

EAT: Towards Long-Tailed Out-of-Distribution Detection

Tong Wei*, Bo-Lin Wang, Min-Ling Zhang

School of Computer Science and Engineering, Southeast University, Nanjing 210096, China
 Key Laboratory of Computer Network and Information Integration (Southeast University), Ministry of Education, China
 {weit, wangbl, zhangml}@seu.edu.cn

Abstract

Despite recent advancements in out-of-distribution (OOD) detection, most current studies assume a class-balanced in-distribution training dataset, which is rarely the case in real-world scenarios. This paper addresses the challenging task of long-tailed OOD detection, where the in-distribution data follows a long-tailed class distribution. The main difficulty lies in distinguishing OOD data from samples belonging to the tail classes, as the ability of a classifier to detect OOD instances is not strongly correlated with its accuracy on the in-distribution classes. To overcome this issue, we propose two simple ideas: (1) Expanding the in-distribution class space by introducing multiple abstention classes. This approach allows us to build a detector with clear decision boundaries by training on OOD data using *virtual labels*. (2) Augmenting the context-limited tail classes by overlaying images onto the context-rich OOD data. This technique encourages the model to pay more attention to the discriminative features of the tail classes. We provide a clue for separating in-distribution and OOD data by analyzing gradient noise. Through extensive experiments, we demonstrate that our method outperforms the current state-of-the-art on various benchmark datasets. Moreover, our method can be used as an add-on for existing long-tail learning approaches, significantly enhancing their OOD detection performance. Code is available at: <https://github.com/Stomach-ache/Long-Tailed-OOD-Detection>.

Introduction

Deep neural networks (DNNs) can achieve high performance in various real-world applications by training on large-scale and well-annotated datasets. Most supervised learning literature makes a common assumption that the training and test data have the same distribution. However, DNNs in deployment often encounter data from an unknown distribution, and it has been shown that DNNs tend to produce wrong predictions on anonymous, or out-of-distribution (OOD) test data with high confidence (Hendrycks and Gimpel 2017; Liang, Li, and Srikant 2018; Hein, Andriushchenko, and Bitterwolf 2019), which can result in severe mistakes in practice.

Recently, OOD detection, which aims to reject OOD test data without classifying them as in-distribution labels, has

caught great attention. Existing state-of-the-art OOD detectors achieve huge success by maximizing the predictive uncertainty (Hendrycks, Mazeika, and Dietterich 2019; Meinke and Hein 2020), energy function (Liu et al. 2020), and abstention class confidence (Mohseni et al. 2020; Chen et al. 2021) for OOD data. However, these approaches assume that the in-distribution data is class-balanced, which is usually violated in real-world tasks (Van Horn and Perona 2017; Liu et al. 2019; Cui et al. 2019; Wei and Li 2020; Wei et al. 2021b; Wei and Gan 2023). In this paper, we consider that the in-distribution training data follows a long-tailed class distribution. Under this setup, directly combining existing OOD detectors with long-tailed learning methods still leads to unsatisfactory performance (Wang et al. 2022). So, a natural question is raised:

Is it possible to effectively distinguish OOD data from tail-class samples?

To answer this question, we propose a novel framework, EAT, which is composed of two key ingredients: (1) *dynamic virtual labels*, which expand the classification space with abstention classes for OOD data and are dynamically assigned by the model in the training process. EAT classifies OOD samples into abstention OOD classes rather than imposing uniform predictive probabilities over inlier classes such as in OE (Hendrycks, Mazeika, and Dietterich 2019), Energy (Liu et al. 2020), and PASCL (Wang et al. 2022). This step is critical because inherent similar OOD samples can be pushed closer if they are classified as an identical OOD class, and the decision boundary between inlier data and OOD data will be clearer. (2) *tail class augmentation*, which augments the tail-class images by pasting them onto the context-rich OOD images to force the model to focus on the foreground objects. Precisely, given an original image from the tail class, it is cropped in various sizes and pasted onto images from OOD data. Then, we can create tail-class images with more diverse contexts by changing the background. The generalization for tail classes can be significantly improved.

To further enhance the classification of inlier data, we propose a method that involves fine-tuning the classifiers exclusively using inlier data. This fine-tuning process employs a class-balanced loss function for a few iterations. Additionally, we illustrate that our method can be seamlessly integrated with existing long-tail learning approaches, leading to a significant improvement in their OOD detection performance.

*Corresponding author

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This is evident from the results presented in Table 6, where our method acts as a valuable plugin to boost the performance of these approaches. These findings contradict the argument put forth by previous work (Vaze et al. 2022) that a classifier performing well on in-distribution data would automatically excel as an OOD detector.

The key **contributions** of this paper are summarized as follows: (1) We tackle the challenging and under-explored problem of long-tailed OOD detection. This problem poses unique difficulties and requires innovative solutions. (2) We propose a novel approach to train OOD data using virtual labels, presenting an alternative to the outlier exposure method specifically designed for long-tailed data. Furthermore, we provide insights into the impact of virtual labels by examining gradient noise, deepening our understanding of their effectiveness. (3) Through extensive experiments conducted on various datasets, we empirically validate the effectiveness of our proposed method. Our results demonstrate an average boost of 2.0% AUROC and 2.9% inlier classification accuracy compared to the previous state-of-the-art method. (4) Our method serves as a versatile add-on for mainstream long-tailed learning methods, significantly enhancing their performance in detecting OOD samples. Importantly, our findings challenge the notion that a strong inlier classifier necessarily implies good OOD detection performance.

Related Work

OOD detection As a representative approach, Outlier Exposure (OE) proposes maximizing the OOD data’s predictive uncertainty as a complementary objective for the in-distribution classification loss. Further, Energy (Liu et al. 2020) improves OE by introducing the energy function as a regularization term and detects OOD samples according to their energy scores. Conversely, SOFL (Mohseni et al. 2020) and ATOM (Chen et al. 2021) attempt to classify OOD samples into abstention classes while in-distribution samples are classified into their true classes. Then, OOD data can be identified according to the model’s outputs on abstention classes. Although existing OOD detectors can achieve high performance, they are typically trained on class-balanced in-distribution datasets and cannot be directly applied to long-tailed tasks.

Long-tailed learning Existing approaches to long-tailed learning can be roughly categorized into three types by modifying: (1) the inputs to a model by re-balancing the training data (He and Garcia 2009; Liu et al. 2019; Zhou et al. 2020); (2) the outputs of a model, for example by posthoc adjustment of the classifier (Kang et al. 2020; Menon et al. 2021; Tang, Huang, and Zhang 2020) and (3) the internals of a model by modifying the loss function (Cui et al. 2019; Cao et al. 2019; Jamal et al. 2020; Ren et al. 2020). Recently, (Yang and Xu 2020) and (Wei et al. 2022) propose using OOD data to improve the performance of long-tailed learning. However, it is noted that these approaches are designed to boost the in-distribution classification performance and cannot be directly employed to detect OOD data.

Long-tailed OOD detection Recently, long-tailed OOD detection has received more and more attention, and several

approaches have been proposed to tackle this challenging problem. PASCL (Wang et al. 2022) optimizes a contrastive objective between tail class samples and OOD data to push each other away in the latent representation space, which can boost the performance of OOD detection. Further, it minimizes the logit adjustment loss to yield a class-balanced performance of inlier classification. HOD (Roy et al. 2022) studies a long-tail OOD detection problem in medical image analysis, which directly trains a binary classifier to discriminate in-distribution data and OOD data. However, HOD assumes that the OOD data is labeled, while we do not make this assumption and only leverage unlabeled OOD data to aid the detection performance. OLTR (Liu et al. 2019) formally studies the OOD detection task in long-tailed learning. It detects OOD inputs in the latent representation space according to the minimum distance between them and the centroids of in-distribution classes. Although OLTR outperforms several OOD detectors such as MSP (Hendrycks and Gimpel 2017), it is outperformed by the state-of-the-art OOD detection methods, suggesting that there remains room for improvement.

The Proposed Approach

Overview

We follow the popular training objective of existing state-of-the-art OOD detection methods, which train the model using both in-distribution data and unlabeled OOD data. Let \mathcal{D}_{in} and \mathcal{D}_{out} denote an in-distribution training set and an unlabeled OOD training set, respectively. Note that \mathcal{D}_{in} follows a long-tailed class distribution in our setup. The training loss function of many existing OOD detection methods (e.g., OE, EnergyOE, ATOM, and PASCL) is defined as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{in}} + \lambda \cdot \mathcal{L}_{\text{out}}, \quad (1)$$

where \mathcal{L}_{in} is the inlier classification loss, \mathcal{L}_{out} is the outlier detection loss, and λ is a trade-off hyperparameter. Typically, we choose to optimize the standard cross entropy loss (denoted by ℓ) for the inlier classification task:

$$\begin{aligned} \mathcal{L}_{\text{in}} &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{\text{in}}} [\ell(f(\mathbf{x}), y)] \\ &= \log[1 + \sum_{y' \neq y} e^{(f_{y'}(\mathbf{x}) - f_y(\mathbf{x}))}] \end{aligned} \quad (2)$$

Here, $f_y(\mathbf{x})$ represents the predicted logit corresponding to label y . For OOD detection, we propose using k abstention classes. The training outlier data is assigned to abstention classes by generating “virtual” labels by the model, and virtual labels may change through training iterations. With this, the training objective for outliers is defined as:

$$\begin{aligned} \mathcal{L}_{\text{out}} &= \mathbb{E}_{\tilde{\mathbf{x}} \sim \mathcal{D}_{\text{out}}} [\ell(f(\tilde{\mathbf{x}}), \tilde{y})] \\ &= \log[1 + \sum_{y' \neq \tilde{y}} e^{(f_{y'}(\tilde{\mathbf{x}}) - f_{\tilde{y}}(\tilde{\mathbf{x}}))}] \\ \text{s.t. } \tilde{y} &= \arg \max_{c \in [C+1, C+k]} f_c(\tilde{\mathbf{x}}) \end{aligned} \quad (3)$$

where \tilde{y} is the virtual label of outlier sample $\tilde{\mathbf{x}}$. Note that our treatment for training outlier data differs from existing methods, including OE, Energy, and PASCL. They attempt to

maximize the predictive uncertainties of outliers. We demonstrate that our approach achieves significantly better results in the experiments by introducing multiple abstention classes. The proposed approach is detailed below.

OOD Samples with Dynamic Virtual Labels

The approach of using abstention OOD classes is motivated by recent works (Abdelzad et al. 2019; Chen et al. 2021; Vernekar et al. 2019) which propose to add a single abstention class for all outlier data. Although this is shown to be effective compared to the outlier exposure method (Hendrycks, Mazeika, and Dietterich 2019), fitting a heterogeneous outlier set to a single class is challenging and problematic. One natural mitigation strategy here is to assign multiple abstention classes as possible outputs, which essentially turns the C -class classification into a $(C+k)$ -class classification problem. Here, we denote C as the number of inlier classes and k as the number of abstention classes added for outliers. Taking CIFAR100-LT as an example, if we use an additional $k = 30$ classes for fitting outliers, the number of neurons in the final fully-connected layer will be 130.

Ultimately, we want our model to classify unseen outliers in the test set into those k abstention classes. This can be achieved by encouraging the model to learn a structured decision boundary for the inliers *vs.* outliers. However, the ground-truth labels for training outlier data are not accessible. Thus, we propose generating virtual labels for the outliers so that the model learns to distinguish them from inliers. Towards this end, we take the predictions of the immediate model as the virtual labels at each training iteration, also known as self-labelling. The model is trained to predict virtual labels by minimizing the cross-entropy loss at the next iteration. The generation of virtual labels coincides with the self-training process, which is a popular framework in semi-supervised learning. At test time, the sum of probabilities for the k abstention classes indicating the OOD score is used. This is because the abstention classes are meaningless and virtual labels do not correspond to their ground-truth labels.

Mathematical Interpretation Exploring the reason behind OOD samples yielding higher scores than in-distribution samples is an intriguing endeavor. One way to comprehend the impact of virtual labels is through the lens of noise in loss gradients (Wei et al. 2021a). We define the trainable parameter of model f as $\theta \in \mathbb{R}^p$. By calculating the gradient of the loss function with respect to θ and updating the parameter accordingly, we gain insight into this phenomenon. Specifically, we represent the output probabilities for an in-distribution sample \mathbf{x} and an OOD sample $\tilde{\mathbf{x}}$ as $\mathbf{z} = \text{Softmax}(f(\mathbf{x}))$ and $\tilde{\mathbf{z}} = \text{Softmax}(f(\tilde{\mathbf{x}}))$ respectively.

Proposition 1. *For the cross-entropy loss, Eq. (3) induces gradient noise $\mathbf{g} = -\frac{\nabla_{\theta}\tilde{\mathbf{z}}_j}{\tilde{\mathbf{z}}_j}$ on $\nabla_{\theta}\ell(\mathbf{z}, y)$, s.t., $\mathbf{g} \in \mathbb{R}^p$, $j = \arg \max_{j \in [C+1, C+k]} \tilde{\mathbf{z}}_j$. While each OOD sample in OE (Hendrycks, Mazeika, and Dietterich 2019) induces gradient noise $\mathbf{g}' = -\frac{1}{C} \sum_{j=1}^C \frac{\nabla_{\theta}\tilde{\mathbf{z}}_j}{\tilde{\mathbf{z}}_j}$ on $\nabla_{\theta}\ell(\mathbf{z}, y)$, where $\frac{\cdot}{\tilde{\mathbf{z}}_j}$ denotes the element-wise division.*

Remark. The detailed proof for the following proposition can be found in the supplementary material. We first draw the

conclusion that our proposed virtual labeling induces gradient noise of $\mathbf{g} = -\frac{\nabla_{\theta}\tilde{\mathbf{z}}_j}{\tilde{\mathbf{z}}_j}$ where $j \in [C+1, C+k]$ is the virtual label for an OOD sample. On the contrary, previous method OE induces gradient noise of $\mathbf{g}' = -\frac{1}{C} \sum_{j=1}^C \frac{\nabla_{\theta}\tilde{\mathbf{z}}_j}{\tilde{\mathbf{z}}_j}$. Therefore, the main advantage of our approach yields gradient noise with dynamic direction depending on the virtual label of each OOD sample, which helps escape local minima during optimization. However, OE induces constant gradient noise so that the optimization of the model always follows the direction of gradient descent. Furthermore, the gradient noise induced by our approach helps the model to produce more conservative in-distribution class (i.e., $[1, C]$) predictions on OOD samples than OE. This is because of the nature of virtual labels which encourages the model to produce confident predictions on virtual classes, i.e., $[C+1, C+k]$.

Context-rich Tail Class Augmentation In our pursuit to enhance generalization, we go beyond the utilization of virtual labels to amplify the distinction between in-distribution and OOD samples. We additionally harness OOD samples to augment the tail classes, leading to improved performance. Our approach involves the implementation of an image-mixing data augmentation technique called CutMix (Yun et al. 2019), which enables us to generate training samples specifically tailored for the tail class. The core concept revolves around leveraging the context-rich nature of the head class and outlier images as backgrounds to create diverse and enriched tail samples. Given a tail-class image \mathbf{x}_f , we combine it with a randomly selected head-class or OOD image represented as \mathbf{x}_b . This merging operation is referred to as the CutMix operator and is defined as follows:

$$\mathbf{x}^{b \odot f} = \mathbf{M} \odot \mathbf{x}^b + (1 - \mathbf{M}) \odot \mathbf{x}^f \quad (4)$$

In this context, we designate \mathbf{x}^b as the background image and \mathbf{x}^f as the foreground image. A binary mask $\mathbf{M} \in \{0, 1\}^{W \times H}$ is employed to indicate the areas to preserve as background. Correspondingly, $(1 - \mathbf{M})$ selects the patch from the foreground image to be pasted onto the background image. Here, 1 represents a matrix filled with ones, and \odot denotes element-wise multiplication. In order to address the limited availability of data for tail classes, we assume that the composite image $\mathbf{x}^b \odot \mathbf{x}^f$ carries the same label as the foreground image \mathbf{x}^f , i.e., $y^{b \odot f} = y^f$. However, it is important to note that this approach can introduce label noise during training. Therefore, we assign lower sample weights to the generated tail-class images to mitigate the adverse impact.

Tailored CutMix offers two notable advantages for both outlier detection and inlier classification. Firstly, by using diverse OOD images as backgrounds, the model is encouraged to differentiate between tail-class images and OOD data based on foreground objects rather than image backgrounds. This aids in enhancing the model's ability to identify and distinguish outliers effectively. Secondly, the inclusion of head-class and OOD data through mixing increases the frequency of tail classes, leading to a more balanced training set. This improved class balance contributes to enhanced generalization capabilities.

It is worth noting that the study conducted by (Park et al. 2022) also incorporates CutMix to generate tail-class samples.

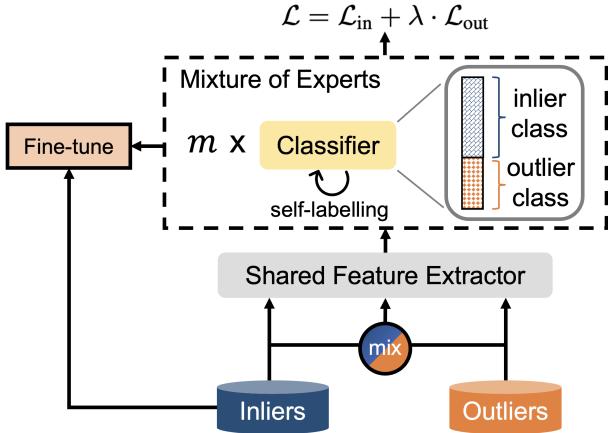


Figure 1: Overview of EAT framework.

However, their approach differs in that they sample image pairs from the original long-tailed data distribution and a tail-class-weighted distribution. Furthermore, their study focuses on improving inlier prediction accuracy rather than OOD detection. As far as our knowledge extends, we are the first to adapt CutMix specifically for long-tailed OOD detection, distinguishing our work in this area.

Improving OOD Separation

To amplify both outlier detection and inlier classification performance, we employ a mixture of experts by integrating multiple classifiers that share a common feature extractor. By training an ensemble of m members with random initializations, we optimize the sum of loss functions for these classifiers, aiming to achieve superior results.

$$\mathcal{L}_{\text{total}} = \sum_{i=1}^m (\mathcal{L}_{\text{in}}^{(i)} + \lambda \cdot \mathcal{L}_{\text{out}}^{(i)}) \quad (5)$$

Given an input \mathbf{x} at test time, we use the average predictions of ensemble members as the OOD score:

$$G(\mathbf{x}) = \frac{1}{m} \sum_{i=1}^m \sum_{j=C+1}^{C+k} z_j^{(i)}, \quad (6)$$

where $\mathbf{z}^{(i)} = \text{Softmax}(f^{(i)}(\mathbf{x}))$. We choose $m = 3$ in our experiments. If \mathbf{x} is not deemed as an OOD input, the prediction will be $\arg \max_{1 \leq c \leq C} \frac{1}{m} \sum_{i=1}^m z_j^{(i)}$.

It is important to note that the performance improvement achieved by deep ensembles relies on the diversity introduced through random initialization of network parameters. In our specific setup, since we employ a shared feature extractor, random initialization is applied solely to the parameters of the last layer. To enhance diversity further, we train ensemble models with virtual labels generated by each classifier. This means that the virtual label for a given sample \mathbf{x} is obtained by selecting $\tilde{y} = \arg \max_{c \in [C+1, C+k]} f_c^{(i)}(\mathbf{x})$ for the i -th classifier. The overall framework of our approach is depicted in Figure 1.

Model Fine-tuning

After training a multi-branch model, we proceed to fine-tune the classifiers using training inlier data, aiming to enhance the classification performance. It is worth mentioning that the two-stage approach, involving representation learning followed by classifier learning, is commonly employed in the field of long-tailed learning. Prominent examples include decoupling (Kang et al. 2020), BBN (Zhou et al. 2020), and MisLAS (Zhong et al. 2021). In this article, we implicitly explore the advantages of utilizing OOD data for representation learning by optimizing supervised objectives during model training. Furthermore, we keep the feature extractor fixed and refine the classifiers for a few iterations to improve inlier classification performance. During the fine-tuning process, we employ the logits adjustment (LA) loss (Menon et al. 2021) to guide the training.

$$\ell_{\text{LA}}(y, f(\mathbf{x})) = \log[1 + \sum_{y' \neq y} e^{\Delta_{yy'}} \cdot e^{(f_{y'}(\mathbf{x}) - f_y(\mathbf{x}))}] \quad (7)$$

Here, the pairwise label margins $\Delta_{yy'} = \log \frac{\pi_y}{\pi_{y'}}$ represents the desired gap between predictive confidence for y and y' depending on the number of each class. π_y denotes the class prior of class y in the training inlier data. We will empirically show that fine-tuning can introduce not only large improvements for the inlier classification task but also boost the performance of OOD detection.

Experiments

Experiment Settings

We verify our approach on commonly used datasets in comparison with the existing state-of-the-art. CIFAR10-LT, CIFAR100-LT (Cao et al. 2019), and ImageNet-LT (Liu et al. 2019) are used as in-distribution training sets (\mathcal{D}_{in}). The standard CIFAR10, CIFAR100, and ImageNet test sets are used as in-distribution test sets ($\mathcal{D}_{\text{in}}^{\text{test}}$). Following (Wang et al. 2022), we set the default imbalance ratio to 100 for CIFAR10-LT and CIFAR100-LT during training. For evaluation measures, we mainly use **AUROC**, **AUPR**, **FPR95**, **ACC@FPRn**, which is ACC when $n = 0$, following (Hendrycks, Mazeika, and Dietterich 2019; Mohseni et al. 2020; Yang et al. 2021; Wang et al. 2022).

OOD datasets for CIFAR-LT We employ 300 thousand samples from TinyImages80M (Torralba, Fergus, and Freeman 2008) as the OOD training images for CIFAT10-LT and CIFAR100-LT following (Hendrycks, Mazeika, and Dietterich 2019; Wang et al. 2022). Of those, 80 Million Tiny Images is a large-scale, diverse dataset of 32×32 natural images. The 300 thousand samples are selected from the 80 Million Tiny Images by (Hendrycks, Mazeika, and Dietterich 2019), not intersected with the CIFAR datasets. For OOD test data, we use Textures (Cimpoi et al. 2014), SVHN (Netzer et al. 2011), Tiny ImageNet (Le and Yang 2015), LSUN (Yu et al. 2015), and Places365 (Zhou et al. 2017) as $\mathcal{D}_{\text{out}}^{\text{test}}$. We use CIFAR-100 as a $\mathcal{D}_{\text{out}}^{\text{test}}$ for CIFAR10-LT and vice-versa.

OOD datasets for ImageNet-LT We use a specifically designed \mathcal{D}_{out} called ImageNet-Extra following (Wang et al. 2022). ImageNet-Extra has 517,711 images belonging to 500

$\mathcal{D}_{\text{out}}^{\text{test}}$	Method	AUROC(\uparrow)	AUPR(\uparrow)	FPR95(\downarrow)
Texture	OE	92.59	83.32	25.10
	PASCL	93.16	84.80	23.26
	Ours	95.44	92.28	21.50
SVHN	OE	95.10	97.14	16.15
	PASCL	96.63	98.06	12.18
	Ours	97.92	99.06	9.87
CIFAR100	OE	83.40	80.93	56.96
	PASCL	84.43	82.99	57.27
	Ours	85.93	86.10	54.13
Tiny ImageNet	OE	86.14	79.33	47.78
	PASCL	87.14	81.54	47.69
	Ours	89.11	85.43	41.75
LSUN	OE	91.35	87.62	27.86
	PASCL	93.17	91.76	26.40
	Ours	95.13	94.12	19.72
Places365	OE	90.07	95.15	34.04
	PASCL	91.43	96.28	33.40
	Ours	93.68	97.42	26.03
Average	OE	89.77	87.25	34.65
	PASCL	90.99	89.24	33.36
	Ours	92.87	92.40	28.83

(a) OOD detection results.

Method	ACC@FPR n (\uparrow)			
	0	0.001	0.01	0.1
OE	73.54	73.90	74.46	78.88
PASCL	77.08	77.13	77.64	81.96
Ours	81.31	81.36	81.81	84.40

(b) In-distribution classification results in terms of ACC@FPR n .

Method	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)	ACC (\uparrow)
ST (MSP)	72.28	70.27	66.07	72.34
OECC	87.28	86.29	45.24	60.16
EnergyOE	89.31	88.92	40.88	74.68
OE	89.77	87.25	34.65	73.84
PASCL	<u>90.99</u>	<u>89.24</u>	<u>33.36</u>	<u>77.08</u>
Ours	92.87	92.40	28.83	81.31

(c) Comparison with other methods.

Table 1: Results on CIFAR10-LT using ResNet18. The best results are shown in bold. Mean over six random runs are reported. “Average” means the results averaged across six different $\mathcal{D}_{\text{out}}^{\text{test}}$ sets.

classes randomly sampled from ImageNet-22k (Deng et al. 2009), but not overlapping with the 1,000 in-distribution classes in ImageNet-LT. For $\mathcal{D}_{\text{out}}^{\text{test}}$, we use ImageNet-1k-OOD constructed by (Wang et al. 2022), which has 50,000 OOD test images from 1,000 classes randomly selected from ImageNet-22k (with 50 images in each class). Considering the fairness of OOD detection, it has the same size as the in-distribution test set. The 1,000 classes in ImageNet-1k-OOD

are not intersecting either the 1,000 in-distribution classes in ImageNet-LT or the 500 OOD training classes in ImageNet-Extra. To ensure the rigor of the experiment, ImageNet-LT $\mathcal{D}_{\text{train}}^{\text{in}}$, ImageNet-Extra $\mathcal{D}_{\text{train}}^{\text{out}}$, ImageNet-1k-OOD $\mathcal{D}_{\text{test}}^{\text{out}}$, and ImageNet $\mathcal{D}_{\text{test}}^{\text{in}}$ are orthogonal.

Model Configuration The current best long-tail OOD detection method is PASCL, and before that it is OE, so we mainly compare the experimental results with these two baseline methods. For experiments on CIFAR10 and CIFAR100, we use the ResNet18 (He et al. 2016) following (Yang et al. 2021). For experiments on CIFAR10-LT and CIFAR100-LT, we train the model for 180 epochs using Adam (Kingma and Ba 2014) optimizer with initial learning rate 1×10^{-3} and batch size 128. We decay the learning rate to 0 using a cosine annealing learning rate scheduler (Loshchilov and Hutter 2016). For fine-tuning, we fine-tune the classifier and BN layers for 10 epochs using Adam optimizer with an initial learning rate 5×10^{-4} . For experiments on ImageNet-LT, we follow the settings in (Wang et al. 2021) and use ResNet50 (He et al. 2016). We train the main branch for 60 epochs using SGD optimizer with an initial learning rate of 0.1 and batch size of 64. We fine-tune the classifier and BN layers for the 1 epoch using SGD optimizer with an initial learning rate of 0.01. In all experiments, we set $\lambda = 0.05$, and the weights for generated tail class samples are set to 0.05 for EAT. For the number of abstention classes, we set $k = 3$ on CIFAR10-LT, $k = 30$ on CIFAR100-LT and ImageNet-LT. For other hyper-parameters in the baseline methods, we use the suggested values in their original papers.

Main Results

Table 1, Table 2, and Table 4 report the results for CIFAR10-LT, CIFAR100-LT, and ImageNet-LT datasets, respectively. For fair comparisons, results of existing methods are directly borrowed from (Wang et al. 2022). There are three sub-tables in Table 1 and Table 2: since performance measures may differ across different $\mathcal{D}_{\text{out}}^{\text{test}}$ datasets, we report AUROC, AUPR, FPR95, and ACC95 on each $\mathcal{D}_{\text{out}}^{\text{test}}$ as well as the average values across six $\mathcal{D}_{\text{out}}^{\text{test}}$ datasets in sub-table (a). In sub-table (b), we compare ACC@FPR n with various n values that are independent of $\mathcal{D}_{\text{out}}^{\text{test}}$. Finally, we put together four main performance measures in terms of both outlier detection and inlier classification in sub-table (c).

From the results, we can see that EAT significantly outperforms OE, PASCL, and other baselines. For instance, on CIFAR10-LT, EAT achieves 1.88% AUROC, 3.16% AUPR, 4.53% FPR95, 0.73% ACC95, and 4.23% in-distribution accuracy improvement than PASCL on average. Likewise, on CIFAR100-LT, our approach achieves 2.13% higher AUROC, 3.69% higher AUPR, 3.43% lower FPR95, and 3.13% higher classification accuracy than PASCL on average.

On ImageNet-LT, our approach achieves the best results in seven cases, while the previous state-of-the-art method PASCL performs the best in only one case. Compared with OE, our approach achieves 3.51% higher AUROC, 0.96% higher AUPR, and 9.19% higher in-distribution accuracy.

Improvements on head and tail classes In Table 3, we show the improvements of our method over OE on head and

$\mathcal{D}_{\text{out}}^{\text{test}}$	Method	AUROC(↑)	AUPR(↑)	FPR95(↓)
Texture	OE	76.71	58.79	68.28
	PASCL	76.01	58.12	67.43
	Ours	80.27	71.76	67.53
SVHN	OE	77.61	86.82	58.04
	PASCL	80.19	88.49	53.45
	Ours	83.11	89.71	47.78
CIFAR10	OE	62.23	57.57	80.64
	PASCL	62.33	57.14	79.55
	Ours	61.62	55.30	77.97
Tiny ImageNet	OE	68.04	51.66	76.66
	PASCL	68.20	51.53	76.11
	Ours	68.34	52.79	74.89
LSUN	OE	77.10	61.42	63.98
	PASCL	77.19	61.27	63.31
	Ours	81.09	67.46	55.02
Places365	OE	75.80	86.68	65.72
	PASCL	76.02	86.52	64.81
	Ours	78.28	88.20	60.85
Average	OE	72.91	67.16	68.89
	PASCL	73.32	67.18	67.44
	Ours	75.45	70.87	64.01

(a) OOD detection results.

Method	ACC@FPRn (↑)			
	0	0.001	0.01	0.1
OE	39.04	39.07	39.38	42.40
PASCL	43.10	43.12	43.39	46.14
Ours	46.23	46.24	46.38	48.39

(b) In-distribution classification results in terms of ACC95.

Method	AUROC (↑)	AUPR (↑)	FPR95 (↓)	ACC (↑)
ST (MSP)	61.00	57.54	82.01	40.97
OECC	70.38	66.87	73.15	32.93
EnergyOE	71.10	67.23	71.78	39.05
OE	72.91	67.16	68.89	39.04
PASCL	73.32	67.18	67.44	43.10
Ours	75.45	70.87	64.01	46.23

(c) Comparison with other methods.

Table 2: Results on CIFAR100-LT using ResNet18. The best results are shown in bold. Mean over six random runs are reported. “Average” means the results averaged across six different $\mathcal{D}_{\text{out}}^{\text{test}}$ sets.

tail in-distribution classes. As we can see, our approach can substantially benefit both the head and tail classes. Compared with PASCL, it is highly biased towards the tail class and the improvement on head classes is marginal, our method achieves a good balance.

Why our method achieves low ACC@TPRn? We show the failure cases. Table 4 shows several cases where the base-

Method	ACC (↑)	
	Head classes	Tail classes
OE	54.29	20.90
PASCL	54.73 (+0.44)	36.26 (+15.36)
Ours	59.46 (+5.17)	34.12 (+13.22)

Table 3: Results on ImageNet-LT.

line is better than our approach concerning ACC@TPRn. We empirically find this is due to more in-distribution samples preserved by our approach than the baselines when a percentage of OOD samples have been successfully detected. As a result, the classification accuracy of the remaining samples may be lower even though our approach correctly classifies more in-distribution samples than baselines. We provide more statistics in the supplementary.

Ablation Study

How do the key components of EAT affect the performance? In Table 5, we study the effects of the four critical components in our EAT approach: (1) virtual labels, (2) classifier fine-tuning, (3) CutMix, and (4) the mixture of experts (MoE), on CIFAR10-LT and CIFAR100-LT datasets. Since the performance of most approaches fluctuates on SVHN, we choose SVHN as $\mathcal{D}_{\text{out}}^{\text{test}}$.

First, on both CIFAR10-LT and CIFAR100-LT datasets, employing the virtual label strategy for OOD data significantly improves OOD detection performance. It improves the AUROC, AUPR, and FPR95 by an average margin of 10%. Note that, although without using virtual labels achieves higher results of ACC95, it is because more in-distribution samples are incorrectly deemed as OOD by the model.

Second, fine-tuning improves the inlier classification while maintaining a competitive OOD detection performance. For instance, the ACC@FPR increases by about 3% on average. Moreover, since we fine-tune the classifiers for only one iteration, it does not introduce much extra computational cost.

Third, we study the role of CutMix. We find it beneficial to use context-rich images as backgrounds to improve the performance of OOD detection and inlier classification. For OOD detection, it improves AUROC and FPR95 by about 4% and 8% on CIFAR100-LT, respectively. For inlier classification, the ACC@FPR is improved by an average of 2% on CIFAR100-LT. Notably, we observe different results on CIFAR100-LT and CIFAR10-LT concerning ACC@FPRn, which is related to the in-distribution accuracy. This may be related to the value of k because a larger k means that more abstention class heads need to be learned, resulting in a little effect on ID accuracy. Furthermore, Figure 2 illustrates the distribution of OOD scores on the other two OOD datasets, and our model distinguishes OOD data from in-distribution data with a clear decision boundary.

A Good Closed-Set Classifier is All You Need? An interesting finding from previous work (Vaze et al. 2022) suggests that utilizing the maximum logit score rule with a highly

Method	AUROC (\uparrow)	AUPR (\uparrow)	FPR@TPR n (\downarrow)				ACC@TPR n (\uparrow)				ACC@FPR n (\uparrow)			
			0.98	0.95	0.90	0.80	0.98	0.95	0.90	0.80	0	0.001	0.01	0.1
ST (MSP)	53.81	51.63	95.38	90.15	83.52	72.97	96.67	92.61	87.43	77.52	39.65	39.68	40.00	43.18
OECC	63.07	63.05	93.15	86.90	78.79	65.23	94.25	88.23	80.12	68.36	38.25	38.28	38.56	41.47
EnergyOE	64.76	64.77	<u>94.15</u>	87.72	78.36	63.71	80.18	74.38	67.65	59.68	38.50	38.52	38.72	40.99
OE	66.33	68.29	95.11	88.22	78.68	65.28	95.46	88.22	78.68	65.28	37.60	37.62	37.79	40.00
PASCL	<u>68.00</u>	70.15	94.38	<u>87.53</u>	<u>78.12</u>	62.48	<u>95.69</u>	<u>89.55</u>	<u>80.88</u>	69.60	45.49	<u>45.51</u>	<u>45.62</u>	47.49
Ours	69.84	<u>69.25</u>	94.34	87.63	77.30	57.81	83.22	77.80	70.84	61.49	46.79	46.79	46.83	48.30

Table 4: Results on ImageNet-LT. The best and second-best are bolded and underlined, respectively.

D_{in}	Virtual label	Fine-tuning	CutMix	MoE	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)	ACC95 (\uparrow)	ACC@FPR n (\uparrow)				
									0	0.001	0.01	0.1	
CIFAR10-LT	X	X	X	***	84.91	91.75	47.46	92.91	77.48	77.52	77.92	81.65	
					96.10	98.02	16.87	83.53	78.49	78.54	79.00	81.85	
					96.62	98.22	14.92	84.79	79.63	79.69	80.19	83.47	
					*	97.08	98.70	13.86	84.09	80.38	80.41	80.72	83.10
					**	97.62	98.93	11.43	83.87	80.49	80.53	80.92	83.49
	EAT				97.92	99.06	9.87	84.39	81.31	81.36	81.81	84.40	
CIFAR100-LT	X	X	X	***	74.04	84.99	63.58	73.98	46.68	46.70	46.91	49.50	
					81.77	89.67	54.53	64.47	43.34	43.36	43.59	46.00	
					79.52	88.11	55.89	64.30	43.93	43.94	44.10	46.24	
					*	80.70	88.27	52.86	64.30	45.82	45.84	45.97	47.94
					**	82.13	89.01	50.51	63.18	46.32	46.32	46.48	48.57
	EAT				83.11	89.71	47.78	61.67	46.23	46.24	46.38	48.39	

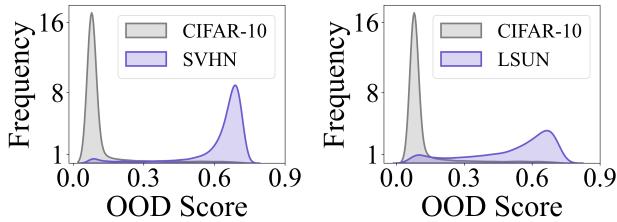
Table 5: The impact of key ingredients for EAT. Experiments are conducted on CIFAR10-LT and CIFAR100-LT ($\rho = 100$). SVHN is used as $\mathcal{D}_{out}^{\text{test}}$. The number of * denotes the ensemble size.

Figure 2: Distribution of OOD scores from our model. The CIFAR10 is used as the in-distribution dataset, and the other two are OOD datasets. It shows that both in-distribution data and OOD data naturally form smooth distributions.

accurate in-distribution classifier can outperform many well-designed OOD detectors. However, we sought to investigate whether this holds true in the context of long-tailed tasks. To validate this, we conducted experiments involving two sophisticated long-tail learning methods, namely RIDE (Wang et al. 2021) and GLMC (Du et al. 2023). We find that their OOD detection performance lagged significantly behind that of our proposed method, providing evidence that a strong classifier alone is insufficient for effective long-tailed OOD detection. Furthermore, when we combined our method with RIDE and GLMC, as shown in Table 6, EAT achieves a substantial improvement with minimal sacrifice.

Method	ACC (\uparrow)	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)
RIDE	48.49	66.18	62.13	77.73
RIDE+Ours	47.94	72.41	69.37	71.99
GLMC	54.51	65.01	62.01	79.65
GLMC+Ours	52.30	73.07	67.14	69.31

Table 6: Combining with other methods. The experiment is conducted on CIFAR-100 (in-distribution) dataset and six OOD datasets.

Conclusion

This paper proposes a novel framework (EAT) to tackle the long-tailed OOD detection problem. Towards this end, EAT presents several general techniques that can easily be applied to mainstream OOD detectors and long-tail learning methods. First, the abstention OOD classes can be used as an alternative to the outlier exposure method. Second, tail-class augmentation can be employed as a universal add-on for existing methods. Third, the classifier ensembling technique can further boost the performance without introducing much additional computational cost. Finally, we evaluate the proposed method on many commonly used datasets, showing that it consistently outperforms the existing state-of-the-art.

Acknowledgments

This research was supported by the National Science Foundation of China (62206049, 62176055). This research work is supported by the Big Data Computing Center of Southeast University.

References

Abdelzad, V.; Czarnecki, K.; Salay, R.; Denouden, T.; Vernekar, S.; and Phan, B. 2019. Detecting out-of-distribution inputs in deep neural networks using an early-layer output. *arXiv preprint arXiv:1910.10307*.

Cao, K.; Wei, C.; Gaidon, A.; Arechiga, N.; and Ma, T. 2019. Learning imbalanced datasets with label-distribution-aware margin loss. In *NeurIPS*, 1567–1578.

Chen, J.; Li, Y.; Wu, X.; Liang, Y.; and Jha, S. 2021. ATOM: Robustifying out-of-distribution detection using outlier mining. In *ECML*, 430–445.

Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing textures in the wild. In *CVPR*, 3606–3613.

Cui, Y.; Jia, M.; Lin, T.-Y.; Song, Y.; and Belongie, S. 2019. Class-balanced loss based on effective number of samples. In *CVPR*, 9268–9277.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *CVPR*, 248–255.

Du, F.; Yang, P.; Jia, Q.; Nan, F.; Chen, X.; and Yang, Y. 2023. Global and Local Mixture Consistency Cumulative Learning for Long-tailed Visual Recognitions. In *CVPR*.

He, H.; and Garcia, E. A. 2009. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9): 1263–1284.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.

Hein, M.; Andriushchenko, M.; and Bitterwolf, J. 2019. Why ReLU networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In *CVPR*, 41–50.

Hendrycks, D.; and Gimpel, K. 2017. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *ICLR*.

Hendrycks, D.; Mazeika, M.; and Dietterich, T. 2019. Deep Anomaly Detection with Outlier Exposure. In *ICLR*.

Jamal, M. A.; Brown, M.; Yang, M.-H.; Wang, L.; and Gong, B. 2020. Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. In *CVPR*, 7610–7619.

Kang, B.; Xie, S.; Rohrbach, M.; Yan, Z.; Gordo, A.; Feng, J.; and Kalantidis, Y. 2020. Decoupling Representation and Classifier for Long-Tailed Recognition. In *ICLR*.

Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Le, Y.; and Yang, X. 2015. Tiny ImageNet Visual Recognition Challenge.

Liang, S.; Li, Y.; and Srikant, R. 2018. Enhancing The Reliability of Out-of-distribution Image Detection in Neural Networks. In *ICLR*.

Liu, W.; Wang, X.; Owens, J.; and Li, Y. 2020. Energy-based Out-of-distribution Detection. In *NeurIPS*.

Liu, Z.; Miao, Z.; Zhan, X.; Wang, J.; Gong, B.; and Yu, S. X. 2019. Large-scale long-tailed recognition in an open world. In *CVPR*, 2537–2546.

Loshchilov, I.; and Hutter, F. 2016. SGDR: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*.

Meinke, A.; and Hein, M. 2020. Towards neural networks that provably know when they don't know. In *ICLR*.

Menon, A. K.; Jayasumana, S.; Rawat, A. S.; Jain, H.; Veit, A.; and Kumar, S. 2021. Long-tail learning via logit adjustment. In *ICLR*.

Mohseni, S.; Pitale, M.; Yadawa, J.; and Wang, Z. 2020. Self-supervised learning for generalizable out-of-distribution detection. In *AAAI*, 5216–5223.

Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. Reading digits in natural images with unsupervised feature learning.

Park, S.; Hong, Y.; Heo, B.; Yun, S.; and Choi, J. Y. 2022. The Majority Can Help The Minority: Context-rich Minority Oversampling for Long-tailed Classification. In *CVPR*.

Ren, J.; Yu, C.; sheng, s.; Ma, X.; Zhao, H.; Yi, S.; and Li, h. 2020. Balanced Meta-Softmax for Long-Tailed Visual Recognition. In *NeurIPS*, 4175–4186.

Roy, A. G.; Ren, J.; Azizi, S.; Loh, A.; Natarajan, V.; Mustafa, B.; Pawlowski, N.; Freyberg, J.; Liu, Y.; Beaver, Z.; et al. 2022. Does your dermatology classifier know what it doesn't know? Detecting the long-tail of unseen conditions. *Medical Image Analysis*, 75: 102274.

Tang, K.; Huang, J.; and Zhang, H. 2020. Long-tailed classification by keeping the good and removing the bad momentum causal effect. *NeurIPS*, 1513–1524.

Torralba, A.; Fergus, R.; and Freeman, W. T. 2008. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE TPAMI*, 30(11): 1958–1970.

Van Horn, G.; and Perona, P. 2017. The devil is in the tails: Fine-grained classification in the wild. *arXiv preprint arXiv:1709.01450*.

Vaze, S.; Han, K.; Vedaldi, A.; and Zisserman, A. 2022. Open-Set Recognition: a Good Closed-Set Classifier is All You Need? In *ICLR*.

Vernekar, S.; Gaurav, A.; Abdelzad, V.; Denouden, T.; Salay, R.; and Czarnecki, K. 2019. Out-of-distribution detection in classifiers via generation. *arXiv preprint arXiv:1910.04241*.

Wang, H.; Zhang, A.; Zhu, Y.; Zheng, S.; Li, M.; Smola, A. J.; and Wang, Z. 2022. Partial and Asymmetric Contrastive Learning for Out-of-Distribution Detection in Long-Tailed Recognition. In *ICML*, 23446–23458.

Wang, X.; Lian, L.; Miao, Z.; Liu, Z.; and Yu, S. 2021. Long-tailed Recognition by Routing Diverse Distribution-Aware Experts. In *ICLR*.

Wei, H.; Tao, L.; Xie, R.; and An, B. 2021a. Open-set Label Noise Can Improve Robustness Against Inherent Label Noise. In *NeurIPS*.

Wei, H.; Tao, L.; Xie, R.; Feng, L.; and An, B. 2022. Open-Sampling: Exploring Out-of-Distribution data for Re-balancing Long-tailed datasets. In *ICML*, 23615–23630.

Wei, T.; and Gan, K. 2023. Towards Realistic Long-Tailed Semi-Supervised Learning: Consistency is All You Need. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

Wei, T.; and Li, Y.-F. 2020. Does tail label help for large-scale multi-label learning? *IEEE Transactions on Neural Networks and Learning Systems*, 31(7): 2315–2324.

Wei, T.; Shi, J.-X.; Tu, W.-W.; and Li, Y.-F. 2021b. Robust Long-Tailed Learning under Label Noise. *arXiv preprint arXiv:2108.11569*.

Yang, J.; Wang, H.; Feng, L.; Yan, X.; Zheng, H.; Zhang, W.; and Liu, Z. 2021. Semantically coherent out-of-distribution detection. In *ICCV*, 8301–8309.

Yang, Y.; and Xu, Z. 2020. Rethinking the Value of Labels for Improving Class-Imbalanced Learning. In *NeurIPS*.

Yu, F.; Seff, A.; Zhang, Y.; Song, S.; Funkhouser, T.; and Xiao, J. 2015. LSUN: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*.

Yun, S.; Han, D.; Chun, S.; Oh, S. J.; Yoo, Y.; and Choe, J. 2019. CutMix: Regularization Strategy to Train Strong Classifiers With Localizable Features. In *ICCV*, 6022–6031.

Zhong, Z.; Cui, J.; Liu, S.; and Jia, J. 2021. Improving Calibration for Long-Tailed Recognition. In *CVPR*, 16489–16498.

Zhou, B.; Cui, Q.; Wei, X.-S.; and Chen, Z.-M. 2020. BBN: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In *CVPR*, 9719–9728.

Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; and Torralba, A. 2017. Places: A 10 million image database for scene recognition. *IEEE TPAMI*, 40(6): 1452–1464.