
Partial and Asymmetric Contrastive Learning for Out-of-Distribution Detection in Long-Tailed Recognition

Haotao Wang¹ Aston Zhang² Yi Zhu² Shuai Zheng² Mu Li² Alex Smola² Zhangyang Wang¹

Abstract

Existing out-of-distribution (OOD) detection methods are typically benchmarked on training sets with balanced class distributions. However, in real-world applications, it is common for the training sets to have long-tailed distributions. In this work, we first demonstrate that existing OOD detection methods commonly suffer from significant performance degradation when the training set is long-tail distributed. Through analysis, we posit that this is because the models struggle to distinguish the minority tail-class in-distribution samples, from the true OOD samples, making the tail classes more prone to be falsely detected as OOD. To solve this problem, we propose *Partial and Asymmetric Supervised Contrastive Learning* (PASCL), which explicitly encourages the model to distinguish between tail-class in-distribution samples and OOD samples. To further boost in-distribution classification accuracy, we propose Auxiliary Branch Finetuning, which uses two separate branches of BN and classification layers for anomaly detection and in-distribution classification, respectively. The intuition is that in-distribution and OOD anomaly data have different underlying distributions. Our method outperforms previous state-of-the-art method by 1.29%, 1.45%, 0.69% anomaly detection false positive rate (FPR) and 3.24%, 4.06%, 7.89% in-distribution classification accuracy on CIFAR10-LT, CIFAR100-LT, and ImageNet-LT, respectively. Code and pre-trained models are available at <https://github.com/amazon-research/>

Work done during the first author’s internship at Amazon Web Services. ¹University of Texas at Austin, Austin, USA ²Amazon Web Services, Santa Clara, USA. Correspondence to: Haotao Wang <htwang@utexas.edu>, Aston Zhang <astonz@amazon.com>, Zhangyang Wang <atlaswang@utexas.edu>.

Proceedings of the 39th International Conference on Machine Learning, Baltimore, Maryland, USA, PMLR 162, 2022. Copyright 2022 by the author(s).

`long-tailed-ood-detection`.

1. Introduction

Deep neural networks (DNNs) are widely known to be overconfident on what they do not know. Given a test sample that does not belong to any training class, DNNs tend to recklessly predict this out-of-distribution (OOD) test sample as one of the training classes with high confidence, which is doomed to be wrong (Hendrycks & Gimpel, 2017; Hein et al., 2019). Numerous OOD detection methods have been proposed to solve this problem with promising results (Hendrycks et al., 2019; Mohseni et al., 2020; Liu et al., 2020). The goal of OOD detection is two-fold: (i) to distinguish OOD samples from in-distribution samples (i.e., *anomaly detection*), and (ii) to still achieve high accuracy on the *in-distribution sample classification*.

Existing OOD detection methods are typically evaluated on balanced training sets. However, the training sets in real-world applications often follow a long-tailed distribution (Van Horn & Perona, 2017; Liu et al., 2019; Cui et al., 2019; Menon et al., 2021). In view of this, we benchmark existing OOD detection methods (including both widely used classic methods and recently published ones) on the long-tailed training sets. We observe significant performance drop in terms of both anomaly detection and in-distribution classification, when compared with the results obtained on the balanced training sets (Table 1). This shows the challenge of OOD detection on long-tailed recognition (LTR) task. Moreover, we show that it cannot be easily fixed by naive solutions such as combining state-of-the-art OOD detection and long-tailed recognition methods (Table 2). Through feature visualization (Fig. 1), we further unveil the underlying cause for such performance drop: The model struggles to distinguish tail-class in-distribution samples from OOD samples. In contrast, the head-class in-distribution samples can be easily distinguished from OOD samples.

Motivated by such observations, we propose *Partial and Asymmetric Supervised Contrastive Learning* (PASCL) as the solution for OOD detection in long-tailed recognition. The key intuition is to push tail-class in-distribution samples away from OOD samples in the feature space through

supervised contrastive learning (SCL). The **core idea** of our technical design is that head-class in-distribution, tail-class in-distribution, and OOD samples play different roles in long-tailed OOD detection, and thus they should be treated differently in contrastive learning. Compared to traditional supervised contrastive learning (Khosla et al., 2020), PASCL is featured by two important insights tailored for OOD detection on long-tailed training sets: *partiality* and *asymmetry*.

▷ **Partiality:** Through feature visualization (Fig. 1), we observe tail-class in-distribution samples to be heavily overlapped with OOD samples in the feature space, while head-class in-distribution samples are relatively well separated from OOD samples. For this reason, we only partially apply contrastive loss on tail-class in-distribution and OOD samples, rather than on all training samples. In other words, we only push OOD data away from tail-class in-distribution data. Otherwise, if we simultaneously pushed OOD samples away from all in-distribution samples, the learning process would bias towards discriminating the head-class in-distribution and OOD data, so that the difference between tail-class in-distribution and OOD data could not be effectively learned. This is because head-class in-distribution and OOD data make up the great majority of the training set. Experimental results show partiality plays an important role in OOD detection under long-tailed recognition.

▷ **Asymmetry:** For traditional supervised contrastive learning, samples from the same classes are considered as positive pairs. However, in OOD detection, the training OOD images are not necessarily from the same class, nor should they be pulled together in the feature space. Indeed, the OOD training set usually has vast diversity (in both low-level visual features and high-level semantics) in order to be representative enough for the open visual world. Therefore, while pushing tail-class in-distribution samples away from OOD samples, we pull tail-class in-distribution samples within the same classes together, but do *not* pull OOD training samples together. Experimental results show that such asymmetric push-and-pull operation on in-distribution and OOD samples is an important knob for good performance.

To further boost the in-distribution classification accuracy in OOD detection, we propose Auxiliary Branch Finetuning (ABF). Current state-of-the-art OOD detection methods train the models using both in-distribution and OOD data (Hendrycks et al., 2019; Liu et al., 2020; Mohseni et al., 2020), where the batch normalization (BN) statistics are estimated using both in-distribution and OOD data. Although such BN statistics of *mixture* distribution between in-distribution and OOD data are beneficial for anomaly detection, they are not the optimal choice for in-distribution classification. This is because the in-distribution and OOD data have different underlying distributions, so that the statistics of the mixture distribution estimated in BN do not match those of the in-distribution test data. Previous work has

shown that even slight mismatch in BN statistics can lead to significant performance decay under covariate shifts (Chang et al., 2019; Benz et al., 2021; Xie & Yuille, 2019). Moreover, the classification layer (i.e., the last fully connected layer) has been empirically shown to play an important role in long-tailed recognition (Kang et al., 2020). Therefore, our auxiliary branch finetuning scheme uses two separate branches of BN and classification layers for anomaly detection and in-distribution classification, while all the other layers are shared across the two tasks.

We evaluate our method on the recently published semantically coherent OOD detection benchmarks (Yang et al., 2021). PASCL outperforms previous state-of-the-art method by 1.29%, 1.45%, 0.69% anomaly detection false positive rate (FPR) and 3.24%, 4.06%, 7.89% in-distribution classification accuracy on CIFAR10-LT, CIFAR100-LT (Cao et al., 2019), and ImageNet-LT (Liu et al., 2019), respectively.

2. Related Work

OOD detection Hendrycks & Gimpel (2017) formally studied the OOD detection problem in deep learning. The authors proposed to use the maximum softmax probability (MSP) as a naive baseline for OOD detection, and observed DNNs are commonly overconfident on OOD test samples. Other OOD detection measures such as Gram matrix (Sastry & Oore, 2020), Mahalanobis distance (Lee et al., 2018), and free energy (Liu et al., 2020) were also studied.

Hendrycks et al. (2019) proposed the classical Outlier Exposure (OE) approach, which utilized unlabeled auxiliary training set as OOD training data. OE assigns uniform distribution as pseudo labels for OOD training data, and then minimizes the cross-entropy loss on both in-distribution and OOD training data. The use of auxiliary OOD training sample greatly boosts OOD detection performance compared with previous methods that only use in-distribution training samples. OE has since become a cornerstone where many following works are based upon. EnergyOE (Liu et al., 2020) maximizes the free-energy of OOD training samples. OECC (Papadopoulos et al., 2021) replaced the cross-entropy loss in OE with total variance loss and further added a confidence-calibration loss term. NTOM (Chen et al., 2021) proposed to classify all OOD samples into one abstaining class, and used hard-sample mining to select the most informative OOD training samples in each epoch.

Lately, Yang et al. (2021) considered a more challenging scenario where the unlabeled auxiliary training set may contain both in-distribution and OOD data. They proposed unsupervised dual grouping (UDG) to split the unlabeled in-distribution and OOD training samples. The authors also pointed out that in-distribution data even widely exist in previous OOD benchmarks. To solve this problem,

the authors constructed a new OOD detection benchmark, named semantically coherent out-of-distribution detection (SC-OOD), which appears more challenging than previous benchmarks (Yang et al., 2021).

Several previous works have shown that contrastive learning (CL) and distance metric learning benefits OOD detection on balanced training sets (Winkens et al., 2020; Zhou et al., 2021a; Sohn et al., 2021; Yang et al., 2020). Our PASCL differs from them since it is tailored for the OOD detection problem on long-tailed dataset, which is a new challenge unaddressed in the prior literature.

Long-tailed recognition The most straight forward solution for long-tailed recognition is to re-balance the training set using undersampling or oversampling (He & Garcia, 2009). However, such methods lead to unsatisfactory performance in large-scaled deep learning (Cui et al., 2019; Wang et al., 2021b). Kang et al. (2020) observed that after the DNN has been trained on imbalanced data, simply retraining the classification layer on a re-balanced dataset can significantly boost long-tailed recognition performance. Recently, Menon et al. (2021) proposed a formal statistical framework for long-tailed recognition and a statistically grounded long-tailed recognition method termed logit adjustment that achieves impressive improvements over previous methods.

Long-tailed OOD detection Roy et al. (2022) studied a long-tail OOD detection problem in medical image analysis, whose problem setting is relevant but intrinsically different with ours. In their work, all the head-classes are in-distribution while all the tail-classes are considered as OOD. Their goal is to distinguish the head-class in-distribution samples from the tail-class OOD samples. In contrast, in our *more general* problem setting, the in-distribution samples come from both head-classes and tail-classes, and the same goes for the OOD samples. OLTR (Liu et al., 2019) is a long-tailed recognition method taking open set classification into consideration. Its main goal is to improve in-distribution classification accuracy in long-tailed recognition instead of OOD detection. Although it achieves better OOD detection performance than naive baselines such as MSP, it is not comparable with state-of-the-art OOD detection methods such as OE (see Appendix C.3 for results). In contrast, our paper focuses on improve OOD detection performance on long-tailed training sets. More related works are discussed in Appendix D.

3. The Challenge of OOD Detection in Long-Tailed Recognition

In this section, we first examine the performance of existing OOD detection methods on long-tailed training sets. Specifically, we benchmark six different OOD detection methods, including both widely used classical methods (e.g., MSP (Hendrycks & Gimpel, 2017), OE (Hendrycks et al.,

2019), EnergyOE (Liu et al., 2020), SOFL (Mohseni et al., 2020) and very recently published new methods (e.g., OECC (Papadopoulos et al., 2021), NTOM (Chen et al., 2021)). We train all methods on CIFAR10 (Krizhevsky, 2009) and CIFAR10-LT (i.e., the long-tailed version of CIFAR10) with imbalance ratio¹ $\rho = 100$.

We then evaluate OOD detection performance on the recently published SC-OOD benchmarks. The results are summarized in Table 1.² Compared with the results obtained on the balanced CIFAR10 dataset, CIFAR10-LT leads to significant performance drop in terms of both anomaly detection and in-distribution classification.

Table 1. The challenge of OOD detection in long-tailed recognition: existing OOD detection methods suffer significant performance drop when trained on long-tailed datasets. For each method, we train two ResNet18 models on CIFAR10 and CIFAR10-LT, respectively. Reported are the average performance across six different OOD test sets in the SC-OOD detection benchmark. The performance drop of each method from CIFAR10 training set to CIFAR10-LT training set are shown in parenthesis. All results are shown in percentage.

Method	Dataset	AUROC (↑)	AUPR (↑)	FPR95 (↓)	ACC (↑)
NT (MSP)	CIFAR10	85.86	84.37	52.52	93.45
	CIFAR10-LT	72.28 (-13.58)	70.27 (-14.10)	66.07 (+13.55)	72.34 (-21.11)
OE	CIFAR10	96.68	96.29	14.59	92.81
	CIFAR10-LT	89.92 (-6.75)	87.71 (-8.58)	34.80 (+20.21)	73.30 (-19.51)
EnergyOE	CIFAR10	96.59	96.37	14.80	93.07
	CIFAR10-LT	89.31 (-7.27)	88.92 (-7.45)	40.88 (+26.08)	74.68 (-18.39)
SOFL	CIFAR10	96.74	96.60	14.57	89.13
	CIFAR10-LT	91.13 (-5.61)	90.49 (-6.10)	34.98 (+20.41)	54.42 (-34.71)
OECC	CIFAR10	96.27	95.41	14.77	91.95
	CIFAR10-LT	87.28 (-8.99)	86.29 (-9.12)	45.24 (+30.47)	60.16 (-31.79)
NTOM	CIFAR10	96.92	96.95	14.95	91.44
	CIFAR10-LT	92.89 (-4.03)	92.31 (-4.65)	29.03 (+14.09)	66.41 (-25.03)

Table 2. Naive combinations of state-of-the-art OOD detection and long-tailed recognition methods cannot address the challenge of OOD detection in long-tailed recognition (LTR). All methods are trained on CIFAR10/100-LT using ResNet18 and evaluated on the SC-OOD benchmarks. The best results are shown in bold and the second-best ones are underlined. Mean and standard deviation over six random runs are shown for our method.

\mathcal{D}_m	OOD Detection Method	LTR Method	AUROC (↑)	AUPR (↑)	FPR95 (↓)	ACC (↑)
CIFAR 10-LT	OE +	None	<u>89.92</u>	<u>87.71</u>	34.80	73.30
		Re-weighting	89.34	86.39	37.09	70.35
		τ -norm	89.58	85.88	33.80	73.33
		LA	89.46	86.39	34.94	73.93
		Our method	90.99 ±0.19	89.24 ±0.34	33.36 ±0.79	77.08 ±1.01
	Our method	90.99 ±0.19	89.24 ±0.34	33.36 ±0.79	77.08 ±1.01	
CIFAR 100-LT	OE +	None	72.62	<u>66.73</u>	68.69	39.33
		Re-weighting	72.07	66.05	70.62	39.42
		τ -norm	<u>72.71</u>	66.59	68.04	40.87
		LA	72.56	66.48	68.24	42.06
		Our method	73.32 ± 0.32	67.18 ± 0.10	67.44 ± 0.58	43.10 ± 0.47
	Our method	73.32 ± 0.32	67.18 ± 0.10	67.44 ± 0.58	43.10 ± 0.47	

We further show that such challenge can not be easily solved

¹Imbalance ratio: the number of training samples in the most frequent class divided by that of the least frequent class.

²The evaluation measures will be described in Section 5.1.

by naive solutions such as combining state-of-the-art OOD detection methods with state-of-the-art long-tailed recognition methods. Specifically, we combine OE with three popular long-tailed recognition methods: re-weighting (Cui et al., 2019), τ -norm (Kang et al., 2020), and logits adjustment (LA) (Menon et al., 2021). The results are shown in Table 2. As we can see, combining OE with long-tailed recognition methods does not necessarily bring benefit compared with the original OE baseline: E.g., on CIFAR10-LT, OE has the best AUROC and AUPR compared with all three combination methods. In contrast, our method (to be introduced later) achieves considerable improvements.

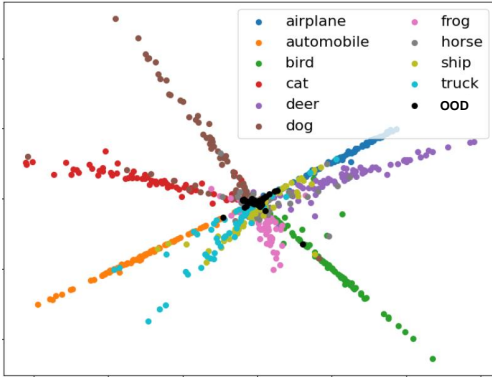


Figure 1. Visualization of in-distribution and OOD samples in the feature space. We visualize the penultimate layer features of ResNet18 trained by OE on CIFAR10-LT, using the same visualization method as in (Pang et al., 2020). 100 test samples from each in-distribution class and 100 OOD samples are visualized. The model struggles to distinguish tail-class in-distribution classes like “horse” (grey) and “ship” (yellow) from OOD samples (black). Note that in CIFAR10-LT, the number of training samples in each class follows the alphabetic order: the “airplane” class has the most training samples and the “truck” class has the least. For comparison, see visualization on the balanced CIFAR10 training set in Appendix B.

Why OE fails on long-tailed recognition Feature distributions of models trained by OE on CIFAR10 and CIFAR10-LT are visualized in Fig. 4 (in Appendix) and 1, respectively. As we can see, on the balanced CIFAR10 dataset, the OOD samples are well separated from all in-distribution classes (Fig. 4). In contrast, when trained on CIFAR10-LT, the features of tail-class in-distribution (e.g., horse, ship) samples and OOD samples heavily overlap, while those of head-class in-distribution (e.g., airplane, automobile, bird) samples are well separated from OOD samples (Fig. 1). This indicates that on long-tailed datasets, models trained by OE struggles to distinguish tail-class in-distribution samples from OOD samples, which is likely to be a major contributing cause of the performance drop of OE on long-tailed recognition. An intuitive explanation for this phenomenon is that the tail-class in-distribution samples have few occurrence in the training set, and are thus more likely to result in low-

confidence predictions like the OOD samples.

4. Method

4.1. Preliminary

State-of-the-art OOD detection methods train the model on both in-distribution and OOD training data. Let \mathcal{D}_{in} and \mathcal{D}_{out} denote an in-distribution training set and an unlabeled OOD training set, respectively. The training objective of the existing OOD detection methods (e.g., OE, EnergyOE, OECC, NTOM, etc.) is formulated as:

$$\mathcal{L} = \mathbb{E}_{x \sim \mathcal{D}_{in}}[\mathcal{L}_{in}(x)] + \lambda \mathbb{E}_{x \sim \mathcal{D}_{out}}[\mathcal{L}_{out}(x)], \quad (1)$$

where \mathcal{L}_{in} is for in-distribution classification and \mathcal{L}_{out} is to assign low prediction confidence to OOD data. For example in OE, $\mathcal{L}_{in} = \text{XEntropy}(x, y(x))$ is the cross-entropy loss on in-distribution data x and the corresponding label $y(x)$, and $\mathcal{L}_{out}(x) = \text{KL}(f(x) \| u)$ is the KL divergence between the model output (i.e., logits) $f(x)$ and the uniform distribution vector u .

4.2. Partial and Asymmetric Contrastive Learning

Figure 2 illustrates the core idea of our partial and asymmetric supervised contrastive learning strategy for OOD detection in long-tailed recognition. Traditional supervised contrastive learning treats all classes equally: pulling is applied within each class and pushing between any two different classes (Figure 2 (a)). In contrast, PASCL treats head-class in-distribution, tail-class in-distribution, and OOD classes differently: PASCL applies pulling only within each tail-class in-distribution class, and applies pushing among tail-class in-distribution and OOD samples (Figure 2 (b)).




Figure 2. Illustration of our partial and asymmetric push-and-pull strategy, in contrast to those in the original supervised contrastive learning (SCL). The ellipses with different colors represent different types of training samples. Grey: tail-class in-distribution classes; blue: head-class in-distribution classes; orange: OOD training samples. The green and red arrows show pulling and pushing operations in contrastive learning, respectively.

Next we formularize our method to achieve the above partial and asymmetric push-and-pull strategy. Specifically, we design the following partial and asymmetric supervised contrastive learning (PASCL) loss:

$$\mathcal{L}_c = \mathbb{E}_{x \in \mathcal{I}}[\mathcal{L}_c(x)], \text{ where } \mathcal{L}_c(x) = \sum_{p \in \mathcal{P}(x)} \frac{1}{|\mathcal{P}(x)|} \log \frac{\exp(z(x)^T z(p)/\tau)}{\sum_{a \in \mathcal{A}(x)} \exp(z(x)^T z(a)/\tau)}, \quad (2)$$

where $z(x)$ is the output of a non-linear projection on the model’s penultimate layer at sample x (as in the original supervised contrastive learning), τ is the scaling temperature, $\mathcal{P}(x) = \{p|y(p) = y(x), p \in \mathcal{D}_{\text{in}} \setminus \{x\}\}$ is the set of positive samples with respect to x (i.e., the samples with identical ground truth labels with x), $\mathcal{A}(x)$ is the set of all contrastive samples with respect to x , and \mathcal{I} is the set of training samples where the PASCL loss is applied on. At first glance, Eq. (2) *seems* identical with the original supervised contrastive learning loss (Khosla et al., 2020), previously used to learn good features for in-distribution classification. However, the *key difference* lies in the definitions of \mathcal{I} and $\mathcal{A}(x)$, as we detail below.

The original supervised contrastive learning (Khosla et al., 2020) All classes play the same role in the original supervised contrastive learning: Pulling is done within every class and pushing between every pair of different classes. Such uniform contrastive loss can be applied either only on in-distribution samples (i.e., $\mathcal{I} = \mathcal{D}_{\text{in}}$, $\mathcal{A}(x) = \mathcal{D}_{\text{in}} \setminus \{x\}$) or on both in-distribution and OOD samples (i.e., $\mathcal{I} = \mathcal{D}_{\text{in}} \cup \mathcal{D}_{\text{out}}$, $\mathcal{A}(x) = \mathcal{D}_{\text{in}} \cup \mathcal{D}_{\text{out}} \setminus \{x\}$). In the latter case, all OOD samples are assumed to form the $(C+1)^{\text{th}}$ class along with the C in-distribution classes. Experimental results show the latter empirically achieves slightly better performance than the former when combined with OE (Appendix C.2).

Partial supervised contrastive learning Here we apply supervised contrastive learning only on tail-class in-distribution and OOD data: $\mathcal{I} = \mathcal{D}_{\text{tail}} \cup \mathcal{D}_{\text{out}}$, $\mathcal{A}(x) = \mathcal{D}_{\text{tail}} \cup \mathcal{D}_{\text{out}} \setminus \{x\}$, where $\mathcal{D}_{\text{tail}}$ is the set of all tail-class in-distribution training data. This design is motivated by our observations in Figure 1, where head-class in-distribution samples are easier to be separated from the OOD samples, since there are enough occurrence of them in the training set. Thus, we do not use an extra supervised contrastive learning loss to explicitly push them away. Otherwise, if we simultaneously pushed OOD samples away from all in-distribution training samples, the difference between head-class in-distribution and OOD data would be over-emphasized, hindering the learning of the difference between tail-class in-distribution and OOD samples.

Asymmetric supervised contrastive learning In this case, OOD training samples are only contained in $\mathcal{A}(x)$ but not in \mathcal{I} : $\mathcal{I} = \mathcal{D}_{\text{in}}$, $\mathcal{A}(x) = \mathcal{D}_{\text{in}} \cup \mathcal{D}_{\text{out}} \setminus \{x\}$. That is, we still push in-distribution samples away from OOD samples and in-distribution samples from different classes, and pull in-distribution samples within the same class together, but do not pull OOD samples together. This is because the OOD training samples are not necessarily from the same

class. Hence they should not be pulled together in the feature space, as the OOD training set typically has huge diversity in order to be representative for the open visual world.

Putting them together: the PASCL framework Combining partiality and asymmetry together, we have

$$\mathcal{I} = \mathcal{D}_{\text{tail}}, \mathcal{A}(x) = \mathcal{D}_{\text{tail}} \cup \mathcal{D}_{\text{out}} \setminus \{x\}. \quad (3)$$

In summary, our final loss function is

$$\mathcal{L} = \mathbb{E}_{x \sim \mathcal{D}_{\text{in}}}[\mathcal{L}_{\text{in}}(x)] + \lambda_1 \mathbb{E}_{x \sim \mathcal{D}_{\text{out}}}[\mathcal{L}_{\text{out}}(x)] + \lambda_2 \mathbb{E}_{x \sim \mathcal{I}}[\mathcal{L}_c(x)], \quad (4)$$

where $\mathcal{L}_{\text{in}}(x)$ and $\mathcal{L}_{\text{out}}(x)$ are identical with those in OE (see Eq. (1)), and $\mathcal{L}_c(x)$ is as defined in Eq. (2) with \mathcal{I} and $\mathcal{A}(x)$ defined in Eq. (3).

4.3. Auxiliary Branch Finetuning (ABF)

As demonstrated in Eq. (1), current state-of-the-art OOD detection methods train the models using both in-distribution and OOD data (Hendrycks et al., 2019; Liu et al., 2020; Mohseni et al., 2020), where the BN statistics are estimated using both in-distribution and OOD data. Such BN statistics of *mixture* distribution between in-distribution and OOD data are beneficial for anomaly detection (i.e., to distinguish in-distribution from OOD samples), but not optimal for in-distribution classification. The ideal case is to use mixture BN statistics for anomaly detection and *pure* in-distribution BN statistics for in-distribution classification. Moreover, the classification layer (i.e., the last fully connected (FC) layer) has been empirically shown to play an important role in long-tailed recognition (Kang et al., 2020). Thus, it may also be beneficial to use a separate classification layer trained only on \mathcal{D}_{in} for better in-distribution classification. To this end, we propose Auxiliary Branch Finetuning (ABF).

The overall framework of auxiliary branch finetuning is shown in Figure 3. We add an auxiliary counterpart for each BN layer and the classification layer in the model. The training process consists of two stages. In the first stage, we train the main branch using both in-distribution and OOD data, by minimizing Eq. (4). In the second stage, we finetune the auxiliary BN and classification layers, while fixing all the other layers, by minimizing the logits-adjustment (LA) cross-entropy loss (Menon et al., 2021) only on in-distribution data³ for a few iterations:

$$\mathbb{E}_{x \sim \mathcal{D}_{\text{in}}}[\mathcal{L}_{\text{in}}^{\text{LA}}(x)]. \quad (5)$$

During test time, we forward each test image twice, using the main and auxiliary branches respectively. The outputs from the main branch are used for anomaly detection, while those from the second branch are used for in-distribution classification.

³Naively combining LA with OE has been shown ineffective in Table 2.

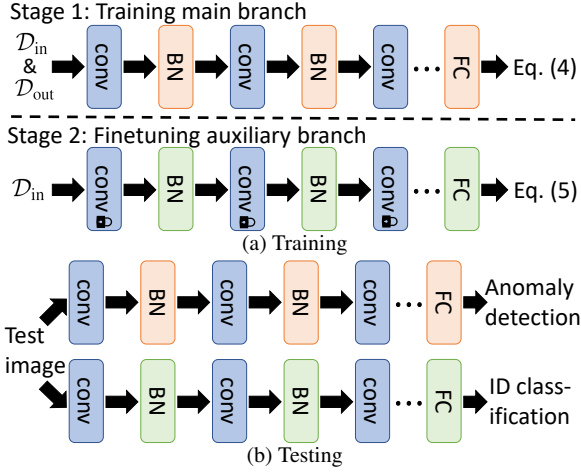


Figure 3. The framework of auxiliary branch finetuning (ABF) at (a) training and (b) test time. Orange layers are trained in stage one. Green layers are finetuned in stage two, using the weights learned in stage one as initialization. Blue layers are trained in stage one and kept fixed (with the lock icon) in stage two.

The final PASCL framework is summarized in Algorithm 1.

Algorithm 1 Partial and Asymmetric Supervised Contrastive Learning (PASCL)

Input: in-distribution training set \mathcal{D}_{in} , OOD training set \mathcal{D}_{out} , main branch training iteration n_1 , auxiliary branch finetuning iteration n_2 .

#Stage 1: Train main branch.

for $i = 1$ to n_1 **do**

 Sample a batch of in-distribution and OOD training samples.

 Update the main branch model by minimizing Eq. (4)

end for

#Stage 2: Finetune auxiliary branch.

Fix all layers except the auxiliary BN and classification layers in the model.

for $i = 1$ to n_2 **do**

 Sample a batch of in-distribution training samples.

 Update the auxiliary BN and classification layers by minimizing Eq. (5).

end for

5. Experiments

5.1. Experiment Settings

In-distribution training and test sets (\mathcal{D}_{in} , \mathcal{D}_{in}^{test}) We use three popular long-tailed image classification datasets, CIFAR10-LT, CIFAR100-LT (Cao et al., 2019), and ImageNet-LT (Liu et al., 2019), as the in-distribution training data (i.e., \mathcal{D}_{in}). Following Menon et al. (2021), we use the default imbalance ratio $\rho = 100$ on CIFAR10-LT and CIFAR100-LT, and we conduct ablation studies on differ-

Table 3. Results on CIFAR10-LT using ResNet18. The best results are shown in bold. Mean and standard deviation over six random runs are reported for OE and our method. ‘‘Average’’ means the results averaged across six different \mathcal{D}_{out}^{test} sets.

(a) OOD detection results and in-distribution classification results in terms of ACC95.

\mathcal{D}_{out}^{test}	Method	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)	ACC95 (\uparrow)
Texture	OE	92.59 \pm 0.42	83.32 \pm 1.67	25.10 \pm 1.08	84.52 \pm 0.76
	Ours	93.16 \pm 0.37	84.80 \pm 1.50	23.26 \pm 0.91	85.86 \pm 0.72
SVHN	OE	95.10 \pm 1.01	97.14 \pm 0.81	16.15 \pm 1.52	81.33 \pm 0.81
	Ours	96.63 \pm 0.90	98.06 \pm 0.56	12.18 \pm 3.33	82.72 \pm 1.51
CIFAR100	OE	83.40 \pm 0.30	80.93 \pm 0.57	56.96 \pm 0.91	94.56 \pm 0.57
	Ours	84.43 \pm 0.23	82.99 \pm 0.48	57.27 \pm 0.88	94.48 \pm 0.31
Tiny ImageNet	OE	86.14 \pm 0.29	79.33 \pm 0.65	47.78 \pm 0.72	91.19 \pm 0.33
	Ours	87.14 \pm 0.18	81.54 \pm 0.38	47.69 \pm 0.59	91.20 \pm 0.35
LSUN	OE	91.35 \pm 0.23	87.62 \pm 0.82	27.86 \pm 0.68	85.49 \pm 0.69
	Ours	93.17 \pm 0.15	91.76 \pm 0.53	26.40 \pm 1.00	86.67 \pm 0.90
Places365	OE	90.07 \pm 0.26	95.15 \pm 0.24	34.04 \pm 0.91	87.07 \pm 0.53
	Ours	91.43 \pm 0.17	96.28 \pm 0.14	33.40 \pm 0.88	87.87 \pm 0.71
Average	OE	89.77 \pm 0.27	87.25 \pm 0.61	34.65 \pm 0.46	87.36 \pm 0.51
	Ours	90.99 \pm 0.19	89.24 \pm 0.34	33.36 \pm 0.79	88.13 \pm 0.56

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

Method	ACC@FPR n (\uparrow)			
	0	0.001	0.01	0.1
OE	73.84 \pm 0.77	73.90 \pm 0.77	74.46 \pm 0.81	78.88 \pm 0.66
Ours	77.08 \pm 1.01	77.13 \pm 1.02	77.64 \pm 0.99	81.96 \pm 0.85

(c) Comparison with other methods.

\mathcal{D}_{out}^{test}	Method	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)	ACC (\uparrow)
Average	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	<u>74.68</u>
	OE	89.77 \pm 0.27	87.25 \pm 0.61	34.65 \pm 0.46	73.84 \pm 0.77
	Ours	90.99 \pm 0.19	89.24 \pm 0.34	33.36 \pm 0.79	77.08 \pm 1.01

ent ρ in Section 5.3. We use the original CIFAR10 and CIFAR100 test sets and the ImageNet validation set as the in-distribution test sets (i.e., \mathcal{D}_{in}^{test}).

OOD training set \mathcal{D}_{out} For experiments on CIFAR10-LT and CIFAR100-LT, we use TinyImages80M (Torralba et al., 2008) as the unlabeled OOD training images (i.e., \mathcal{D}_{out}) following (Hendrycks et al., 2019; Liu et al., 2020). For experiments on ImageNet-LT, we construct our own \mathcal{D}_{out} named ImageNet-Extra. Specifically, ImageNet-Extra contains 517, 711 images belonging to 500 classes randomly sampled from ImageNet-22k (Deng et al., 2009) but not overlapping with the 1, 000 in-distribution classes in ImageNet-LT.

OOD test set \mathcal{D}_{out}^{test} For experiments on CIFAR10-LT and CIFAR100-LT, we use the recently published semantically coherent out-of-distribution detection (SC-OOD) benchmark datasets (Yang et al., 2021) as OOD test sets \mathcal{D}_{out}^{test} . Specifically, for experiments on CIFAR10-LT, we use Textures (Cimpoi et al., 2014), SVHN (Netzer et al., 2011), CIFAR100, Places365 (Zhou et al., 2017), LSUN (Yu et al., 2015), and Tiny ImageNet (Le & Yang, 2015) as OOD evaluation datasets \mathcal{D}_{out}^{test} ; for experiments on CIFAR100-LT, we use Textures, SVHN, CIFAR10, Places365, LSUN, and

Table 4. Results on CIFAR100-LT using ResNet18. The best results are shown in bold. Mean and standard deviation over six random runs are reported for OE and our method. ‘‘Average’’ means the results averaged across six different $\mathcal{D}_{\text{out}}^{\text{test}}$ sets.

(a) OOD detection results and in-distribution classification results in terms of ACC95.

$\mathcal{D}_{\text{out}}^{\text{test}}$	Method	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)	ACC95 (\uparrow)
Texture	OE	76.71 \pm 1.20	58.79 \pm 1.39	68.28 \pm 1.53	71.43 \pm 1.58
	Ours	76.01 \pm 0.66	58.12 \pm 1.06	67.43 \pm 1.93	73.11 \pm 1.55
SVHN	OE	77.61 \pm 3.26	86.82 \pm 2.50	58.04 \pm 4.82	64.27 \pm 3.26
	Ours	80.19 \pm 2.19	88.49 \pm 1.59	53.45 \pm 3.60	64.50 \pm 1.87
CIFAR10	OE	62.23 \pm 0.30	57.57 \pm 0.34	80.64 \pm 0.98	82.67 \pm 0.99
	Ours	62.33 \pm 0.38	57.14 \pm 0.20	79.55 \pm 0.84	82.30 \pm 1.07
Tiny ImageNet	OE	68.04 \pm 0.37	51.66 \pm 0.51	76.66 \pm 0.47	76.22 \pm 0.61
	Ours	68.20 \pm 0.37	51.53 \pm 0.42	76.11 \pm 0.80	77.56 \pm 1.15
LSUN	OE	77.10 \pm 0.64	61.42 \pm 0.99	63.98 \pm 1.38	65.64 \pm 1.03
	Ours	77.19 \pm 0.44	61.27 \pm 0.72	63.31 \pm 0.87	68.05 \pm 1.24
Places365	OE	75.80 \pm 0.45	86.68 \pm 0.38	65.72 \pm 0.92	67.04 \pm 0.49
	Ours	76.02 \pm 0.21	86.52 \pm 0.29	64.81 \pm 0.27	69.04 \pm 0.90
Average	OE	72.91 \pm 0.68	67.16 \pm 0.57	68.89 \pm 1.07	71.21 \pm 0.84
	Ours	73.32 \pm 0.32	67.18 \pm 0.10	67.44 \pm 0.58	72.43 \pm 0.66

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

Method	ACC@FPR n (\uparrow)			
	0	0.001	0.01	0.1
OE	39.04 \pm 0.37	39.07 \pm 0.38	39.38 \pm 0.38	42.40 \pm 0.44
Ours	43.10 \pm 0.47	43.12 \pm 0.47	43.39 \pm 0.48	46.14 \pm 0.38

(c) Comparison with other methods.

$\mathcal{D}_{\text{out}}^{\text{test}}$	Method	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)	ACC (\uparrow)
Average	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 \pm 0.68	67.16 \pm 0.57	68.89 \pm 1.07	39.04 \pm 0.37
	Ours	73.32 \pm 0.32	67.18 \pm 0.10	67.44 \pm 0.58	43.10 \pm 0.47

Tiny ImageNet as $\mathcal{D}_{\text{out}}^{\text{test}}$.⁴ For experiments on ImageNet-LT, we construct ImageNet-1k-OOD as $\mathcal{D}_{\text{out}}^{\text{test}}$. Specifically, ImageNet-1k-OOD contains 50,000 OOD test images from 1,000 classes randomly selected from ImageNet-22k (with 50 images in each class), which is of the same size as the in-distribution test set. The 1,000 classes in ImageNet-1k-OOD are not overlapped with either the 1,000 in-distribution classes in ImageNet-LT or the 500 OOD training classes in ImageNet-Extra.

Evaluation measures Following Hendrycks et al. (2019); Mohseni et al. (2020); Yang et al. (2021), we use the below evaluation measures:

- **AUROC**: The area under the receiver operating characteristic curve. AUROC is equivalent to the probability that a positive example has a larger detector score than a negative example. A perfect detector has 100% AUROC, while a random guess leads to 50% AUROC.
- **AUPR**: The area under precision-recall curve. This is also known as the average precision over all recall values.

⁴According to Yang et al. (2021), not all images in those datasets are OOD images with respect to CIFAR10 and CIFAR100. We use the benchmarks released by Yang et al. (2021) where each test image is categorized into either $\mathcal{D}_{\text{in}}^{\text{test}}$ or $\mathcal{D}_{\text{out}}^{\text{test}}$ according to their ground-truth semantic meaning.

- **FPR@TPR n** : The false positive rate (FPR) when n (in percentage) OOD samples have been successfully detected (i.e., when the true positive rate (TPR) is n). Previous papers (Hendrycks & Gimpel, 2017) mainly use the measure **FPR95**, which is the abbreviation for FPR@TPR95%. They set n to large values such as 95% since they care about the FPR when most OOD samples are successfully detected.
- **ACC@TPR n** : The classification accuracy on the remaining in-distribution data when n (in percentage) OOD samples have been successfully detected. **ACC95** is the abbreviation for ACC@TPR95%.
- **ACC@FPR n** : The classification accuracy on the remaining in-distribution data when n (in percentage) in-distribution samples are mistakenly detected as OOD samples (i.e., when FPR is n). **ACC@FPR0**, also abbreviated as **ACC**, is the accuracy on the entire in-distribution test set. Previous papers (Yang et al., 2021) set n to small values such as 1% since they care about the accuracy when few in-distribution samples are falsely detected as OOD. By definition, the value of ACC@FPR n is irrelevant to the choice of $\mathcal{D}_{\text{out}}^{\text{test}}$.

Methods in comparison As shown in Table 1, among all those OOD detection methods we evaluated, OE has arguably the best OOD detection performance on long-tailed datasets. Thus we use OE as the main baseline method. We also compare with two more recent methods: EnergyOE (Liu et al., 2020) and OECC (Papadopoulos et al., 2021). We also show results of MSP on standard training (ST) models.

Models For experiments on CIFAR10 and CIFAR100, we use the standard ResNet18 (He et al., 2016) following Yang et al. (2021). For experiments on ImageNet, we use ResNet50 (He et al., 2016). More details on experimental settings (e.g., learning rate, etc.) are in Appendix A.

5.2. Main Results

Results on CIFAR10-LT, CIFAR100-LT, and ImageNet-LT are shown in Table 3, 4, and 5, respectively. Table 3 and 4 both contain three sub-tables: In sub-table (a), we show AUROC, AUPR, FPR95 and ACC95 on each \mathcal{D}_{out} and also the average value on all six $\mathcal{D}_{\text{out}}^{\text{test}}$ datasets, since these four measures may vary among different $\mathcal{D}_{\text{out}}^{\text{test}}$. In sub-table (b), we show ACC@FPR n with different n values, which remains the same regardless of the choice of $\mathcal{D}_{\text{out}}^{\text{test}}$. The results of other baseline methods (in terms of both anomaly detection and in-distribution classification) are summarized in sub-table (c) due to space limit. The results on ImageNet-LT can be reported within one table, since there is only one $\mathcal{D}_{\text{out}}^{\text{test}}$.

As we can see, our method outperforms OE and also other baseline methods by a considerable margin. For example, on CIFAR10-LT, our method achieves 1.22% higher average AUROC, 1.99% higher average AUPR, 1.29% lower

Table 5. Results on ImageNet-LT using ResNet50. The best and second-best results are bolded and underlined, respectively. Improvements of our method over OE are shown in parentheses.

\mathcal{D}_{out}^{test}	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
ImageNet -1k-OOD	ST (MSP)	53.81	51.63	95.38	90.15	83.52	72.97	96.67	92.61	87.43	77.52	<u>39.65</u>	<u>39.68</u>	<u>40.00</u>	<u>43.18</u>
	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	94.15	87.72	78.36	63.71	80.18	74.38	67.65	59.68	38.50	38.52	38.72	40.99
	OE	<u>66.33</u>	<u>68.29</u>	95.11	88.22	78.68	65.28	95.46	88.22	78.68	65.28	37.60	37.62	37.79	40.00
	Ours	68.00 (+1.67)	70.15 (+1.86)	94.38 (-0.73)	<u>87.53</u> (-0.69)	78.12 (-0.56)	62.48 (-2.80)	<u>95.69</u> (+0.23)	<u>89.55</u> (+1.33)	<u>80.88</u> (+2.20)	<u>69.60</u> (+4.32)	45.49 (+7.89)	45.51 (+7.89)	45.62 (+7.83)	47.49 (+7.49)

Table 6. The importance of each component (i.e., partiality, asymmetry, and the auxiliary branch) in PASCL. Experiments are conducted on CIFAR10-LT and CIFAR100-LT (both with $\rho = 100$) with ResNet18. The first row in each block, where no supervised contrastive learning is used, is the OE baseline. SVHN is used as \mathcal{D}_{out}^{test} . Mean and standard deviation over six random runs are reported.

\mathcal{D}_{in}	Asymmetry	Partiality	ABF	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)	ACC95 (\uparrow)	ACC@FPR $_n$ (\uparrow)			
								0	0.001	0.01	0.1
CIFAR10-LT	No contrastive loss (OE)			95.10 \pm 1.01	97.14 \pm 0.81	16.15 \pm 1.52	81.33 \pm 0.81	73.84 \pm 0.77	73.90 \pm 0.77	74.46 \pm 0.81	78.88 \pm 0.66
	\times	\times	\times	<u>95.34</u> \pm 1.58	<u>97.30</u> \pm 1.20	<u>15.12</u> \pm 3.07	81.94 \pm 1.28	75.03 \pm 1.46	75.09 \pm 1.45	75.60 \pm 1.44	80.02 \pm 1.10
	\times	\checkmark	\times	95.01 \pm 1.25	96.74 \pm 0.78	15.31 \pm 4.35	<u>82.34</u> \pm 1.56	74.46 \pm 1.80	74.52 \pm 1.80	75.04 \pm 1.76	80.21 \pm 0.99
	\checkmark	\times	\times	94.91 \pm 1.43	96.86 \pm 1.47	15.57 \pm 1.19	82.08 \pm 0.47	75.24 \pm 0.99	75.29 \pm 0.98	75.77 \pm 0.98	79.85 \pm 0.77
	\checkmark	\checkmark	\times	96.63 \pm 0.90	98.06 \pm 0.56	12.18 \pm 3.33	81.70 \pm 1.21	<u>76.20</u> \pm 0.79	<u>76.26</u> \pm 0.79	<u>76.85</u> \pm 0.81	<u>81.07</u> \pm 0.58
	\checkmark	\checkmark	\checkmark	96.63 \pm 0.90	98.06 \pm 0.56	12.18 \pm 3.33	82.72 \pm 1.51	77.08 \pm 1.01	77.13 \pm 1.02	77.64 \pm 0.99	81.96 \pm 0.85
CIFAR100-LT	No contrastive loss (OE)			77.61 \pm 3.26	86.82 \pm 2.50	58.04 \pm 4.82	64.27 \pm 3.26	39.04 \pm 0.37	39.07 \pm 0.38	39.38 \pm 0.38	42.40 \pm 0.44
	\times	\times	\times	78.05 \pm 2.12	87.18 \pm 0.87	59.10 \pm 5.03	66.44 \pm 3.90	40.21 \pm 0.43	40.25 \pm 0.43	40.56 \pm 0.45	43.71 \pm 0.42
	\times	\checkmark	\times	79.46 \pm 1.83	<u>88.01</u> \pm 1.90	54.59 \pm 3.34	63.86 \pm 2.52	40.24 \pm 0.53	40.28 \pm 0.53	40.60 \pm 0.55	<u>43.93</u> \pm 0.57
	\checkmark	\times	\times	<u>79.54</u> \pm 2.38	87.68 \pm 1.51	<u>54.27</u> \pm 3.69	63.33 \pm 2.87	40.00 \pm 0.42	40.04 \pm 0.41	40.36 \pm 0.42	43.60 \pm 0.42
	\checkmark	\checkmark	\times	80.19 \pm 2.19	88.49 \pm 1.59	53.45 \pm 3.60	63.10 \pm 1.87	<u>40.33</u> \pm 0.20	<u>40.36</u> \pm 0.20	<u>40.66</u> \pm 0.18	43.79 \pm 0.22
	\checkmark	\checkmark	\checkmark	80.19 \pm 2.19	88.49 \pm 1.59	53.45 \pm 3.60	64.50 \pm 1.87	43.10 \pm 0.47	43.12 \pm 0.47	43.39 \pm 0.48	46.14 \pm 0.38

average FPR95, 0.77% higher average ACC95, and 3.24% higher in-distribution accuracy than OE. On CIFAR100-LT, our method achieves 1.45% lower average FPR95, 1.22% higher average ACC95 and 4.06% higher in-distribution accuracy than OE. On ImageNet-LT, our method achieves 1.67% higher AUROC, 1.86% higher AUPR, 1.33% higher average ACC95, and 7.89% higher in-distribution accuracy than OE.

Simultaneously improving FPR95 and ACC95 Previous work has shown it extremely challenging to simultaneously improve FPR95 and ACC95 (Mohseni et al., 2020). The reason is simple: When the FPR95 is high, many in-distribution samples are falsely detected as OOD, including those hard or corner-case in-distribution samples. In result, the remaining in-distribution samples are mostly easy to classify, leading to high ACC95. We highlight that our method simultaneously achieves better FPR95 and ACC95 than OE. This crucial improvement shows our method learns better features for both the anomaly detection and in-distribution classification tasks.

Improvements on head and tail in-distribution classes

In Table 7, we show the improvements of our method over OE on head and tail in-distribution classes separately. As we can see, our method mainly benefits the tail classes.

Table 7. The results on head and tail classes of ImageNet-LT. The improvements of our method over OE are shown in parentheses.

Method	ACC (\uparrow)	
	Head classes	Tail classes
OE	54.29	20.90
Ours	54.73 (+0.44)	36.26 (+15.36)

5.3. Ablation Study

On partiality, asymmetry, and auxiliary branch finetuning As introduced in Section 4, there are three building blocks in our PASCL framework: (i) partiality and (ii) asymmetry in supervised contrastive learning and the (iii) auxiliary branch finetuning (ABF). In Table 6, we report the results of ablation studies on the three components using both CIFAR10-LT and CIFAR100-LT, to show the importance of each one. Here we use SVHN as \mathcal{D}_{out}^{test} since our method has relatively large performance gain over OE on SVHN, making it easier to identify where the performance gain come from. First, on both datasets, simply applying the original supervised contrastive learning onto the OE framework (i.e., the second rows in each block in Table 6) only brings marginal improvements over the OE baseline (i.e., the first rows in each block in Table 6), especially in terms of AUROC and AUPR. Second, partiality and asymmetry should always be used together for better performance. For example, on CIFAR10-LT, neither asymmetric nor partiality brings benefit over the original supervised contrastive learning when applied alone, while great improvements are made when they are used together. Third, the auxiliary branch finetuning can significantly boost in-distribution classification results (i.e., the last row in each block in Table 6).

On the percentage of tail classes As mentioned in Section 4, we apply contrastive loss on tail-class samples to push them away from OOD samples. A key hyper-parameter here is the threshold to define the separation of head and tail classes: the k (in percentage) classes with the least training

samples are defined as tail classes. Ablation study on k is provided in Table 8. As we can see, on both CIFAR10-LT and CIFAR100-LT, the best results are achieved around $k = 50\%$ (the default value in our experiments). The results at $k = 50\%$ are considerably better than those at $k = 100\%$ (without partiality in supervised contrastive learning) and $k = 0\%$ (without supervised contrastive learning, i.e., the OE baseline), showing the importance of our partiality design in supervised contrastive learning. Also, the performance of our method is stable with respect to k within a large range (e.g, $k \in [40\%, 60\%]$).

Table 8. Ablation study on the percentage (k) of tail classes. Experiments are conducted on CIFAR10/100-LT (both with $\rho = 100$) using ResNet18. The first row ($k = 100\%$) means to apply the asymmetric supervised contrastive learning on all in-distribution training samples (without partiality). The last row ($k = 0\%$) is the OE baseline where no supervised contrastive learning is used. Average results over six \mathcal{D}_{out}^{test} in SC-OOD benchmarks are reported.

\mathcal{D}_{in}	k	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)
CIFAR10-LT	100%	89.73	87.05	33.05
	60%	90.44	88.51	32.90
	50%	91.10	89.01	32.63
	40%	90.66	88.90	33.30
	0% (OE)	89.45	86.85	35.21
CIFAR100-LT	100%	72.95	66.36	67.86
	60%	73.24	67.26	67.59
	50%	73.54	67.29	66.92
	40%	72.76	67.07	68.37
	0% (OE)	72.62	66.73	68.69

On auxiliary branch finetuning layers As described in Section 4, auxiliary branch finetuning updates all BN layers and the classification (CLF) layer. In this section, we experiment with the following choices of finetuned layers: none of the layers (denoted as None), all BN layers (denoted as BN), the classification layer (denoted as CLF), all BN layers and the classification layer (denoted as BN & CLF), and every layer in the model (denoted as All). The results are shown in Table 9. As we can see, finetuning all BN layers or the classification layer both helps improve accuracy. Finetuning BN and classification layers simultaneously (our default setting) achieves the best accuracy, even outperforming finetuning all layers. A likely explanation is that the convolutional layers from stage one have learned features from OOD training data that are transferable to in-distribution data. Also, auxiliary branch finetuning (on BN & CLF layers) adds only a tiny overhead on model size.

More experimental results (e.g., ablation study on imbalance ratio, model structure, etc.) can be found in Appendix C.

6. Summary

To address the new challenging problem of OOD detection in long-tailed recognition, we proposed a novel PASCL

Table 9. Ablation study on auxiliary branch finetuning layers using ImageNet-LT and ResNet50. Model size in Million Bytes (MBs).

Layers	None	BN	CLF	BN & CLF	All
ACC (%)	41.32	43.42	44.57	45.49	45.12
Model size (MB)	94.11	94.30	94.20	94.38	188.23

framework to explicitly encourage the model to distinguish tail-class in-distribution samples from OOD samples. The core idea of PASCL is that head-class in-distribution, tail-class in-distribution, and OOD samples play different roles in OOD detection under long-tailed recognition and thus should be treated differently in contrastive learning. Experiments on long-tailed image classification datasets demonstrated the empirical effectiveness of PASCL.

Acknowledgement

Z.W. is supported by the U.S. Army Research Laboratory Cooperative Research Agreement W911NF17-2-0196 (IOBT REIGN) and an Amazon Research Award.

References

Benz, P., Zhang, C., Karjauv, A., and Kweon, I. S. Revisiting batch normalization for improving corruption robustness. In *WACV*, pp. 494–503, 2021.

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

Chang, W.-G., You, T., Seo, S., Kwak, S., and Han, B. Domain-specific batch normalization for unsupervised domain adaptation. In *CVPR*, pp. 7354–7362, 2019.

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

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

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

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

Fort, S., Ren, J., and Lakshminarayanan, B. Exploring the limits of out-of-distribution detection. In *NeurIPS*, 2021.

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

- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *CVPR*, pp. 770–778, 2016.
- Hein, M., Andriushchenko, M., and Bitterwolf, J. Why ReLU networks yield high-confidence predictions far away from the training data and how to mitigate the problem. In *CVPR*, pp. 41–50, 2019.
- Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *ICLR*, 2017.
- Hendrycks, D., Mazeika, M., and Dietterich, T. Deep anomaly detection with outlier exposure. In *ICLR*, 2019.
- Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., and Kalantidis, Y. Decoupling representation and classifier for long-tailed recognition. In *ICLR*, 2020.
- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., and Krishnan, D. Supervised contrastive learning. In *NeurIPS*, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Master’s thesis, University of Toronto, 2009.
- Le, Y. and Yang, X. Tiny imagenet visual recognition challenge. 2015. URL <http://cs231n.stanford.edu/tiny-imagenet-200.zip>.
- Lee, K., Lee, K., Lee, H., and Shin, J. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *NeurIPS*, 2018.
- Liang, S., Li, Y., and Srikant, R. Enhancing the reliability of out-of-distribution image detection in neural networks. In *ICLR*, 2018.
- Liu, W., Wang, X., Owens, J., and Li, Y. Energy-based out-of-distribution detection. In *NeurIPS*, 2020.
- Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., and Yu, S. X. Large-scale long-tailed recognition in an open world. In *CVPR*, pp. 2537–2546, 2019.
- Loshchilov, I. and Hutter, F. SGDR: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- Meinke, A. and Hein, M. Towards neural networks that provably know when they don’t know. In *ICLR*, 2020.
- Menon, A. K., Jayasumana, S., Rawat, A. S., Jain, H., Veit, A., and Kumar, S. Long-tail learning via logit adjustment. In *ICLR*, 2021.
- Mohseni, S., Pitale, M., Yadawa, J., and Wang, Z. Self-supervised learning for generalizable out-of-distribution detection. In *AAAI*, pp. 5216–5223, 2020.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. 2011.
- Pang, T., Xu, K., Dong, Y., Du, C., Chen, N., and Zhu, J. Rethinking softmax cross-entropy loss for adversarial robustness. In *ICLR*, 2020.
- Papadopoulos, A.-A., Rajati, M. R., Shaikh, N., and Wang, J. Outlier exposure with confidence control for out-of-distribution detection. *Neurocomputing*, 441:138–150, 2021.
- Roy, A. G., Ren, J., Azizi, S., Loh, A., Natarajan, V., Mustafa, B., Pawlowski, N., Freyberg, J., Liu, Y., Beaver, Z., et al. Does your dermatology classifier know what it doesn’t know? Detecting the long-tail of unseen conditions. *Medical Image Analysis*, 75:102274, 2022.
- Sastry, C. S. and Oore, S. Detecting out-of-distribution examples with Gram matrices. In *ICML*, pp. 8491–8501, 2020.
- Sohn, K., Li, C.-L., Yoon, J., Jin, M., and Pfister, T. Learning and evaluating representations for deep one-class classification. In *ICLR*, 2021.
- Tang, K., Miao, D., Peng, W., Wu, J., Shi, Y., Gu, Z., Tian, Z., and Wang, W. CODEs: Chamfer out-of-distribution examples against overconfidence issue. In *ICCV*, pp. 1153–1162, 2021.
- Torralba, A., Fergus, R., and Freeman, W. T. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE TPAMI*, 30(11):1958–1970, 2008.
- Van Horn, G. and Perona, P. The devil is in the tails: Fine-grained classification in the wild. *arXiv preprint arXiv:1709.01450*, 2017.
- Wang, H., Liu, W., Bocchieri, A., and Li, Y. Can multi-label classification networks know what they don’t know? In *NeurIPS*, 2021a.
- Wang, X., Lian, L., Miao, Z., Liu, Z., and Yu, S. Long-tailed recognition by routing diverse distribution-aware experts. In *ICLR*, 2021b.
- Winkens, J., Bunel, R., Roy, A. G., Stanforth, R., Natarajan, V., Ledsam, J. R., MacWilliams, P., Kohli, P., Karthikesalingam, A., Kohl, S., et al. Contrastive training for improved out-of-distribution detection. *arXiv preprint arXiv:2007.05566*, 2020.

- Xiang, L., Ding, G., and Han, J. Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In *ECCV*, pp. 247–263. Springer, 2020.
- Xie, C. and Yuille, A. Intriguing properties of adversarial training at scale. In *ICLR*, 2019.
- Yang, D., Shin, I., Kien, M. N., Kim, H., Yu, C., and Hwang, M. Out-of-distribution detection based on distance metric learning. In *International Conference on Smart Media and Applications*, pp. 214–218, 2020.
- Yang, J., Wang, H., Feng, L., Yan, X., Zheng, H., Zhang, W., and Liu, Z. Semantically coherent out-of-distribution detection. In *ICCV*, pp. 8301–8309, 2021.
- Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., and Xiao, J. LSUN: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
- Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., and Torralba, A. Places: A 10 million image database for scene recognition. *IEEE TPAMI*, 40(6):1452–1464, 2017.
- Zhou, B., Cui, Q., Wei, X.-S., and Chen, Z.-M. BBN: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In *CVPR*, pp. 9719–9728, 2020.
- Zhou, W., Liu, F., and Chen, M. Contrastive out-of-distribution detection for pretrained transformers. In *EMNLP*, 2021a.
- Zhou, Z., Guo, L.-Z., Cheng, Z., Li, Y.-F., and Pu, S. STEP: Out-of-distribution detection in the presence of limited in-distribution labeled data. In *NeurIPS*, 2021b.

A. More Details on Experimental Settings

In this section, we provide more details on experimental settings in addition to those in Section 5.1. For experiments on CIFAR10-LT and CIFAR100-LT, we train the main branch (i.e., stage 1 in Algorithm 1) for 200 epochs using Adam (Kingma & Ba, 2014) optimizer with initial learning rate 1×10^{-3} and batch size 256. We decay the learning rate to 0 using a cosine annealing learning rate scheduler (Loshchilov & Hutter, 2016). For auxiliary branch finetuning (i.e., stage 2 in Algorithm 1), we finetune the auxiliary branch for 3 epochs using Adam optimizer with initial learning rate 5×10^{-4} . Other hyper-parameters are the same as in main branch training. For experiments on ImageNet-LT, we follow the settings in (Wang et al., 2021b). Specifically, we train the main branch for 100 epochs using SGD optimizer with initial learning rate 0.1 and batch size 256. We decay the learning rate by a factor of 10 at epoch 60 and 80. For auxiliary branch finetuning, we finetune the auxiliary branch for 3 epochs using SGD optimizer with initial learning rate 0.01, which is decayed by a factor of 10 after each finetune epoch. On all datasets, we set $\tau = 0.1$ following (Khosla et al., 2020), $\lambda_1 = 0.5$ following (Hendrycks et al., 2019), and empirically set $\lambda_2 = 0.1$ for PASCL. The total training epochs are kept the same across all compared methods. For other hyper-parameters in the baseline methods, we use the suggested values in the original papers.

B. Feature Visualization on CIFAR10 and CIFAR10-LT

In Figure 1 (in Section 3), we showed the feature distribution of OE trained on CIFAR10-LT. In this section, we compare the feature distribution of OE trained on the balanced CIFAR10 (Figure 4 (a)) and CIFAR10-LT (Figure 4 (b)). As we mentioned in Section 3, we can see that, on the balanced CIFAR10 dataset, the OOD samples are well separated from all in-distribution classes. In contrast, when trained on CIFAR10-LT, the features of tail-class in-distribution (e.g., horse, ship) samples and OOD samples heavily overlap, while those of head-class in-distribution (e.g., airplane, automobile, bird) samples are well separated from OOD samples.

C. More Experimental Results

C.1. More Ablation Study

On imbalance ratio ρ In Section 5.2, we use imbalance ratio $\rho = 100$ on both CIFAR10-LT and CIFAR100-LT. In this section, we show that our method can work well under different imbalance ratios. Specifically, we conduct experiments on CIFAR10-LT with $\rho = 50$. The results are shown in Table 10. Our method also outperforms the OE baseline by a considerable margin when $\rho = 50$.

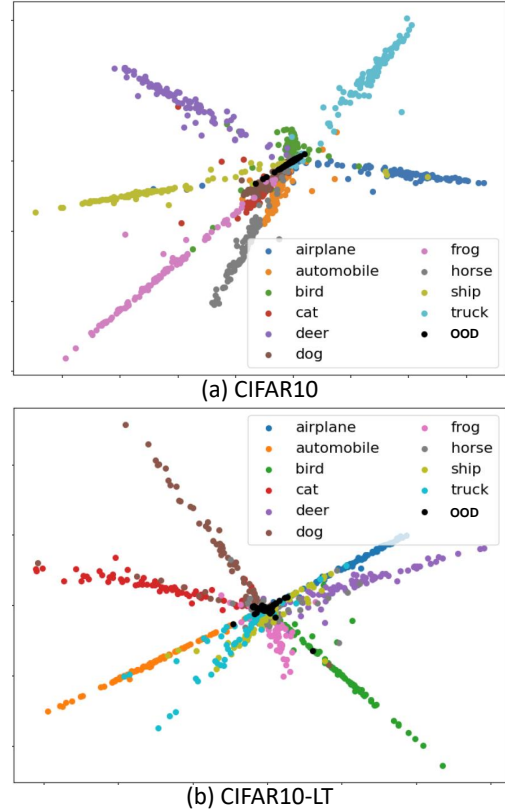


Figure 4. Visualization of in-distribution and OOD samples in the feature space. We visualize the penultimate layer features of ResNet18 trained by OE on (a) CIFAR10 and (b) CIFAR10-LT, using the same visualization method as in (Pang et al., 2020). 100 test samples from each in-distribution class and 100 OOD samples are visualized. The model trained on CIFAR10-LT struggles to distinguish tail in-distribution classes like “horse” (grey) and “ship” (yellow) from OOD samples (black). Note that in CIFAR10-LT, the number of training samples in each class follows the alphabetic order: the “airplane” class has the most training samples and the “truck” class has the least.

On model structures In Section 5.2, we use the standard ResNet18 as the backbone model. In this section, we show that our method can work well under different model structures, but conducting experiments using the standard ResNet34 (He et al., 2016). The results are shown in Table 11. Our method also outperforms the OE baseline by a considerable margin on ResNet34.

Ablation study on λ_2 The results are shown in Table 12. The performance of our method is stable with respect to different λ_2 values.

C.2. Combining OE with the Original Supervised Contrastive Learning on \mathcal{D}_{in} or $\mathcal{D}_{in} \cup \mathcal{D}_{out}$

In Section 4, we mentioned that when combined with OE, the original supervised contrastive learning achieves slightly

Table 10. Results on CIFAR10-LT ($\rho = 50$) using ResNet18.

\mathcal{D}_{out}^{test}	Method	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)	ACC (\uparrow)
Average	OE	93.13	91.06	24.73	83.34
	Ours	93.94	92.79	22.80	85.44

Table 11. Results on CIFAR10-LT ($\rho = 100$) using ResNet34.

\mathcal{D}_{out}^{test}	Method	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)	ACC (\uparrow)
Average	OE	89.86	87.28	33.66	73.39
	Ours	91.11	89.28	33.21	75.34

better OOD detection performance if applied on $\mathcal{D}_{in} \cup \mathcal{D}_{out}$ than on \mathcal{D}_{in} . In this section, we provide experimental results to support this. Specifically, we compare the two methods on CIFAR10-LT and CIFAR100-LT, and show the results in Table 13. As we can see, these two methods achieve almost identical performance, although those on $\mathcal{D}_{in} \cup \mathcal{D}_{out}$ are only slightly better. For that reason, we use the second method (i.e., supervised contrastive learning on $\mathcal{D}_{in} \cup \mathcal{D}_{out}$) as “the original SCL” baseline in Table 9 (i.e., the second rows in each block in Table 9).

C.3. OLTR Results

As discussed in Section 2, OLTR (Liu et al., 2019) focuses on in-distribution classification in long-tailed recognition instead of OOD detection. Although it achieves better OOD detection performance than naive baselines such as MSP, it is not comparable with state-of-the-art OOD detection methods such as OE. In contrast, our paper focuses on improve OOD detection performance on long-tailed training sets. The results of OLTR and other methods (including our PASCL) on ImageNet-LT are compared in Table 14.

D. More Related Works

In this section, we discuss more related works along with those in Section 2.

OOD detection ODIN (Liang et al., 2018) improved MSP by adding temperature scaling and adversarial attacks to enlarge the differences in MSP between in-distribution and OOD data. Hein et al. (2019) theoretically explained why piecewise linear models tend to give high confident predictions on OOD data, and proposed a new robust learning method to prevent such overconfidence issue. Meinke & Hein (2020) proposed Certified Certain Uncertainty (CCU), a robust learning method which achieves provably low confidence predictions on OOD samples. Self-Supervised Feature Learning (SOFL) (Mohseni et al., 2020) proposed to classify all OOD samples into several abstaining classes, where the pseudo labels for OOD training samples are assigned in a self-supervised manner. Tang et al. (2021) proposed to generate Chamfer OOD samples (i.e., the OOD

Table 12. Ablation study results on λ_2 . CIFAR100-LT is used as in-distribution training set. SVHN is used as OOD test set.

λ_2	0	0.05	0.1	0.5
AUPR (\uparrow)	86.82	88.12	88.49	88.46

Table 13. Combining OE with the original supervised contrastive learning (SCL) on \mathcal{D}_{in} and $\mathcal{D}_{in} \cup \mathcal{D}_{out}$. Experiments are conducted on CIFAR10/100-LT (both with $\rho = 100$) using ResNet18. SVHN is used as \mathcal{D}_{out}^{test} .

\mathcal{D}_{in}	Method	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)
CIFAR10-LT	OE + SCL on \mathcal{D}_{in}	95.30 \pm 1.42	97.19 \pm 1.02	15.08 \pm 2.54
	OE+SCL on $\mathcal{D}_{in} \cup \mathcal{D}_{out}$	95.34 \pm 1.58	97.30 \pm 1.20	15.12 \pm 3.07
CIFAR100-LT	OE + SCL on \mathcal{D}_{in}	78.01 \pm 1.95	87.16 \pm 0.62	59.36 \pm 4.52
	OE+SCL on $\mathcal{D}_{in} \cup \mathcal{D}_{out}$	78.05 \pm 2.12	87.18 \pm 0.87	59.10 \pm 5.03

Table 14. Comparing OLTR with our PASCL and other methods. Experiments are conducted on ImageNet-LT using ResNet50.

Method	AUROC (\uparrow)	AUPR (\uparrow)	FPR95 (\downarrow)
ST (MSP)	53.81	51.63	90.15
OLTR	55.68	54.02	90.02
OE	<u>66.33</u>	<u>68.29</u>	<u>88.22</u>
Ours	68.00	70.15	87.53

samples that are close to the in-distribution samples), when there is no OOD samples directly available for training. Fort et al. (2021) observed that large-scale vision transformers achieves significantly better OOD detection performance than CNNs. They further proposed a simple but surprisingly effective OOD detection method based on multi-model transformers. Zhou et al. (2021b) studied a new problem setting where the number of labeled in-distribution samples are limited. Wang et al. (2021a) studied the OOD detection problem in multi-label classification.

Long-tailed recognition Cui et al. (2019) proposed a loss reweighing method based on the effective numbers of training samples in each class. BBN (Zhou et al., 2020) used two separate branches which focus on learning universal and tail-class-specific features, respectively. Xiang et al. (2020) split the unbalanced training set into several relatively balanced subsets, and then trained a separate model on each subset, whose knowledge is later distilled into a single model. RIDE (Wang et al., 2021b) used mixture of experts and a novel dynamic expert routing module to largely boost long-tailed recognition performance.