This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# **OODD: Test-time Out-of-Distribution Detection with Dynamic Dictionary**

Yifeng Yang<sup>1</sup>, Lin Zhu<sup>1</sup>, Zewen Sun<sup>2</sup>, Hengyu Liu<sup>3</sup>, Qinying Gu<sup>4</sup>, Nanyang Ye<sup>1\*</sup> <sup>1</sup> Shanghai Jiao Tong University <sup>2</sup> Tianjin University

<sup>3</sup> The Chinese University of Hong Kong <sup>4</sup> Shanghai Artificial Intelligence Laboratory

{zxk1212, zhulin\_sjtu, ynylincoln}@sjtu.edu.cn, 3022244294@tju.edu.cn,

piang.lhy@link.cuhk.edu.hk,guqinying@pjlab.org.cn

### Abstract

Out-of-distribution (OOD) detection remains challenging for deep learning models, particularly when test-time OOD samples differ significantly from training outliers. We propose OODD, a novel test-time OOD detection method that dynamically maintains and updates an OOD dictionary without fine-tuning. Our approach leverages a priority queue-based dictionary that accumulates representative OOD features during testing, combined with an informative inlier sampling strategy for in-distribution (ID) samples. To ensure stable performance during early testing, we propose a dual OOD stabilization mechanism that leverages strategically generated outliers derived from ID data. To our best knowledge, extensive experiments on the OpenOOD benchmark demonstrate that OODD significantly outperforms existing methods, achieving a 26.0% improvement in FPR95 on CIFAR-100 Far OOD detection compared to the state-of-the-art approach. Furthermore, we present an optimized variant of the KNN-based OOD detection framework that achieves a 3x speedup while maintaining detection performance. Our code is available at https://github.com/zxk1212/00DD.

### 1. Introduction

Deep learning has demonstrated impressive capabilities but often exhibits unpredictable behavior when faced with unknown situations, such as receiving data unrelated to its training tasks. What's worse, the model tends to incorrectly classify unknown out-of-distribution (OOD) data into one of the in-distribution (ID) classes with high confidence [10, 26]. Existing solutions in OOD detection areas can be broadly categorized into three main types: (i) post-hoc methods, (ii) fine-tuning methods, and (iii) test-time calibrated methods.

Post-hoc methods typically design an OOD score function to analyze model outputs (e.g., logits), features, or layer-wise



Figure 1. Before calibration, we dynamically feed the lower tail of the OOD score distribution into an OOD dictionary through a priority queue and then use this dictionary to calibrate the OOD scores.

statistics [7, 13, 21–23, 31]. During inference, these methods discriminate between ID and OOD samples based on the OOD scores of test samples. In contrast, fine-tuning methods [4, 14, 37, 41] require fine-tuning the model that has been pre-trained on ID data, using a combination of an auxiliary dataset with outliers and ID samples. These methods are often called outlier exposure (OE) based methods, aiming to make the classifier outputs more uniformly distributed when exposed to outliers. However, when the test-time OOD samples differ significantly from the training outliers, the model may still assign high confidence to test-time OOD samples [35]. While recent works [1, 8, 36] attempt to calibrate OOD detection during test time, their batch-wise parameter updates can lead to catastrophic forgetting, unstable detection performance, and even impact ID classification. We consider that the key to test-time OOD detection is utilizing the latent OOD features. Particularly in the medical field, when dealing with unknown diseases during the early stages of an epidemic, there is an urgent need to specifically collect the features of these unknown diseases [9].

Given this, we propose a method called **OODD** (*test-time* **Out-Of-D**istribution detection with dynamic **D**ictionary), which does not impose any burden during the training phase, nor does it require optimization or fine-tuning at test time. It only requires the addition of a dynamic OOD dictionary

<sup>\*</sup>Nanyang Ye is the corresponding author.

managed by a priority queue, enabling calibration at test time with just a single matrix multiplication with the keys in the OOD dictionary. The intuition behind our method is illustrated in Figure 1. Specifically, based on the OOD scores at test time, we collect features of OOD samples in real time, focusing on those with lower scores in the left tail of the OOD score distribution. These features are then stored in the dynamic OOD dictionary by a priority queue. During testing, we calculate the similarity between the features of the test sample and the features stored in the dynamic OOD dictionary. Then, according to the similarity between the sample's feature and the features stored in the dynamic OOD dictionary, we can adaptively calibrate the OOD score. Additionally, by using a priority queue to manage the features in the OOD dictionary, our approach mitigates catastrophic forgetting and detection fluctuation in other methods by updating parameters with minibatches. Notably, we designed a strategy to initialize this priority queue with outliers generated from the training ID data, alleviating the detection fluctuations caused by inaccurate estimates during the early testing phase. To evaluate the effectiveness of our method, we conducted comprehensive experiments on the OpenOOD benchmark [38], which is one of the authoritative platforms for testing OOD detection methods.

Extensive experiments confirm that our method can effectively calibrate the OOD score during testing. Compared to the most relevant baseline TULIP [1], our method significantly improves the FPR95 from 58.17% to 24.74% on CIFAR-100 Far OOD. Our key contributions are as follows:

- We propose a novel test-time OOD detection method based on a dynamic OOD dictionary. Unlike previous approaches, it doesn't require prior outlier samples from auxiliary datasets. Additionally, we devise a specific initialization for the OOD dictionary to alleviate the initial test-time fluctuations.
- Our approach achieves state-of-the-art performance on the OpenOOD benchmark, surpassing current baselines. Moreover, it demonstrates strong complementarity with existing post-hoc methods, effectively enhancing their performance when used in combination.
- To optimize computational efficiency, we use cosine similarity to replace the Euclidean distance used in KNN-based OOD detection methods without comprising performance. We provide a rigorous theoretical analysis establishing the relationship between these distance metrics and justify their equivalence in discriminative ability for OOD detection.

# 2. Related Works

Learning from auxiliary data. According to auxiliary data type, learning from auxiliary data methods can be classified into two categories: active learning and outlier exposure. In the first category, Scone [2] addresses OOD detection by uti-

lizing auxiliary data consisting of a mixture of ID and OOD samples. They assign pseudo-labels to this auxiliary data, enabling the model to adapt to these pseudo-labeled samples. Other notable works in this category include AL-OOD [28], DUL [39]. The second category, known as Outlier Exposure (OE) [14], uses auxiliary data where the class space usually does not overlap with ID or test-time OOD samples. By leveraging such data, the aim is to make the model more robust against test-time OOD samples [4, 14, 37, 41].

Test-Time Adaptation. Test-time adaptation methods can generally be categorized as: (i) Test-Time Domain Adaptation (TTDA): In this situation, we can obtain samples of unlabeled data in test domain once before deployment [29]. However, for OOD detection scenarios, it's difficult to gather all potential test domain data beforehand, particularly in the dynamic open world. (ii) Test-Time Batch Adaptation (TTBA): the model undergoes adaptive updates each time it processes a new batch, helping it handle distributional shifts over time. RTL++ [8] assumes a linear correlation between OOD scores and features, but the feature distribution overlap between hard OOD and ID samples challenges this simple assumption. AUTO [36] directly modifies the model at test time using stochastic gradient descent to produce lowconfidence predictions for potential OOD samples. Our work also belongs to this category, but we don't require strong assumptions or fine-tuning. (iii) Test-Time Prior Adaptation (TTPA): This adjusts prior information within the model, such as class or feature distribution at test time. For example, TUIP [1] introduces priors for ID and OOD samples at test time to calibrate model uncertainty.

**Deep K-Nearest Neighbor for OOD detection.** [31] proposes using K-nearest neighbors (KNN) for OOD detection. The approach first normalizes the penultimate layer features of a neural network and computes the Euclidean distance between a test sample's feature and all training sample features. The opposite of the distance to the K-th nearest neighbor is then used as the OOD score of this test sample. Despite its simplicity, this KNN-based method nearly outperforms other post-hoc methods and even fine-tuning methods in OpenOOD benchmark [38].

# 3. Preliminaries

Setup. We consider a supervised multi-class classification setting where the input space is denoted by  $\mathcal{X}$  and the label space by  $\mathcal{Y} = \{1, 2, ..., C\}$ . The training set  $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ is drawn *i.i.d.* from the joint distribution  $P_{\mathcal{X}\mathcal{Y}}$ . Let  $f : \mathcal{X} \mapsto \mathbb{R}^{|\mathcal{Y}|}$  be a model parameterized by  $\theta$ , trained on samples drawn from  $P_{\mathcal{X}\mathcal{Y}}$ , which outputs logits. Then the logits are transformed by a softmax layer:  $\hat{p}_k = \frac{e^{f_k}}{\sum_{j=1}^C e^{f_j}}$  for  $k \in \mathcal{Y}$ , where the class probability of a sample  $\mathbf{x}$  is defined as  $P(y = k \mid \theta, \mathbf{x}) = \hat{p}_k$ , and the confidence is defined as  $Conf(\mathbf{x}) = \max_y P(y \mid \theta, \mathbf{x})$ . Importantly, we could use the latent feature  $\mathbf{z} = \phi(\mathbf{x})$  for OOD detection, where  $\phi$ :



Figure 2. Overview of the proposed OODD framework, which performs real-time, batch-wise updates to the OOD dictionary using a priority queue based on the left-tail distribution of OOD score. Using training ID samples, we generate multiple random crops for each sample [30], which are then filtered at the patch and class levels. High-confidence inliers are selected to compute the latent OOD score, while low-confidence outliers initialize the priority queue, allowing for a more robust and adaptive OOD detection process.

 $\mathcal{X} \mapsto \mathbb{R}^d$  is an image encoder, and d is the dimensionality of the latent space.

**Problem Statement.** At test time, our goal is to find a good OOD detector, which can distinguish whether a sample  $\mathbf{x} \in \mathcal{X}$  is from ID or not OOD. The decision function  $D_{\tau}$  is usually defined as follows:

$$D_{\tau}(\mathbf{x}) = \begin{cases} ID & S(\mathbf{x}) \ge \tau\\ OOD & S(\mathbf{x}) < \tau \end{cases},$$
(1)

where samples with lower OOD scores  $S(\mathbf{x})$  are classified as OOD and vice versa, and  $\tau$  is a handcraft threshold, which is commonly chosen so that a high fraction (e.g., 95%) of ID data is correctly distinguished.

### 4. Methodology

In this section, we present our proposed method for informative inlier and dynamic OOD dictionary construction. We first describe the process of informative inlier sampling (IIS) in Sec 4.1, where key patches with high confidence are selected to create a representative ID dictionary. In Sec 4.2, we extend the KNN-based method to build a dynamic OOD dictionary using a priority queue that is progressively updated with latent OOD samples to calibrate OOD score. To stabilize the OOD detection process, we introduce a Dual OOD Stabilization (DOS) strategy. Finally, we propose three approaches for obtaining outliers, which are crucial for initializing the OOD dictionary and stabilizing the detection process. An overview of the proposed method is shown in Figure 2.

### 4.1. Informative Inlier Sampling

Unlike the KNN-based OOD detection [31], which uses all or randomly selected training data as an ID dictionary to

compute the similarity between the query (latent features) of a test sample, we propose that the ID dictionary should be more informative and representative. To sample the informative inliers from training ID dataset, we capture the key patch of samples with high OOD scores at the patch level and the class level respectively. Specifically, to sufficiently explore the vicinal space of training samples, we perform multiple random cropping [30] on each training ID sample  $\mathbf{x}_i$  to obtain the set  $X_i^{crop} = {\{\mathbf{x}_{i,m}^{crop}\}}_{m=1}^M$ , where M is the number of random cropping. Then, based on the confidence of each patch  $\operatorname{Conf}(\mathbf{x}_{i,m}^{crop})$  for each sample we select the patch with the line  $X_i^{crop}$ . ple, we select the patch with the highest confidence. After selecting the highest-confidence patches for all samples, we further choose the top  $\alpha\%$  of patches based on their confidence within each class. The latent features of the final selected patches are then saved in ID dictionary, denoted as  $\mathcal{K}_{n'}^{id} = \{k_1^{id}, k_2^{id}, \dots, k_{n'}^{id}\}$ , where  $n' = [\alpha \% \cdot n], n$  is the number of training samples and  $[\cdot]$  denotes the rounding function.

# 4.2. Dynamic OOD Dictionary

In a similar way to constructing the ID dictionary, we can easily extend this approach to construct the OOD dictionary, but the OOD dictionary should be dynamic in the sense that OOD samples randomly appear and are sampled as keys by latent OOD scores  $S_{in}(\mathbf{x})$  during testing. Our hypothesis is that good OOD latent features can be saved in the OOD dictionary that covers a rich set of OOD samples, while the OOD dictionary keys are kept as consistent as possible despite its evolution.

**Latent OOD Score.** we derive the query  $q(\mathbf{x}^*) = \Phi(\mathbf{x}^*)$ for a test sample  $\mathbf{x}^*$  and compute the cosine similarity  $\cos(k_{id}^{id}, q)$  with each key  $k_{id}^{id}$  in ID dictionary  $\mathcal{K}_{n'}^{id}$ . Then, we denote the sorted list of these similarities in descending order as  $\cos(k_{(1)}^{id}, q) \ge \cos(k_{(2)}^{id}, q) \ge \cdots \ge \cos(k_{(n')}^{id}, q)$ . Finally, the latent OOD score is given by:

$$S_{in}(\mathbf{x}^*) = \cos(k_{(\mathbb{K})}^{id}, q) \tag{2}$$

Where the  $\mathbb{K}$ -th largest cosine similarity serves as the latent OOD score  $S_{in}(\mathbf{x}^*)$  for the test-time sample.

OOD dictionary as a priority queue. During test time, the core of our motivation is maintaining the OOD dictionary as a priority queue of potential OOD sample keys (or latent features). This configuration enables the reuse of encoded keys from previous mini-batches by leveraging a priority queue structure that separates the dictionary size from the minibatch size. Consequently, the OOD dictionary can be much larger than the mini-batch, and the OOD dictionary size can be easily adjusted as a hyperparameter. Dictionary samples are progressively updated, with lower potential OOD scores (i.e., more likely to be OOD) replacing existing samples. The priority queue maintains the sample with the highest OOD score at the front. A new sample from the mini-batch is enqueued only if it has a lower OOD score than the queue front, and the front needs to be dequeued first when the dictionary is full. Further, we denote this OOD dictionary as  $\mathcal{K}_{l}^{ood} = \{k_{1}^{ood}, k_{2}^{ooid}, \dots, k_{l'}^{ood}\}$  where l represents the size of the proposed priority queue and  $l \ge l'$ . Thus, the OOD dictionary  $\mathcal{K}_{l}^{ood}$  represents a dynamically sampled latent OOD subset of the test-time data, and keeping the OOD dictionary up to date incurs little computational cost.

**Dual OOD Stabilization.** Although subsequent experiments show that initializing the OOD dictionary as empty has little impact on OOD detection results in most scenarios, the empty OOD dictionary may lead to highly unstable scores during the early stage. To maintain historical information inertia, we propose the Dual OOD Stabilization (DOS) strategy to mitigate substantial fluctuations in calibrated OOD scores caused by dramatic changes in latent OOD data within the priority queue. Specifically, we utilize some outliers (with their obtained approaches further detailed later) as a memory bank and others to initialize the priority queue. Under this setup, the OOD dictionary encompasses not only the priority queue but also the memory bank, which can be represented as:  $\mathcal{K}_{total}^{ood} = \mathcal{K}_{l}^{ood} \cup \mathcal{K}_{mb}^{ood}$  where  $\mathcal{K}_{mb}^{ood}$  represents samples in the memory bank, and the memory bank size is denoted as mb. Similar to Equation 2, we can obtain the negative cosine similarity between the query and each of the keys in the OOD dictionary during testing, as follows:

$$S_{out}(\mathbf{x}^*) = -\cos(k_{(\hat{\mathbb{K}})}^{ood}, q) \tag{3}$$

where  $\hat{\mathbb{K}}$ -th largest cosine similarity serves as the calibrated OOD score  $S_{out}(\mathbf{x}^*)$  for the test-time sample. Finally,the OOD score in our method is integrated as follows:

$$S(\mathbf{x}^*) = S_{in}(\mathbf{x}^*) + S_{out}(\mathbf{x}^*)$$
(4)

However, a critical yet fully underexplored question thus arises:

#### How do we obtain these outliers?

We explore three approaches for obtaining outliers in our subsequent experiments:

- C-Out (Cropping Outliers): Use multiple random cropping augmentation [30] on the training ID data to generate low-confidence outliers that originate from ID data.
- ▷ T-Out (Target Outliers): Introduce a small amount of data that shares the same distribution as the test-time OOD data as outliers.
- ▷ D-Out (Different test-time Outliers): This type of outlier is derived from a distribution different from both the ID and test-time OOD data.

Notably, our proposed C-Out method does not introduce any prior knowledge beyond training ID. The process it uses to generate outliers is the reverse of what is described in Sec 4.1. In contrast, T-Out incorporates priors aligned with the test-time OOD distribution, while D-Out needs to leverage samples from external datasets. Subsequent experiments demonstrate that C-Out alone can achieve strong performance, even without using priors beyond the training ID data.

### 5. Experiments

#### 5.1. Experiments Setup

**Datasets.** We use four popular ID datasets CIFAR-10/100 [18], ImageNet-200/1K [6]. Following OpenOOD benchmark [38], the OOD testing datasets are categorized into two groups: Near OOD and Far OOD. Specifically, for CIFAR-10/100 [18] benchmarks, the Far OOD group includes MNIST [20], SVHN [25], Textures [5], Places365 [40], and the Near OOD group comprises CIFAR-100/10 and Tiny ImageNet-200 [19]. For ImageNet-200/1K, the Near OOD group includes SSB-hard [33], NINCO [3], and the Far OOD group comprises iNaturalist [32], Textures [5], and OpenImage-O [34].

**Evaluation Metrics.** We use three metrics to evaluate the results: (1) FPR95, the false positive rate for OOD when the ID true positive rate is 95%; (2) AUROC, the area under the receiver operating characteristic curve; and (3) ID ACC, the ID classification accuracy. Our reported results are averaged over three independent runs with different random seeds.

**Pre-trained Model.** For CIFAR-10/100 and ImageNet-200, we employed the ResNet-18 model [11] trained with empirical risk minimization on each ID training set provided by OpenOOD [38]. For ImageNet-1K, we use a pre-trained ResNet-50 model [12] from the PyTorch.

**Baseline Methods.** We compared our method with various post-hoc methods, including MaxLogits [15], ODIN [22], Mahalanobis Distance (MD)[21], Energy [23], KNN

			Near	OOD	Far OOD		
Dataset	Method	ID ACC↑	AUROC↑	FPR95↓	AUROC↑	FPR95↓	
CIEAD 10	MaxLogits [15]	95.06 ± 0.30	87.52 ± 0.47	$61.32 \pm 4.62$	$91.10 \pm 0.89$	$41.68 \pm 5.27$	
	ODIN [22]	95.06 ± 0.30	82.87 ± 1.85	$76.19 \pm 6.08$	$87.96 \pm 0.61$	$57.62 \pm 4.24$	
	MD [21]	95.06 ± 0.30	$84.20 \pm 2.40$	$49.90 \pm 3.98$	$89.72 \pm 1.36$	$32.22 \pm 3.40$	
	Energy [23]	95.06 ± 0.30	$87.58 \pm 0.46$	$61.34 \pm 4.63$	$91.21 \pm 0.92$	$41.69 \pm 5.32$	
CIFAR-10	KNN [31]	95.06 ± 0.30	$90.64 \pm 0.20$	$34.01 \pm 0.38$	$92.96 \pm 0.14$	$24.27 \pm 0.40$	
	ViM [7]	95.06 ± 0.30	$88.68 \pm 0.28$	$44.84 \pm 2.31$	$93.48 \pm 0.24$	$25.05 \pm 0.52$	
	RTL++ [8]	95.06 ± 0.30	$88.76 \pm 0.01$	$54.03 \pm 0.21$	$87.06 \pm 0.07$	$64.30 \pm 0.21$	
	TULIP [1]	95.06 ± 0.30	89.67 ± 0.24	$33.80 \pm 0.59$	$92.55 \pm 0.13$	$24.43 \pm 0.17$	
	OODD	$95.06 \pm 0.30$	$90.96 \pm 0.21$	$36.01 \pm 0.44$	$95.77\pm0.12$	$17.44 \pm 0.44$	
	MaxLogits [15]	77.25 ± 0.10	81.05 ± 0.07	55.47 ± 0.66	79.67 ± 0.57	56.73 ± 1.33	
	ODIN [22]	$77.25 \pm 0.10$	$79.90 \pm 0.11$	$57.91 \pm 0.51$	$79.28 \pm 0.21$	$58.86 \pm 0.79$	
	MD [21]	$77.25 \pm 0.10$	$58.69 \pm 0.09$	$83.53 \pm 0.60$	$69.39 \pm 1.39$	$72.26 \pm 1.56$	
CIEAD 100	Energy [23]	$77.25 \pm 0.10$	$80.91 \pm 0.08$	$55.62 \pm 0.61$	79.77 ± 0.61	$56.59 \pm 1.38$	
CIFAK-100	KNN [31]	$77.25 \pm 0.10$	$80.18 \pm 0.15$	$61.22 \pm 0.14$	$82.40 \pm 0.17$	$53.65 \pm 0.28$	
	ViM [7]	$77.25 \pm 0.10$	$74.98 \pm 0.13$	$62.63 \pm 0.27$	$81.70 \pm 0.62$	$50.74 \pm 1.00$	
	RTL++ [8]	77.25 ± 0.10	$79.24 \pm 0.03$	$63.92 \pm 0.27$	$80.76 \pm 0.03$	$67.14 \pm 0.07$	
	TULIP [1]	77.25 ± 0.10	$81.29 \pm 0.26$	$55.07 \pm 0.73$	$79.63 \pm 0.94$	$58.17 \pm 1.78$	
	OODD	$77.25 \pm 0.10$	$82.10 \pm 0.20$	$56.98 \pm 1.09$	93.64 ± 0.31	$24.74 \pm 1.30$	

Table 1. Comparison with state-of-the-art methods on the CIFAR benchmarks, where the best results are in bold.

Table 2. Comparison with state-of-the-art methods on the ImageNet benchmarks, where the best results are in bold.

			Near	OOD	Far OOD		
Dataset	Method	ID ACC↑	AUROC↑	FPR95↓	AUROC↑	FPR95↓	
	MaxLogits [15]	86.37± 0.08	$82.90 \pm 0.04$	59.76 ± 0.59	91.11 ± 0.19	$34.03 \pm 1.21$	
ImageNet-200	ODIN [22]	86.37 ± 0.08	$80.27 \pm 0.08$	$66.76 \pm 0.26$	$91.71 \pm 0.19$	$34.23 \pm 1.05$	
	MD [21]	86.37± 0.08	$61.93 \pm 0.51$	$79.11 \pm 0.31$	$74.72 \pm 0.26$	$61.66 \pm 0.27$	
	Energy [23]	86.37± 0.08	$82.50 \pm 0.05$	$60.24 \pm 0.57$	$90.86 \pm 0.21$	$34.86 \pm 1.30$	
	KNN [31]	86.37± 0.08	81.57 ± 0.17	$60.18 \pm 0.52$	$93.16 \pm 0.22$	$27.27 \pm 0.75$	
	ViM [7]	86.37± 0.08	$78.68 \pm 0.24$	$59.19 \pm 0.71$	$91.26 \pm 0.19$	$27.20 \pm 0.30$	
	RTL++ [8]	86.37± 0.08	$81.02 \pm 0.08$	$63.99 \pm 0.29$	$93.11 \pm 0.02$	$30.38 \pm 0.10$	
	TULIP [1]	86.37± 0.08	$83.77 \pm 0.06$	$54.51 \pm 0.35$	$91.03 \pm 0.09$	$33.94 \pm 0.51$	
	OODD	86.37± 0.08	$85.74 \pm 0.03$	$53.70 \pm 0.14$	$95.47 \pm 0.03$	$20.89 \pm 0.12$	
	MaxLogits [15]	76.18	76.46	67.82	89.57	38.22	
	ODIN [22]	76.18	74.75	72.50	89.47	43.96	
	MD [21]	76.18	55.44	85.45	74.25	62.92	
ImagaNat 1V	Energy [23]	76.18	75.89	68.56	89.47	38.39	
Imageinet-1K	KNN [31]	76.18	71.10	70.87	90.18	34.13	
	ViM [7]	76.18	72.08	71.35	92.68	24.67	
	RTL++ [8]	76.18	76.49 ± 0.49	$71.49 \pm 0.35$	$90.53 \pm 0.39$	$44.39 \pm 1.19$	
	TULIP [1]	76.18	77.52 ± 0.06	$64.96 \pm 0.08$	$88.03 \pm 0.02$	$48.01 \pm 0.02$	
	OODD	76.18	$72.66 \pm 0.16$	$71.41 \pm 0.27$	$93.84 \pm 0.44$	$27.63 \pm 1.16$	

[31], ViM [7], RTL++ [8], and TULIP [1]. In addition, to the best of our knowledge, TULIP—the test-time method achieving state-of-the-art performance on the OpenOOD benchmark—was also included in our comparison.

**Remark on the implementation.** We implemented our method in PyTorch and conducted all experiments on a single NVIDIA GeForce RTX 4090 GPU. The ID dictionary only needs to be extracted once before deploying the trained model. Additionally, unlike KNN-based OOD scoring, which uses faiss.IndexFlatL2 with Euclidean distance, we compute the OOD score through GPU-accelerated matrix multiplication in PyTorch to obtain cosine similarity scores directly, so our approach is more memory-efficient and faster than the former. In all experiments, we set the batch size to 512. In the main results presented in Tables 1 and 2, we use an inlier sampling ratio of  $\alpha\% = 50\%$  with M = 4,  $\hat{\mathbb{K}} = 5$ . For CIFAR-10/100, due to the smaller dataset size, we set the priority queue size l to 128/512 respectively, and the memory bank size mb is 5. Additionally, for CIFAR-10, we

use  $\mathbb{K} = 5$  and for CIFAR-100,  $\mathbb{K} = 10$ . For larger datasets ImageNet-200/1K, we increase l to 2048 and mb to 128, and  $\mathbb{K} = 100$ .

# 6. Main Results.

**CIFAR Benchmarks.** Firstly, we compare the proposed methods using C-Out (Cropping Outliers) with state-of-theart OOD detection methods. The main experimental results on CIFAR-10/100 are shown in Table 1, from which the following conclusions can be drawn: (i) The proposed method OODD outperforms other methods in most cases. For example, on CIFAR-100, the Far OOD AUROC of our method reaches 93.64%, while the second-best method (KNN) achieves 82.40%. (ii) Our method maintains a comparable ability for Near OOD detection. While the tail distributions of Near OOD and ID samples overlap considerably, our proposed approach does not significantly impair Near OOD detection performance. On CIFAR-100, our method even achieves state-of-the-art AUROC, though it does not lead in FPR95. However, AUROC provides a more comprehensive assessment than FPR95, which evaluates only a single point.

**ImageNet Benchmarks.** For ImageNet-200/1K, prior works [16, 35] show that many advanced methods developed on CI-FAR benchmarks struggle with the large semantic space of ImageNet. Nevertheless, as shown in Table 2, our method achieves leading performance in both Near OOD and Far OOD detection on ImageNet-200. Additionally, although KNN performs moderately well on ImageNet-1K, our method significantly enhances performance in Far OOD without compromising its near OOD detection capability.

#### 6.1. Ablation Study

In this section, we conduct ablation experiments on ImageNet-200 to analyze the effectiveness of each module and explore the impact of various hyperparameters. The main ablation study results are shown in Table 3, which demonstrate the effectiveness of our proposed informative inlier sampling and dynamic OOD dictionary strategies.

Table 3. The main ablation study results on ImageNet-200. IIS denotes informative inlier sampling, PR denotes priority queue, and MB denotes memory bank.

		Near	: 00D	Far OOD		
IIS	PR	MB	FPR95	AUROC	FPR95	AUROC
X	X	X	60.57	80.73	30.35	91.47
1	X	×	58.92	81.67	26.97	92.83
X	1	×	53.85	85.69	21.47	95.45
X	X	1	59.14	81.30	37.09	88.89
✓	1	1	53.70	85.74	20.89	95.47

The effectiveness of informative inlier sampling. To examine this, we employ a naive KNN OOD detection [31], adjusting both the sampling ratio  $\alpha$ % and the hyper-parameter K, which defines the K-th largest cosine similarity in  $S_{in}(\mathbf{x}^*)$ . When choosing sampling ratio  $\alpha$ % for training data for nearest neighbor search, K is selected from {1, 5, 10, 20, 50, 100, 200, 500, 1000} accordingly to the validation dataset OpenImage-O. As shown in Figure 3, the ablation results highlight how variations in  $\alpha$  influence performance.

The effectiveness of the priority queue. In Figure 4, we systematically analyze the effect of the priority queue size l under the conditions of not using a memory bank, with a batch size of 512. We vary the queue size l. Several key observations can be made: (i) In non-extreme cases, our method is not sensitive to l, with the results for both Near OOD and Far OOD showing a significantly lower FPR95 compared to the naive KNN method. (ii) When l is particularly large, the effect of the priority queue becomes less prominent. Under this situation, the accumulation of ID data



Figure 3. Results of varying  $\alpha$  on ImageNet-200 ID. IIS refers to the implementation of Informative Inlier Sampling based on the naive KNN method [31]. The value of  $\alpha$  is on the horizontal axis, and the FPR95 is on the vertical axis.

being added to the OOD dictionary increases, which leads to a decline in OOD detection.



Figure 4. Comparison of AUROC detection performance using different queue sizes.



Figure 5. AUROC comparison using different types of outliers (C-Out, T-Out, D-Out, None) to initialize the OOD dictionary.

The different initializations of OOD dictionary. In Figure 5, we show results for initializing the OOD dictionary on the ImageNet-200 benchmark with various types of outliers (C-Out, T-Out, D-Out, None). The experimental results indicate that the choice of initialization method has minimal impact on overall performance. To further investigate how initialization methods affect test-time performance, we set up a new experiment: selecting 10,000 samples from ImageNet-200 as ID dataset and mixing them with 100 OOD samples from OpenImage-O, shuffled randomly for testing. This experiment was repeated five times, and we recorded the AUROC for each of the first 10 iterations, as shown in Figure 6. The results reveal that, without outliers to initialize the prior queue, OOD detection exhibits larger fluctuations in the first



Figure 6. AUROC of test-time performance over the first 10 iterations for 10,000 ID samples from ImageNet-200 mixed with 100 OOD samples from OpenImage-O, averaged across five experimental runs with standard deviation.

few iterations. However, as the iterations progress, AUROC values stabilize across different initialization methods, showing minimal differences. Notably, while the lack of prior OOD knowledge results in lower early performance compared to T-Out, our proposed C-Out initialization method achieves the most stable performance. More experimental results can be seen in Table 8 of the supplementary material.

How to choose  $\mathbb{K}$  and  $\hat{\mathbb{K}}$ ? Following prior work [31] and the OpenOOD benchmark [38], we first select the parameter  $\mathbb{K}$  based on the OOD validation set provided by OpenOOD. Specifically, for CIFAR-10/100, the OOD validation set is tiny ImageNet-200, while for ImageNet-200/1K, the OOD validation set is OpenImage-O. All validation set data does not overlap with the test set data. Then, we are surprised to find that a relatively small value of  $\hat{\mathbb{K}}$  yielded strong performance across all experiments. As a result, we fixed  $\hat{\mathbb{K}} = 5$  for all our experiments.

Table 4. Time cost comparison of different OOD detection methods.

Method	CIFAR-10	ImageNet-200	Time cost factor
MaxLogits	1.20 s	1.45 s	1x
KNN	9.61	30.17 s	15x
RTL++	82.07 s	103.80 s	70x
OODD	3.98 s	10.13 s	5x

**Time Cost.** To provide a more intuitive and fair comparison of computational overhead in OOD detection, we specifically measured the time cost from obtaining the image encoder's features to generating the final OOD scores. As shown in Table 4, Maxlogits exhibits the shortest processing time as it only involves a single linear classifier operation. RTL++ incurs the longest time cost due to the need to solve a Lasso regression between the latent features and the OOD scores [8]. Meanwhile, KNN demands 15 times more computational time than Maxlogits. In comparison to KNN, our proposed method reduces the time cost by a factor of 3, highlighting its practical applicability.

## 7. Further Discussion

#### 7.1. Integrate with the CLIP-based methods

When using the Contrastive Language-Image Pre-training (CLIP) model as the encoder [27], our method can be effectively integrated into various CLIP-based OOD detection methods, such as MCM [24] and NegLabel [17]. As shown in Table 5, when combined with these CLIP-based methods, OODD significantly enhances performance across most datasets.

Table 5. OOD detection performance for ImageNet-1K as ID

	iNaturalist		SUN		Places		Texture	
Method	$FPR95\downarrow$	AUROC↑	$FPR95\downarrow$	AUROC↑	$FPR95\downarrow$	AUROC↑	$FPR95\downarrow$	AUROC↑
MCM	30.91	94.61	37.59	92.57	44.69	89.77	57.77	86.11
+OODD	2.22	99.36	21.49	95.01	44.76	87.10	30.69	93.27
Improve	∇-28.69	△+4.75	∇-16.10	△+2.44	∆+0.07	∇-2.67	∇-27.08	△+7.16
NegLabel	1.91	99.49	20.53	95.49	35.59	91.64	43.56	90.22
+OODD	0.85	99.79	12.94	97.17	30.68	92.51	30.67	94.51
Improve	∇-1.06	∆+0.30	∇-7.59	$\triangle +1.68$	∇-4.91	△+0.87	∇-12.89	△+4.29

#### 7.2. The Temporal Drift OOD Scenarios

Since OODD is a test-time method, it is crucial to address temporal drift where OOD scenarios evolve over time. For example, consider a temporal drift with time points  $t_0$ ,  $t_1$ ,  $t_2$ , and  $t_3$ . The temporal drift of OOD scenarios goes from Textures ( $t_0$ - $t_1$ ) to Places365 ( $t_1$ - $t_2$ ) and then to SVHN ( $t_2$  $t_3$ ), denoted as  $T \rightarrow P \rightarrow S$  in Table 6. After incorporating our OODD method under the temporal drift OOD scenarios, we observe improvements across all datasets. Furthermore, the earlier a dataset is encountered, the greater the gain, demonstrating the effectiveness of our approach.

Table 6. Results for the temporal drift OOD scenario settings using CIFAR-100 as ID dataset. T represents Texture, P for Places365, and S for SVHN.

	Textures		Places365		SVHN		Mean	
Temporal Drift	$FPR95\downarrow$	AUROC↑	$FPR95\downarrow$	AUROC↑	$FPR95\downarrow$	AUROC↑	$FPR95\downarrow$	AUROC↑
$T \rightarrow P \rightarrow S$	35.33	91.80	53.61	83.27	18.68	97.13	35.87	90.73
$P \rightarrow S \rightarrow T$	45.42	86.25	50.60	85.16	17.76	97.31	37.93	89.57
$S \to T \to P$	42.40	87.77	58.63	80.19	8.12	98.69	36.39	88.88
w/o OODD	53.56	83.66	60.70	79.43	51.75	84.15	55.34	82.41

#### 7.3. Combine with More Post-hoc Methods

**Our approach further enhances the effectiveness of various post-hoc OOD detection methods.** We explore multiple post-hoc methods to identify latent OOD features. As shown in Table 7, models equipped with the OOD dictionary consistently outperform their counterparts (i.e., original OOD detection performance), validating the effectiveness of the OOD dictionary mechanism. Notably, KNN with the OOD dictionary demonstrates a significant improvement, while the gains for MaxLogits and Energy methods are relatively modest. This discrepancy may be attributed to our OOD dictionary storing latent features, which have limited benefits for logits-based OOD detection methods. We hope that future work will focus on designing a more comprehensive OOD dictionary to enhance these methods further.

Table 7. AUROC performance comparison for whether the model is equipped with our proposed OOD dictionary.

In Dataset	CIFA	R-10	CIFA	R-100	Imagel	Net-200	Imagel	Net-1K
Method	Near	Far	Near	Far	Near	Far	Near	Far
MaxLogits	87.52	91.10	81.05	79.67	82.90	91.11	76.46	89.57
+Ours	88.34	91.36	81.01	86.56	83.87	92.03	77.44	90.08
Improve	∆+0.82	∆+0.26	∇-0.04	$\triangle$ +6.89	$\triangle$ +0.97	∆+0.92	∆+0.98	△+0.51
Energy	87.58	91.21	80.91	79.77	82.50	90.86	75.89	89.47
+Ours	88.41	91.59	80.98	87.74	83.61	91.86	76.92	89.97
Improve	∆+0.83	∆+0.38	∆+0.07	△+7.97	$\triangle +1.11$	$\triangle +1.00$	△+1.03	$\triangle +0.50$
KNN	90.64	92.96	80.18	82.40	81.57	93.16	71.10	90.18
+Ours	90.96	95.77	82.10	93.64	85.74	95.47	72.66	93.84
Improve	△+0.32	△+2.81	△+1.92	△+11.24	△+4.17	△+2.31	△+1.56	∆+3.66

### 7.4. A Closer Look at KNN-based OOD Detection

For OOD detection, cosine similarity equals Euclidean distance after feature normalization but is more efficient. The KNN-based OOD detection [31] shows that ID data tends to have larger Euclidean distances before normalization, which undermines assumptions about KNN distances in high-dimensional space. Fortunately, the latent features of ID data have larger  $\ell_2$ -norm values than those of OOD data, which effectively mitigates the dimension curse when calculating Euclidean distances. Thus, given two vectors  $z_1$  and  $\mathbf{z}_2$ , we first apply  $\ell_2$ -norm normalization to each vector to obtain their normalized forms,  $\hat{\mathbf{z}}_1$  and  $\hat{\mathbf{z}}_2$ , where  $\hat{\mathbf{z}}_i = \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|_2}$ for i = 1, 2. After normalization, the Euclidean distance and cosine similarity between  $\hat{z}_1$  and  $\hat{z}_2$  are closely related. For the above two normalized vectors, the Euclidean distance  $d_E$  can be expressed as a function of their cosine similarity as follows:

$$d_E(\hat{\mathbf{z}}_1, \hat{\mathbf{z}}_2) = \|\hat{\mathbf{z}}_1 - \hat{\mathbf{z}}_2\|_2 = \sqrt{2 - 2\cos(\mathbf{z}_1, \mathbf{z}_2)} \quad (5)$$

We are surprised to find that the cosine similarity  $\cos (\mathbf{z}_1, \mathbf{z}_2)$ used in our method has a monotonic relationship with OOD score  $-d_E (\hat{\mathbf{z}}_1, \hat{\mathbf{z}}_2)$  employed by KNN. For OOD detection, this monotonic relationship implies that the two functions are equivalent at the decision level. Moreover, after normalization, directly using  $\cos (\mathbf{z}_1, \mathbf{z}_2)$  only requires dot product operations, significantly reducing computational overhead, especially when calculated in parallel.



Figure 7. The density of the obtained ID and OOD score with the proposed method (left) and KNN [31] (right). CIFAR-100 is ID, Near OOD is CIFAR-10, and SVHN is Far OOD.

OOD Score Density: Our Method vs. KNN. We further use CIFAR-100 as ID and plot the density of the OOD scores using a kernel density estimate (KDE) to analyze the impact of Near OOD (CIFAR-10) and Far OOD (SVHN). Besides, we provide a theoretical analysis comparing our method with KNN as an OOD score function for OOD detection in Section B of the Supplementary Material. We observe that the proposed method effectively captures the features of OOD samples. Specifically, in Figure 7, we make the following observations: (i) Thanks to the calibration of OOD scores by the OOD dictionary, the distribution of ID scores is more concentrated. (ii) In Figure 7(b), we observe that for the lefttail distribution, the intersection over union (IoU) between Far OOD and ID is smaller. As a result, the overall distribution of Far OOD scores shifts more to the left, achieving greater calibration gains.

### 8. Conclusion

We present a test-time OOD detection method called OODD that leverages a dynamic dictionary of OOD features to achieve high detection performance without fine-tuning or complex calibration processes. By integrating a priority queue and employing inlier and outlier sampling strategies, OODD effectively adapts to diverse OOD distributions. Our method not only outperforms the state-of-the-art OOD detection method but also demonstrates strong computational efficiency, as verified on the OpenOOD benchmark. These findings highlight OODD's potential to enhance model uncertainty awareness at test time—an essential capability for deploying robust models in open-world scenarios.

### Acknowledgement

This research is supported by Shanghai Artificial Intelligence Laboratory, National Natural Science Foundation of China under Grant (No.62106139) and State Key Laboratory of High Performance Computing, National University of DT (No.202401-11). The authors also gratefully acknowledge the insightful comments provided by the anonymous reviewers.

### References

- Anonymous. TULip: Test-time uncertainty estimation via linearization and weight perturbation. In Submitted to The Thirteenth International Conference on Learning Representations, 2024. under review. 1, 2, 5
- [2] Haoyue Bai, Gregory Canal, Xuefeng Du, Jeongyeol Kwon, Robert D Nowak, and Yixuan Li. Feed two birds with one scone: Exploiting wild data for both out-of-distribution generalization and detection. In *International Conference on Machine Learning*, pages 1454–1471. PMLR, 2023. 2
- [3] Julian Bitterwolf, Maximilian Müller, and Matthias Hein. In or out? fixing imagenet out-of-distribution detection evaluation. arXiv preprint arXiv:2306.00826, 2023. 4
- [4] Jiefeng Chen, Yixuan Li, Xi Wu, Yingyu Liang, and Somesh Jha. Atom: Robustifying out-of-distribution detection using outlier mining. In Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part III 21, pages 430–445. Springer, 2021. 1, 2
- [5] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer* vision and pattern recognition, pages 3606–3613, 2014. 4
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009. 4
- [7] Andrija Djurisic, Nebojsa Bozanic, Arjun Ashok, and Rosanne Liu. Extremely simple activation shaping for outof-distribution detection. *arXiv preprint arXiv:2209.09858*, 2022. 1, 5
- [8] Ke Fan, Tong Liu, Xingyu Qiu, Yikai Wang, Lian Huai, Zeyu Shangguan, Shuang Gou, Fengjian Liu, Yuqian Fu, Yanwei Fu, et al. Test-time linear out-of-distribution detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23752–23761, 2024. 1, 2, 5, 7
- [9] Camila González, Karol Gotkowski, Moritz Fuchs, Andreas Bucher, Armin Dadras, Ricarda Fischbach, Isabel Jasmin Kaltenborn, and Anirban Mukhopadhyay. Distance-based detection of out-of-distribution silent failures for covid-19 lung lesion segmentation. *Medical image analysis*, 82:102596, 2022. 1
- [10] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572, 2014. 1
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14, pages 630–645. Springer, 2016. 4
- [13] Dan Hendrycks and Kevin Gimpel. A baseline for detect-

ing misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016. 1

- [14] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:1812.04606*, 2018. 1, 2
- [15] Dan Hendrycks, Steven Basart, Mantas Mazeika, Andy Zou, Joe Kwon, Mohammadreza Mostajabi, Jacob Steinhardt, and Dawn Song. Scaling out-of-distribution detection for realworld settings. arXiv preprint arXiv:1911.11132, 2019. 4, 5
- [16] Rui Huang and Yixuan Li. Mos: Towards scaling out-ofdistribution detection for large semantic space. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8710–8719, 2021. 6
- [17] Xue Jiang, Feng Liu, Zhen Fang, Hong Chen, Tongliang Liu, Feng Zheng, and Bo Han. Negative label guided ood detection with pretrained vision-language models. *arXiv preprint arXiv:2403.20078*, 2024. 7
- [18] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 4
- [19] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015. 4
- [20] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [21] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 31, 2018. 1, 4, 5
- [22] Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. arXiv preprint arXiv:1706.02690, 2017. 4, 5
- [23] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. Advances in neural information processing systems, 33:21464–21475, 2020. 1, 4, 5
- [24] Yifei Ming, Ziyang Cai, Jiuxiang Gu, Yiyou Sun, Wei Li, and Yixuan Li. Delving into out-of-distribution detection with vision-language representations. Advances in neural information processing systems, 35:35087–35102, 2022. 7
- [25] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011. 4
- [26] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 427–436, 2015. 1
- [27] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 7
- [28] Sebastian Schmidt, Leonard Schenk, Leo Schwinn, and Stephan Günnemann. A unified approach towards active

learning and out-of-distribution detection. *arXiv preprint* arXiv:2405.11337, 2024. 2

- [29] Manli Shu, Weili Nie, De-An Huang, Zhiding Yu, Tom Goldstein, Anima Anandkumar, and Chaowei Xiao. Test-time prompt tuning for zero-shot generalization in vision-language models. *Advances in Neural Information Processing Systems*, 35:14274–14289, 2022. 2
- [30] Peng Sun, Bei Shi, Daiwei Yu, and Tao Lin. On the diversity and realism of distilled dataset: An efficient dataset distillation paradigm. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 9390– 9399, 2024. 3, 4
- [31] Yiyou Sun, Yifei Ming, Xiaojin Zhu, and Yixuan Li. Outof-distribution detection with deep nearest neighbors. In *International Conference on Machine Learning*, pages 20827– 20840. PMLR, 2022. 1, 2, 3, 5, 6, 7, 8
- [32] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8769–8778, 2018, 4
- [33] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Open-set recognition: A good closed-set classifier is all you need? arXiv preprint arXiv:2110.06207, 2021. 4
- [34] Haoqi Wang, Zhizhong Li, Litong Feng, and Wayne Zhang. Vim: Out-of-distribution with virtual-logit matching. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 4921–4930, 2022. 4
- [35] Qizhou Wang, Junjie Ye, Feng Liu, Quanyu Dai, Marcus Kalander, Tongliang Liu, Jianye Hao, and Bo Han. Outof-distribution detection with implicit outlier transformation. *arXiv preprint arXiv:2303.05033*, 2023. 1, 6
- [36] Puning Yang, Jian Liang, Jie Cao, and Ran He. Auto: Adaptive outlier optimization for online test-time ood detection. *arXiv preprint arXiv:2303.12267*, 2023. 1, 2
- [37] Jingyang Zhang, Nathan Inkawhich, Randolph Linderman, Yiran Chen, and Hai Li. Mixture outlier exposure: Towards out-of-distribution detection in fine-grained environments. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pages 5531–5540, 2023. 1, 2
- [38] Jingyang Zhang, Jingkang Yang, Pengyun Wang, Haoqi Wang, Yueqian Lin, Haoran Zhang, Yiyou Sun, Xuefeng Du, Kaiyang Zhou, Wayne Zhang, et al. Openood v1. 5: Enhanced benchmark for out-of-distribution detection. *arXiv* preprint arXiv:2306.09301, 2023. 2, 4, 7, 3
- [39] Qingyang Zhang, Qiuxuan Feng, Joey Tianyi Zhou, Yatao Bian, Qinghua Hu, and Changqing Zhang. The best of both worlds: On the dilemma of out-of-distribution detection. arXiv preprint arXiv:2410.11576, 2024. 2
- [40] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis* and machine intelligence, 40(6):1452–1464, 2017. 4
- [41] Jianing Zhu, Geng Yu, Jiangchao Yao, Tongliang Liu, Gang Niu, Masashi Sugiyama, and Bo Han. Diversified outlier exposure for out-of-distribution detection via informative extrapolation. arXiv preprint arXiv:2310.13923, 2023. 1, 2