasets, including the validation dataset of ImageNet, 3 input-space-shifted datasets, and 4 label-space-shifted datasets. "Average" refers to the average performance on 8 datasets. Due to the page width limits, in this table, "Mahalanobis score" is abbreviated as "Maha.", "KL Matching" is abbreviated as "KL.M", and "GradNorm" is abbreviated as "Grad.". The best performances are indicated in bold. All entries are percentages

Table 2 Continued for Table 1. OOD detection performance (DER99 || DER95) on different models that were pre-trained on the ImageNet training dataset

Method	Textures	iNaturalist	Places	SUN	ImageNet-Val	ImageNet-C	ImageNetV2	ImageNet-R	Average
SENet154	MSP	35.60  20.64	41.85  23.09	38.58  19.63	<b>15.97  19.98</b>	21.82   <b>24.07</b>	21.27  23.36	32.11  31.53	28.71  21.65
	Maha	14.27  7.07	42.56  19.60	38.64  14.84	17.97  21.32	23.30  27.20	24.36  24.53	25.06  27.88	24.56   <b>18.23</b>
	Energy	45.32  30.71	36.56  19.19	53.81  33.30	18.14  21.68	24.99  26.39	24.93  26.52	33.36  32.54	36.70  28.49
	KLM	28.07  15.34	15.94  7.61	<b>36.35</b>   19.12	16.12  20.26	<b>21.32</b>   24.09	<b>20.88</b>    <b>22.93</b>	29.60  29.16	25.29  19.30
	ReAct	21.79  11.38	12.19  5.73	40.84  23.47	17.81  21.23	22.83  26.02	23.40  24.85	28.07  29.11	25.12  19.84
	Grad	97.38  95.69	99.92  99.74	99.92  99.62	99.85  99.52	37.41  38.40	31.76  33.06	58.10  58.23	68.06  68.30
	KNN	20.64  9.63	12.25  4.23	40.31   <b>18.05</b>	37.74   <b>14.04</b>	22.85  26.53	23.42  24.67	25.02   <b>27.81</b>	24.79  18.27
	ViM	<b>12.57</b>    <b>6.52</b>	10.52  3.67	41.47  20.32	36.15  14.47	23.19  27.47	23.84  24.75	<b>24.86</b>   27.82	<b>23.81</b>   18.29
	DICE	99.80  99.38	99.88  99.78	99.99  99.85	99.99  99.86	37.45  38.29	31.77  33.12	58.09  58.21	68.40  68.86
	MSP	59.43  32.93	<b>43.30</b>   17.21	<b>65.49</b>    <b>35.86</b>	64.09  34.62	<b>33.42</b>    <b>24.20</b>	<b>31.30</b>    <b>24.50</b>	<b>40.41</b>   29.10	<b>45.17</b>    <b>27.36</b>
MobileNetV2	Maha	84.65  64.33	99.95  99.48	99.91  99.13	30.53  32.20	57.38  56.73	43.92  44.78	69.54  67.88	73.23  70.49
	Energy	55.94  29.52	47.03  16.31	67.15  38.88	27.43  25.42	37.46  28.66	37.07  30.67	43.74  28.29	47.14  28.74
	KLM	63.28  30.16	48.93  16.81	76.92  44.11	78.21  43.22	37.41  25.42	35.35  26.48	48.53  30.16	51.88  29.80
	ReAct	53.10  25.57	44.86   <b>14.51</b>	66.19  37.43	59.61  30.40	27.40  25.35	36.99  30.59	43.55   <b>27.77</b>	46.12  27.51
	Grad	45.25  23.95	45.53  21.96	68.04  42.77	<b>53.70</b>    <b>29.25</b>	29.43  28.88	41.66  38.00	57.63  43.20	49.01  33.69
	KNN	92.39  81.95	99.91  99.51	99.90  99.09	99.96  99.60	30.59  32.11	43.94  44.64	69.06  66.40	73.68  71.41
	ViM	<b>33.63</b>    <b>11.49</b>	96.66  78.54	96.97  83.05	97.10  84.13	29.48  27.56	42.00  35.70	55.17  31.19	61.79  47.78
	DICE	52.48  31.12	56.75  34.21	76.37  53.41	60.98  37.98	30.22  30.86	43.33  42.22	64.67  56.26	55.10  42.40
	MSP	34.63  20.05	10.51  4.30	<b>38.76</b>   20.41	35.69  17.54	<b>13.66</b>   19.38	<b>18.93</b>    <b>23.67</b>	31.33  37.43	25.85  23.12
	Maha	35.41  15.21	1.81  0.45	47.17  20.31	44.83  16.02	15.76  21.28	21.86  25.60	31.35  40.38	28.20  22.31
Swin-B	Energy	46.38  30.60	12.69  7.03	51.70  36.10	44.99  29.89	30.17  43.71	23.71  27.26	33.45  35.71	32.35  28.82
	KLM	31.21  17.04	7.99  2.66	39.40   <b>18.01</b>	<b>35.09</b>   15.72	26.63  42.28	19.66  23.71	31.82  38.11	25.78  22.18
	ReAct	38.42  23.63	6.75  3.20	44.40  28.76	37.94  23.03	15.64  20.55	22.84  26.95	31.83  35.70	28.47  25.85
	Grad	96.15  94.17	89.20  86.92	94.07  92.74	91.86  90.15	15.43   <b>18.01</b>	25.33  27.81	39.16  39.20	59.88  60.14
	KNN	46.52  19.98	10.70  1.18	53.58  24.80	53.82  21.12	15.74  21.24	22.01  25.12	<b>29.49</b>   37.94	29.94  22.87
	ViM	<b>25.48</b>    <b>11.45</b>	<b>0.96</b>    <b>0.24</b>	41.16  18.79	36.38   <b>12.98</b>	15.86  21.24	31.69  40.84	32.21  41.42	<b>25.71</b>    <b>21.60</b>
	DICE	44.24  23.30	45.71  24.59	52.01  30.44	44.88  25.60	15.55  19.15	23.85  25.63	35.60   <b>35.58</b>	36.90  28.28

Table 2 continued

Method	Textures	iNaturalist	Places	SUN	ImageNet-Val	ImageNet-C	ImageNetV2	ImageNet-R	Average
ViT-B16	35.41 19.91	8.16 3.23	39.12 20.02	36.12 18.40	<b>12.82</b>  18.17	<b>17.20</b>  23.94	<b>18.33</b>  22.88	31.63 38.82	24.85 20.67
Maha	<b>19.79</b>  9.59	<b>1.18</b>  0.41	43.36 18.91	32.21 11.43	15.13 20.93	21.12 30.98	21.80 25.80	32.42 41.43	23.38 19.94
Energy	21.91 11.65	2.66 0.91	<b>33.49</b>  17.35	<b>25.58</b>  12.54	15.28 20.58	21.86 29.77	21.42 25.25	35.98 44.46	<b>22.27</b>  20.31
KLM	28.67 16.97	5.03 2.08	37.45 19.13	33.74 16.68	13.99 18.84	18.93 24.37	19.76 23.64	34.03 39.69	23.95 20.17
ReAct	22.11 11.90	2.35 0.82	36.56 18.43	28.27 12.74	15.29 20.29	21.14 29.02	21.17 25.09	33.74 42.84	22.58  <b>20.14</b>
Grad	24.82 17.07	3.22 1.80	36.37 23.31	28.47 17.68	14.73 18.19	21.14 27.12	21.31 24.25	34.33 39.87	23.05 21.16
KNN	24.54 12.34	2.05 0.68	42.93 20.02	37.41 14.70	15.71 21.75	22.31 31.81	22.10 27.70	34.30 43.23	24.27 20.73
ViM	21.77 10.60	1.87 0.58	46.04 22.18	33.66 14.28	14.86 19.11	21.00 28.65	21.75 24.31	<b>29.85</b>   <b>38.81</b>	23.85 19.81
DICE	92.82 80.83	99.75 98.53	99.60 97.30	99.12 96.00	14.61 17.33	24.53 32.70	24.19 27.24	39.14 39.30	61.72 61.15

All entries are percentages

The best performances are indicated in bold

**Datasets.** We evaluate the OOD detection methods by detecting the risk caused by both the input-space-shift and label-space-shift cases. We regard the correctly classified images in the training dataset of ImageNet (Deng et al., 2009) as the in-distribution dataset. We consider 8 test datasets, including Textures (Cimpoi et al., 2014) which contains 47 categories of textures, iNaturalist (Van et al., 2018) which contains 5089 natural fine-grained categories, Places (Zhou et al., 2017) which contains more than 400 different types of scene environments, SUN (Xiao et al., 2010) which contains 397 categories of scene environments, ImageNetV2 (Recht et al., 2019) which is independent of existing models, ImageNet-C (Hendrycks & Dietterich, 2018) which contains different corruption versions (weather, noise, brightness, et al.) of ImageNet, ImageNet-R (Hendrycks et al., 2021a) which contains different stylized versions (art, sketch, cartoon, et al.) of ImageNet and ImageNet-Val which represents the validation dataset of ImageNet. The label  $y_{cor}$  (in Eq. (9)) for the label-space-shifted datasets Textures, iNaturalist, Places, and SUN is 0 because the label space of these datasets is disjoint from the training dataset of the deep model and the deep model can not make correct predictions. The label  $y_{cor}$  for the datasets ImageNet-C, ImageNet-R, ImageNetV2, and ImageNet-Val indicates whether the model's prediction meets the ground-truth label.

**Model Architectures.** We extensively conduct experiments on 45 deep models with various model architectures, from Multilayer Perceptron (MLP) based model ResMLP (Touvron et al., 2022), classical CNN models (VGG (Simonyan & Zisserman, 2015), ResNet (He et al., 2016), DenseNet (Huang et al., 2017), MobileNet-V2 (Sandler et al., 2018), NASNet (Zoph et al., 2018), the InceptionV4 (Szegedy et al., 2017) and the squeeze-and-excitation network (Hu et al., 2018)) to the latest transformer models (ViT (Dosovitskiy et al., 2021), Swin (Liu et al., 2021) and DeiT (Touvron et al., 2021)). The pre-trained CNN models are provided by torchvision 0.10.0 (Paszke et al., 2019), and other pre-trained models are provided by timm 0.4.12 (Wightman, 2019).

## 4.2 Results on Different Architectures

Prevalent OOD detection evaluations exclusively focus on the risk caused by a part of distribution shifts (label space shift) and ignore the risk caused by the input space shift. For the first time, we evaluate the existing detection methods on a variety of models and different distribution-shift cases, hoping to draw general methodological conclusions.

Tables 1 and 2 show the detection performance results on different model architectures. For each model architecture, we report the OOD detection performance of different methods on each test dataset and the average performance on eight datasets. We draw two main conclusions from our results:

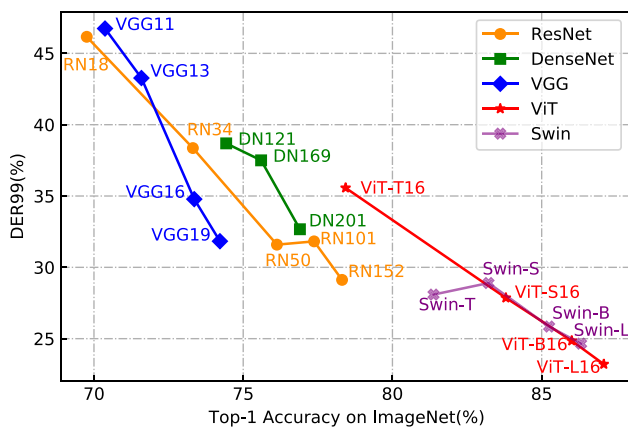
### **From a human-centric perspective, no method can consistently outperform MSP by a significant margin.**

Research on OOD detection originates from a simple baseline MSP (Hendrycks & Gimpel, 2017; Yang et al., 2021) and receives increased attention. However, our experiment results reveal that the development of OOD detection in the past five years may be overestimated. In Table 1, the baseline method MSP achieves state-of-the-art performance in the average DER99 on VGG19, ResNet-50, and MobileNet-V2. When considering VGG19 as the classifier, MSP surpasses the recently proposed ReAct and ViM by 1.87% and 36.5% in DER99, respectively. On other models, although MSP is not the best one, it only lags behind the best method to a small extent. Specifically, when considering Swin-B as the classifier, MSP lags behind the best method ViM by 0.14% but surpasses ReAct and GradNorm by 2.62% and 34.03%, respectively.

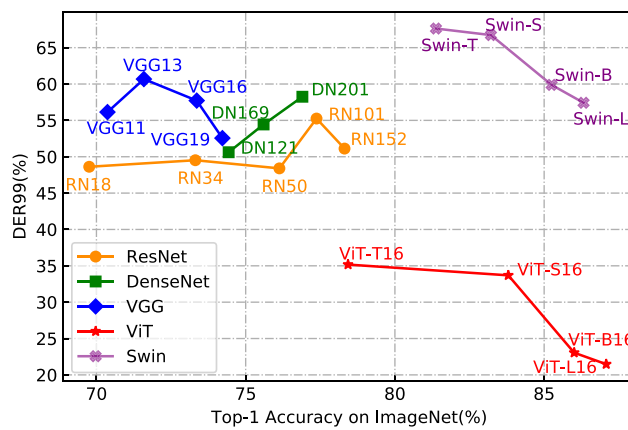
### **Model architecture matters in detection methods.**

Besides ignoring the performance of the model, the existing methods also ignore the influence of the model architecture. The existing methods are always evaluated on several models and then claim to be universal across different models. We conduct experiments on a variety of model architectures and find that the impact of model architecture on detection methods can be fatal. For instance, KNN surpasses MSP by 4.9% with SENet154 in the average DER99 but its performance drops dramatically with ResNet50. That is to say, although KNN performs well with the SENet154, it's unreliable to use KNN as the OOD detection algorithm to improve the security of the ResNet50 model. Similarly, ViM performs poorly with VGG19, and the detection error rate (DER99) is even over 98% when evaluated on the iNaturalist dataset, but it can achieve the best performance in the Swin model. Contrary to prior works that arbitrarily claim their method is universal and performs well across different model architectures, we suggest the proposed OOD detection method takes the model selection strategy into consideration.

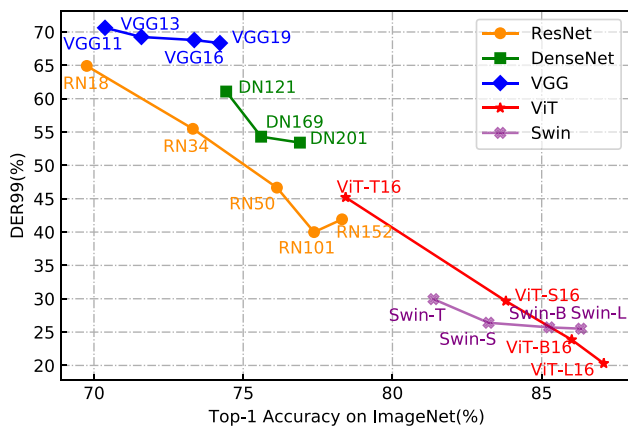
We do not claim that any of these recently proposed methods cannot possibly improve the detection performance, but getting detection performance improvements is challenging and should regard the model architecture as an integral part of the proposed method. The newly proposed OOD detection methods should remain true to their original aspiration, aiming to improve the safety and reliability of the deep models rather than distinguishing the difference between datasets. The future research direction of OOD detection can be to design detection methods that are more in line with human needs and guarantee the safety of the deep models facing different distribution-shifted cases.



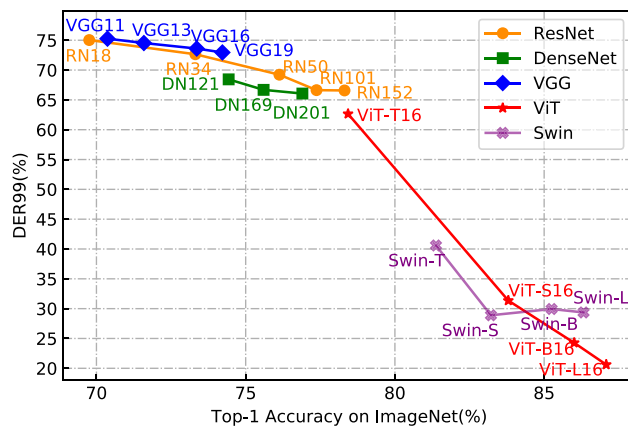
(a) MSP



(b) GradNorm



(c) ViM



(d) KNN

**Fig. 4** The influence of the model’s capacity on OOD detection methods (Lower DER99 indicates better detection performance). We choose VGG, ResNet, DenseNet, ViT, and Swin with different capacities. The

performance is averaged on eight test datasets. The horizontal axis represents the accuracy of different models

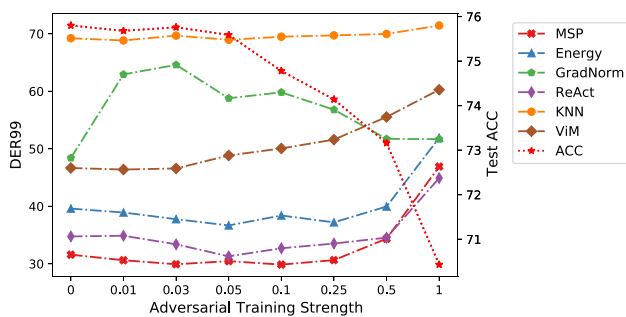
### 4.3 The Influence of Model Capacity

Generally, the model capacity is positively correlated with the classification accuracy. For example, the classification accuracy of models ResNet-18, ResNet-34, ResNet-50, ResNet-101, and ResNet-152 increases. Then, how does the model’s capacity affect the performance of OOD detection methods? In this subsection, we investigate the influence of model capacity on the OOD detection methods.

Figure 4 illustrates the performance of OOD detection methods in models with different capacities. We found an interesting phenomenon that there is no monotonous relationship between the performance of detection methods and model capacity. When using the DenseNet model as the classifier, the performance of GradNorm becomes worse with the increase in model capacity. When using ViM to conduct OOD detection, its performance with the model ResNet101 is 6.67% better than with the smaller model ResNet50 and 1.92% better than with the larger model ResNet152.

Additionally, KNN in the Swin-S(mall), Swin-B(ase), Swin-L(arge) performs better than in the Swin-T(iny), but KNN in the Swin-B(ase) performs worse than in the Swin-S(mall).

On the whole, the accuracy of the model and the OOD detection performance are positively correlated. Models with higher accuracy tend to achieve a lower OOD detection error rate (DER99). The model ViT-L16 which obtains the best Top-1 accuracy, also performs the best in OOD detection. The model ResNet-18, whose Top-1 accuracy is the worst, performs poorly in OOD detection. However, there are also some anomalies. ViT-T16 can achieve better Top-1 accuracy than ResNet-101. MSP and ViM perform better in ResNet-101 than in ViT-T16, while GradNorm and KNN perform worse in ResNet-101 than in ViT-T16.



**Fig. 5** The DER99 for different methods (lower is better) on different robust models. The horizontal axis represents the perturbation strength in training the model, e.g. "0.1" represents the robust model trained with  $\ell_2$  perturbation  $\epsilon = 0.1$ . The red dotted line indicates the test accuracy of the robust model. The performance is averaged on eight test datasets

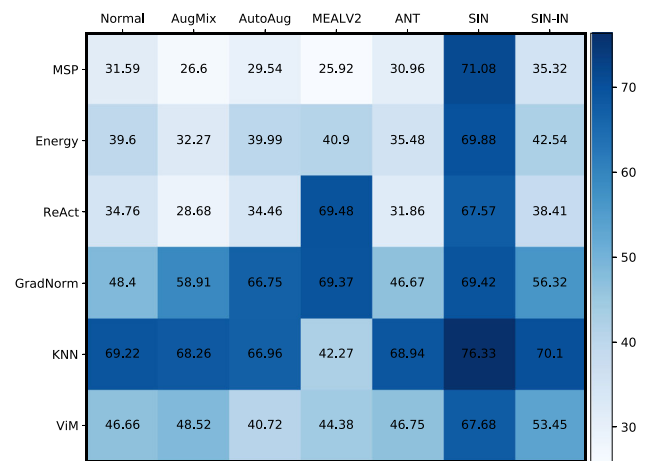
## 4.4 The Influence of Training Regimes

### 4.4.1 Does Adversarial Robustness Contribute to OOD Detection?

Adversarial training (Salman et al., 2020) is efficient in improving the adversarial robustness and the safety of deep neural networks. To find out whether adversarial training can improve the OOD detection performance of the model, we conduct experiments in terms of different perturbation strengths (constraints of  $\ell_2$  norm bound). In Fig. 5, we evaluate the OOD detection performance on eight adversarially pre-trained ResNet-50. When the strength of adversarial training is small, it slightly impacts OOD detection. However, when the strength of adversarial training is high, the OOD detection performance of this model decreases dramatically. Specifically, the DER99 of MSP on the model adversarially trained with strength  $\epsilon = 1$  is 15.32% worse than that on the normal model. MSP almost performs the best on different robust models.

### 4.4.2 Does Training Strategy Contribute to OOD Detection?

Various studies have focused on training strategies to increase the model's performance in classification tasks. To find out the influence of the training strategies on OOD detection, we further consider different strategies, including training with Styled ImageNet (SIN) (Geirhos et al., 2019), training with the mixture of Styled and natural ImageNet (SIN-IN) (Geirhos et al., 2019), training with the augmentation method Augmix (Hendrycks et al., 2020), training with the augmentation method AutoAugment (Lim et al., 2019), training with additive Gaussian and Speckle noise (ANT (Rusak et al., 2020)), and training with the knowledge distillation method MEALV2 (Shen & Savvides, 2020) that achieves 80%+ Top-1 accuracy on ResNet-50. As shown in Fig. 6, the most beneficial training strategies for different OOD detec-



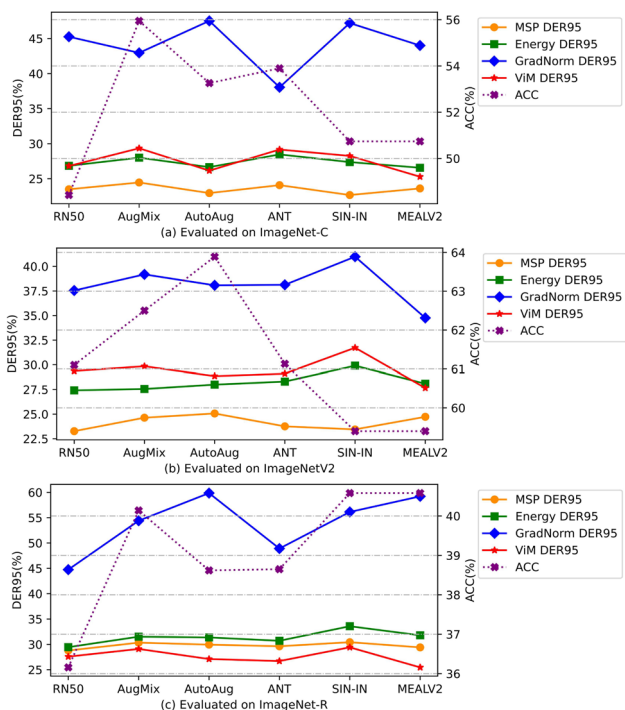
**Fig. 6** The influence of the training strategies on OOD detection. Lighter color indicates better performance (lower DER99). The performance is averaged on eight test datasets

tion methods are different. MSP and KNN perform the best in the model trained with the MEALV2 strategy. GradNorm outperforms ViM in the ANT strategy-trained model, whereas ViM outperforms GradNorm in the AutoAug strategy-trained model.

Furthermore, we have conducted experiments to investigate the relationship between the accuracy of models trained with different strategies on the input-space-shifted test dataset and the OOD detection performance when using these models. The results are illustrated in Fig. 7. When using ImageNet-C as the test dataset, the model trained with the Augmix strategy achieves higher accuracy than the normal ResNet-50. However, the MSP, Energy, and ViM detection algorithms perform worse with the model trained with the AugMix strategy than with the normal ResNet-50. When using ImageNetV2 as the test dataset, models trained with the SIN-IN and MEALV2 training strategies achieve lower accuracy than normal ResNet-50. However, when using these two models with lower accuracy for OOD detection, the performance of the detection algorithms is comparable to or even better than using the normal ResNet-50. To sum up, models trained with different strategies exhibit varying classification abilities on input-space-shifted examples, and the relationship between this classification ability and OOD detection capability is not necessarily monotonous. We think that this is because detection metrics are derived from the model's output space, feature space, or gradient space and do not depend on the model's classification performance.

### 4.4.3 Does Knowledge Distillation Contribute to OOD Detection?

In Fig. 8, we investigate whether knowledge distillation can improve the detection performance of different methods.



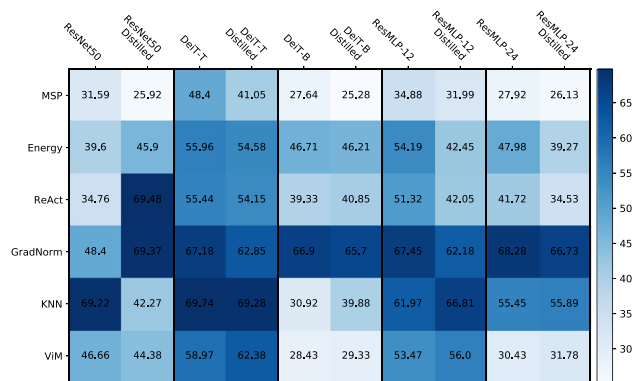
**Fig. 7** The relationship between the model’s classification performance on the input-space-shifted examples and the OOD detection performance

We consider the CNN-based model ResNet-50, MLP-based models ResMLP-12 and ResMLP-24, transformer-based models DeiT-T and DeiT-B, and their distilled versions. We find that MSP’s performance on the distilled model is better than that on the normal model. The influence of the distillation on other methods is inconsistent in different models. For instance, the performance of KNN improves by 26.95% when using the distilled ResNet-50 instead of the normal ResNet-50, while its performance degrades by 8.96% when using the distilled DeiT-B instead of the normal DeiT-B.

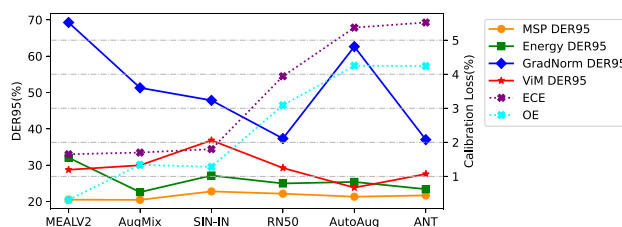
These results shed light on the fact that the model architecture and training regime should also be considered when designing the OOD detection algorithm.

### 4.5 The Influence of Model Calibration

Model calibration (Guo et al., 2017) refers to the accuracy with which the probability output by the deep model reflects its predictive uncertainty, aiming to ensure that the output class probabilities are consistent with what would naturally occur. Model calibration can be measured by Expected Calibration Error (ECE) (Naeini et al., 2015) and Overconfidence Error (OE) (Thulasidasan et al., 2019). The OOD detection task aims to derive a detection metric from the model’s output space (Hendrycks & Gimpel, 2017), features space (Wang et al., 2022), or gradient space (Huang et al., 2021) and set a threshold to determine whether an example is in-distribution



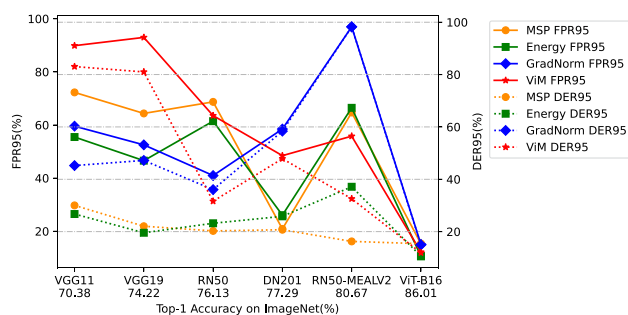
**Fig. 8** The influence of the distillation on OOD detection. Lighter color indicates better performance (lower DER99). The performance is averaged on eight test datasets



**Fig. 9** The influence of the model calibration on OOD detection methods. We use the Expected Calibration Error (ECE) and Overconfidence Error (OE) as calibration metrics

or out-of-distribution. It is essentially a binary classification task (Hendrycks & Gimpel, 2017) and does not require the calibration of the model’s output. The conventional evaluation method evaluates whether the detection algorithm can identify the label-space-shifted examples as OOD examples, while our human-centric evaluation method requires the detection algorithm to identify both label-space-shifted examples and misclassified input-space-shifted examples as OOD examples.

We conduct experiments to investigate how model calibration affects the performance of OOD detection methods. As shown in Fig. 9, models with better calibration may perform better when using metrics derived from the model’s output space for OOD detection, but their performance may not necessarily be better when using other types of OOD detection metrics. Specifically, models trained with the MEALV2 and AugMix strategies exhibit better calibration, with ECE values of 1.65% and 1.69%, respectively. Simultaneously, their detection performance is superior to models obtained through other training strategies when using the MSP detection metric derived from the model’s output space. However, when using the ViM detection metric derived from the model’s features space, the model trained with MEALV2 performs worse than the model trained with AutoAug, even though its ECE is 3.72% lower than the model trained with AutoAug.



**Fig. 10** The influence of model selection on OOD detection performance under conventional evaluation metric FPR95 and our DER95. The performance is averaged on four label-space-shifted test datasets

#### 4.6 The Influence of Model Selection in Conventional Evaluation

In this section, we investigate whether model selection also significantly influences the detection performance of different detection algorithms under conventional evaluation metrics. As shown in Fig. 10, we simultaneously evaluate different detection methods using the conventional evaluation metric FPR95 and our human-centric evaluation metric DER95. Considering that the conventional evaluation metric FPR95 is only applicable to label-space-shifted examples, we conduct experiments on four label-space-shifted datasets: iNaturalist, Places, SUN, and Textures, and illustrate the average results.

From the experimental results, we can obtain the following findings: Firstly, whether using the conventional evaluation metric FPR95 or our proposed DER95, the influence of model accuracy on OOD detection performance does not exhibit a monotonic relationship. In other words, a model with higher accuracy may not perform better in OOD detection tasks. Secondly, when using the conventional evaluation metric FPR95, an OOD detection algorithm that performs well on a particular deep model may not generalize well to another model. For example, when using ResNet-50, the FPR95 of GradNorm is lower than other detection algorithms, indicating better performance than other algorithms. However, when using DenseNet-201, the detection performance of the GradNorm algorithm is inferior to that of other algorithms. Thirdly, when assessed with different evaluation metrics, the model selection has varying effects on the evaluation results of detection algorithms. When evaluated using the conventional evaluation metric FPR95, the ViM detection algorithm performs better with DenseNet-201 compared to ResNet-50. However, when evaluated using our DER95 metric, the ViM detection algorithm with DenseNet-201 is inferior to that with ResNet-50. These findings demonstrate the importance of model selection on the performance of the detection algorithm, which not only holds true when using our proposed

DER as an evaluation metric but also holds true when using conventional evaluation metrics.

## 5 Conclusion

Current deep learning systems yield uncertain and unreliable predictions when facing distribution-shifted examples. Various OOD detection methods have been proposed to improve the reliability and safety of deep models in real-world applications over the years. This work does not propose a novel detection method but analyzes the drawbacks of the conventional evaluation and proposes a human-centric evaluation method that is more in line with the essential goal of OOD detection to investigate how practical existing OOD detection methods are in aiding the safety problem of human needs. We have conducted extensive experiments on nine OOD detection methods with forty-five different models. Our results draw two major conclusions. First, though existing works have made substantial progress under the conventional evaluation, the newly proposed detection algorithms have not shown a clear advantage over the baseline method when considering a human-centric evaluation. Second, model architectures and training regimes matter in OOD detection and should be considered integral when designing new detection methods. We hope our findings motivate researchers to rethink OOD detection from a human-centric perspective and develop OOD detection methods that can reject risks both from the label-space shifts and input-space shifts.

**Author Contributions** All authors contributed to the research conception and design. Methodology: [Yao Zhu], [Yuefeng Chen]; Material preparation: [Bolun Zheng], [Rong Zhang]; Formal analysis and investigation: [Xiang Tian], [Rongxin Jiang]; Writing - original draft preparation: [Yao Zhu]; Writing - review and editing: [Yuefeng Chen], [Xiaodan Li]; Funding acquisition: [Hui Xue], [Yaowu Chen].

**Funding** This work was supported by the National Natural Science Foundation of China (Grant No. U21B2024), the Fundamental Research Funds for the Central Universities, the Fundamental Research Funds for the Provincial Universities of Zhejiang (Grants No. GK239909299001-013), the Key R&D Program of Zhejiang (Grant No. 2023C01044), and Alibaba Group through the Alibaba Research Intern Program.

**Data Availability** The datasets analyzed during the current study are available in the ImageNet repository (Deng et al., 2009), <https://www.image-net.org/about.php>.

**Code Availability** We share some codes in the link below to evaluate the performance of different detection methods with our human-centric evaluation. We plan to open-source more codes for the community in the future. We hope our finds can provide insights for future research on OOD detection. [https://drive.google.com/file/d/1avkZrzMoDjqQPoBL6MkYW\\_javVGrOzln/view?usp=share\\_link](https://drive.google.com/file/d/1avkZrzMoDjqQPoBL6MkYW_javVGrOzln/view?usp=share_link)



[https://drive.google.com/file/d/1avkZrzMoDjqQPoBL6MkYW/\\_javVGrOzln/view?usp=share\\_link](https://drive.google.com/file/d/1avkZrzMoDjqQPoBL6MkYW/_javVGrOzln/view?usp=share_link).

## Declarations

**Conflict of interests** Yao Zhu received a Ph.D. stipend from Zhejiang University. Yuefeng Chen, Xiaodan Li, Rong Zhang, and Hui Xue received salaries from the Alibaba group. Xiang Tian and Yaowu Chen received salaries from Zhejiang University. Bolun Zheng receives a salary from Hangzhou Dianzi University.

**Consent to Participate** Not applicable.

## References

- Chen, J., Li, Y., Wu, X., Liang, Y., & Jha, S. (2021). Atom: Robustifying out-of-distribution detection using outlier mining. *In: Joint European Conference On Machine Learning And Knowledge Discovery In Databases, Springer*, pp 430–445
- Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., & Vedaldi, A. (2014). Describing textures in the wild. *In: Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 3606–3613
- Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. *In: 2009 IEEE conference on computer vision and pattern recognition*, Ieee, pp 248–255
- Dosovitskiy, A., Beyer, L., Kolesnikov, A. et al. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. International Conference on Learning Representations
- Galesso, S., Argus, M., & Brox, T. (2022). Far away in the deep space: Nearest-neighbor-based dense out-of-distribution detection. arXiv preprint [arXiv:2211.06660](https://arxiv.org/abs/2211.06660)
- Geirhos, R., Rubisch, P., Michaelis, C., et al. (2019). Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. *In: International Conference on Learning Representations*, <https://openreview.net/forum?id=Bygh9j09KX>
- Guo, C., Pleiss, G., Sun, Y., et al. (2017). On calibration of modern neural networks. *In: International conference on machine learning*, PMLR, pp 1321–1330
- He, K., Zhang, X., Ren, S. et al. (2016). Deep residual learning for image recognition. *In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp 770–778
- Hein, M., Andriushchenko, M., Bitterwolf, J. (2019). 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
- Hendrycks, D., & Dietterich, T. (2018). Benchmarking neural network robustness to common corruptions and perturbations. *In: International Conference on Learning Representations*, pp 1–13
- Hendrycks, D., Gimpel, K. (2017). A baseline for detecting misclassified and out-of-distribution examples in neural networks. *Proceedings of International Conference on Learning Representations* pp 1–12
- Hendrycks, D., Mazeika, M., & Dietterich, T. (2018). Deep anomaly detection with outlier exposure. *In: International Conference on Learning Representations*, pp 1–13
- Hendrycks\*, D., Mu\*, N., Cubuk, E.D. et al. (2020). Augmix: A simple method to improve robustness and uncertainty under data shift. *In: International Conference On Learning Representations*, pp 1–13
- Hendrycks, D., Basart, S., Mu, N., et al. (2021a). The many faces of robustness: A critical analysis of out-of-distribution generalization. *In: Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp 8340–8349
- Hendrycks, D., Zhao, K., Basart, S. et al. (2021b). Natural adversarial examples. *In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp 15,262–15,271
- Hendrycks, D., Basart, S., Mazeika, M. et al. (2022). Scaling out-of-distribution detection for real-world settings. *In: International Conference on Machine Learning*, PMLR, pp 8759–8773
- Hu, J., Shen, L., & Sun, G. (2018). Squeeze-and-excitation networks. *In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp 7132–7141
- Huang, G., Liu, Z., Van Der Maaten, L. et al. (2017). Densely connected convolutional networks. *In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp 4700–4708
- Huang, R., & Li, Y. (2021). Mos: Towards scaling out-of-distribution detection for large semantic space. *In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp 8710–8719
- Huang, R., Geng, A., & Li, Y. (2021). On the importance of gradients for detecting distributional shifts in the wild. *Advances in Neural Information Processing Systems*, 34, 677–689.
- Kirichenko, P., Izmailov, P., & Wilson, A. G. (2020). Why normalizing flows fail to detect out-of-distribution data. *Advances in neural information processing systems*, 33, 20578–20589.
- Kobyzev, I., Prince, S. J., & Brubaker, M. A. (2020). Normalizing flows: An introduction and review of current methods. *IEEE transactions on pattern analysis and machine intelligence*, 43(11), 3964–3979.
- Lee, K., Lee, K., Lee, H. et al. (2018). A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *In: Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pp 7167–7177
- Liang, S., Li, Y., Srikant, R. (2018). Enhancing the reliability of out-of-distribution image detection in neural networks. *In: International Conference on Learning Representations*, pp 1–12
- Lim, S., Kim, I., Kim, T. et al. (2019). Fast autoaugment. *In: Advances in Neural Information Processing Systems (NeurIPS)*
- Liu, W., Wang, X., Owens, J. et al. (2020). Energy-based out-of-distribution detection. *Advances in Neural Information Processing Systems* 33:21,464–21,475
- Liu, Z., Lin, Y., Cao, Y. et al. (2021). Swin transformer: Hierarchical vision transformer using shifted windows. *In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*
- Naeini, M.P., Cooper, G., & Hauskrecht, M. (2015). Obtaining well calibrated probabilities using bayesian binning. *In: Proceedings of the AAAI conference on artificial intelligence*
- Nalisnick, E., Matsukawa, A., Teh, Y.W. et al. (2018). Do deep generative models know what they don't know? *In: International Conference on Learning Representations*, pp 1–13
- Papadopoulos, A. A., Rajati, M. R., Shaikh, N., et al. (2021). Outlier exposure with confidence control for out-of-distribution detection. *Neurocomputing*, 441, 138–150.
- Paszke, A., Gross, S., Massa, F., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32, 8024–8035.
- Recht, B., Roelofs, R., Schmidt, L. et al. (2019). Do imagenet classifiers generalize to imagenet? *In: International Conference on Machine Learning*, pp 1–13
- Ren, J., Liu, P.J., Fertig, E. et al. (2019). Likelihood ratios for out-of-distribution detection. *In: Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pp 14,707–14,718
- Rusak, E., Schott, L., Zimmermann, R.S. et al. (2020). A simple way to make neural networks robust against diverse image corruptions. *In: European Conference on Computer Vision*, Springer, pp 53–69

- Salman, H., Ilyas, A., Engstrom, L., et al. (2020). Do adversarially robust imagenet models transfer better? *Advances in Neural Information Processing Systems*, 33, 3533–3545.
- Sandler, M., Howard, A., Zhu, M. et al. (2018). Mobilenetv2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp 4510–4520
- Serrà, J., Álvarez, D., Gómez, V., et al. (2019). Input complexity and out-of-distribution detection with likelihood-based generative models. In: International Conference on Learning Representations, pp 1–13
- Shen, Z., Savvides, M. (2020). Meal v2: Boosting vanilla resnet-50 to 80%+ top-1 accuracy on imagenet without tricks. arXiv preprint [arXiv:2009.08453](https://arxiv.org/abs/2009.08453)
- Simonyan, K., Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. arXiv preprint 14091556 <https://arxiv.org/abs/1409.1556> [cs.CV]
- Sun, Y., Li, Y. (2022). Dice: Leveraging sparsification for out-of-distribution detection. In: European Conference on Computer Vision, Springer, pp 691–708
- Sun, Y., Guo, C., & Li, Y. (2021). React: Out-of-distribution detection with rectified activations. *Advances in Neural Information Processing Systems*, 34, 144–157.
- Sun, Y., Ming, Y., Zhu, X. et al. (2022). Out-of-distribution detection with deep nearest neighbors. International Conference on Machine Learning pp 1–13
- Szegedy, C., Ioffe, S., Vanhoucke, V. et al. (2017). Inception-v4, inception-resnet and the impact of residual connections on learning. In: Thirty-first AAAI conference on artificial intelligence
- Thulasidasan, S., Chennupati, G., Bilmes, J. et al. (2019). On mixup training: improved calibration and predictive uncertainty for deep neural networks. In: Proceedings of the 33rd International Conference on Neural Information Processing Systems, pp 13,911–13,922
- Touvron, H., Cord, M., Douze, M., et al. (2021) Training data-efficient image transformers distillation through attention. In: International Conference on Machine Learning, pp 10,347–10,357
- Touvron, H., Bojanowski, P., Caron, M., et al (2022) Resmlp: Feedforward networks for image classification with data-efficient training. IEEE Transactions on Pattern Analysis Machine Intelligence pp 1–9. <https://doi.org/10.1109/TPAMI.2022.3206148>
- Van Horn, G., Mac Aodha, O., Song, Y. et al. (2018). The inaturalist species classification and detection dataset. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp 8769–8778
- Wang, H., Li, Z., Feng, L. et al. (2022). Vim: Out-of-distribution with virtual-logit matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp 4921–4930
- Wightman, R. (2019). Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, <https://doi.org/10.5281/zenodo.4414861>
- Wu, Z.F., Wei, T., Jiang, J. et al. (2021). Ngc: a unified framework for learning with open-world noisy data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp 62–71
- Xiao, J., Hays, J., Ehinger, K.A., et al. (2010). Sun database: Large-scale scene recognition from abbey to zoo. In: 2010 IEEE computer society conference on computer vision and pattern recognition, IEEE, pp 3485–3492
- Xiao, Z., Yan, Q., Amit, Y. (2020). Likelihood regret: An out-of-distribution detection score for variational auto-encoder. *Advances in neural information processing systems* 33:20,685–20,696
- Yang, J., Zhou, K., Li, Y. et al. (2021). Generalized out-of-distribution detection: A survey. arXiv preprint [arXiv:2110.11334](https://arxiv.org/abs/2110.11334)
- Yang, J., Wang, P., Zou, D. et al. (2022). OpenOOD: Benchmarking generalized out-of-distribution detection. In: Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track, [https://openreview.net/forum?id=gT6j4\\_tskUt](https://openreview.net/forum?id=gT6j4_tskUt)
- Zhou, B., Lapedriza, A., Khosla, A., et al. (2017). Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6), 1452–1464.
- Zhu, Y., Chen, Y., Xie, C. et al. (2022). Boosting out-of-distribution detection with typical features. In: Oh AH, Agarwal A, Belgrave D, et al (eds) *Advances in Neural Information Processing Systems*, pp 1–20. <https://openreview.net/forum?id=4maAiUt0A4>
- Zisselman, E., & Tamar, A. (2020) Deep residual flow for out of distribution detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp 13,994–14,003
- Zoph, B., Vasudevan, V., Shlens, J. et al. (2018). Learning transferable architectures for scalable image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 8697–8710

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.