

# Decoupled Contrastive Learning for Long-Tailed Recognition

Shiyu Xuan, Shiliang Zhang

National Key Laboratory for Multimedia Information Processing  
School of Computer Science, Peking University, Beijing, China  
shiyu\_xuan@stu.pku.edu.cn, slzhang.jdl@pku.edu.cn

## Abstract

Supervised Contrastive Loss (SCL) is popular in visual representation learning. Given an anchor image, SCL pulls two types of positive samples, *i.e.*, its augmentation and other images from the same class together, while pushes negative images apart to optimize the learned embedding. In the scenario of long-tailed recognition, where the number of samples in each class is imbalanced, treating two types of positive samples equally leads to the biased optimization for intra-category distance. In addition, similarity relationship among negative samples, that are ignored by SCL, also presents meaningful semantic cues. To improve the performance on long-tailed recognition, this paper addresses those two issues of SCL by decoupling the training objective. Specifically, it decouples two types of positives in SCL and optimizes their relations toward different objectives to alleviate the influence of the imbalanced dataset. We further propose a patch-based self distillation to transfer knowledge from head to tail classes to relieve the under-representation of tail classes. It uses patch-based features to mine shared visual patterns among different instances and leverages a self distillation procedure to transfer such knowledge. Experiments on different long-tailed classification benchmarks demonstrate the superiority of our method. For instance, it achieves the 57.7% top-1 accuracy on the ImageNet-LT dataset. Combined with the ensemble-based method, the performance can be further boosted to 59.7%, which substantially outperforms many recent works. Our code will be released.

## Introduction

Thanks to the powerful deep learning methods, the performance of various vision tasks (Russakovsky et al. 2015; Long, Shelhamer, and Darrell 2015) on manually balanced dataset has been significantly boosted. In real-world applications, training samples commonly exhibit a long-tailed distribution, where a few head classes contribute most of the observations, while lots of tail classes are associated with only a few samples (Van Horn et al. 2018). Long-tailed distribution leads to two challenges for visual recognition: (a) the loss function designed for the balanced dataset can be easily biased toward the head classes. (b) each of tail classes contains too few samples to represent visual variances, leading to the under-representation of the tail classes.

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

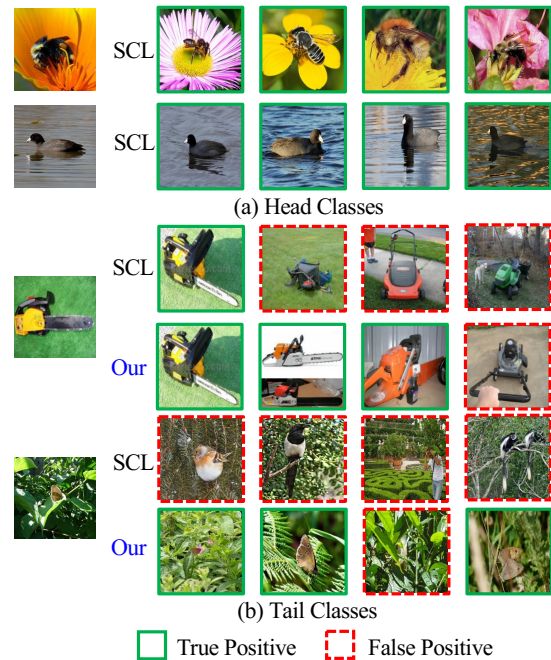


Figure 1: Examples of retrieval results using features learned by SCL on head classes in (a) and tail classes in (b). In (b), features learned by SCL are biased to low-level appearance cues, while features learned by our method are more discriminative to semantic cues.

By optimizing the intra-inter category distance, Supervised Contrastive Loss (SCL) (Khosla et al. 2020) has achieved impressive performance on balanced datasets. Given an anchor image, SCL pulls two kinds of positive samples together, *i.e.*, (a) different views of the anchor image generated by the data augmentation, and (b) other images from the same classes. Those two types of positives supervise the model to learn different representations, *i.e.*, images from the same categories enforce the learning of semantic cues, while samples augmented by appearance variances mostly lead to the learning of low-level appearance cues. Fig. 1 (a) shows that, SCL effectively learns semantic features for head classes, *e.g.*, the learned semantic “bee” is robust to cluttered backgrounds.

As shown in Fig. 1 (b), representations learned by SCL for tail classes are more discriminative to low-level appearance cues like shape, texture, and color.

Our theoretical analysis in Section Overview indicates that, SCL poses imbalanced gradients on two kinds of positive samples, resulting a biased optimization for head and tail classes. We hence propose the Decoupled Contrastive Learning, which adopts the Decoupled Supervised Contrastive Loss (DSCL) to handle this issue. Specifically, DSCL decouples two kinds of positive samples to re-formulate the optimization of intra-category distance. It alleviates the imbalanced gradients of two kinds of positive samples. We also give a theoretical proof that DSCL prevents the learning of a biased intra-category distance. In Fig. 1 (b), features learned by our method are discriminative to semantic cues, and substantially boost the retrieval performance on tail classes.

To further alleviate the challenge of long-tailed distribution, we propose the Patch-based Self Distillation (PBSD) to leverage head classes to facilitate the representation learning in tail classes. PBSD adopts a self distillation strategy to better optimize the inter-category distance, through mining shared visual patterns between different classes and transferring knowledge from head to tail classes. We introduce patch-based features to represent visual patterns from an object. The similarity between patch-based features and instance-level features is calculated to mine the shared visual patterns, *i.e.*, if an instance shares visual patterns with a patch-based feature, they will have high similarity. We leverage the self distillation loss to maintain the similarity relationship among samples, and integrate the knowledge into the training.

DSCL and PBSD are easy to implement, and substantially boosts the long-tailed recognition performance. We evaluate our method on several long-tailed datasets including ImageNet-LT (Liu et al. 2019), iNaturaList 2018 (Van Horn et al. 2018), and Places-LT (Liu et al. 2019). Experimental results show that our method improves the SCL by 6.5% and achieves superior performance compared with recent works. For example, it outperforms a recent contrastive learning based method TSC (Li et al. 2021) by 5.3% on ImageNet-LT. Our method can be flexibly combined with ensemble-based methods like RIDE (Wang et al. 2020), which achieves the overall accuracy of 74.9% on the iNaturaList 2018, outperforming the recent work CR (Ma et al. 2023) by 1.4% in overall accuracy.

To the best of our knowledge, this is an original contribution decoupling two kinds of positives and using patch-based self distillation to boost the performance of SCL on long-tail recognition. The proposed DSCL decouples different types of positive samples to pursue a more balanced intra-category distance optimization across head and tail classes. It also introduces the similarity relationship cues to leverage shared patterns in head classes for the optimization in tail classes. Extensive experiments on three commonly used datasets have shown its promising performance. Our method is easy to implement and the code will be released to benefit the future research on long-tailed visual recognition.

## Related Work

**Long-tailed recognition** aims to address the problem of the model training in the situation where a small portion of classes have massive samples but the others are associated with only a few samples. Current research can be summarized into four categories, *e.g.*, re-balancing methods, decoupling methods, transfer learning methods and ensemble-based methods, respectively.

Re-balancing methods use re-sampling or re-weighting to deal with long-tailed recognition. Re-sampling methods typically include over-sampling for the tail classes (Byrd and Lipton 2019) or under-sampling for the head classes (Japkowicz and Stephen 2002). Besides re-sampling, re-weighting the loss function is also an effective solution. For instance, Balanced-Softmax (Ren et al. 2020) presents the unbiased extension of Softmax based on the Bayesian estimation. Re-balancing methods could be harmful to the discriminative power of the learned backbone (Kang et al. 2019). Therefore, decoupling methods propose the two-stage training to decouple the representation learning and the classifier training. Transfer learning methods enhance the performance of the model by transferring knowledge from head classes to tail classes. BatchFormer (Hou, Yu, and Tao 2022) introduces a one-layer Transformer (Vaswani et al. 2017) to transfer knowledge by learning the sample relationship from each mini-batch. Ensemble-based methods leverage multi experts to solve long-tailed visual learning. RIDE (Wang et al. 2020) proposes a multi-branch network to learn diverse classifiers in parallel. Although ensemble-based method achieves superior performance, the introduction of multi experts increases the number of parameters and computational complexity.

**Contrastive learning** has received much attention because of its superior performance on representation learning (He et al. 2020). Contrastive learning aims to find a feature space that can encode the semantic similarities by pulling the positive pairs together while pushing negative pairs apart. Some researchers have leveraged contrastive learning in the long-tailed recognition. For example, KCL (Kang et al. 2020) finds that the self-supervised learning based on contrastive learning can learn a balanced feature space. To leverage the useful label information, they extend SCL (Khosla et al. 2020) by introducing a k-positive sampling method. TSC (Li et al. 2021) improves the uniformity of the feature distribution by making features of different classes converge to a pre-defined uniformly distributed targets. Some methods (Yun et al. 2022; Zhang et al. 2023) extend contrastive learning with localized information to benefit the dense prediction tasks.

This work differs with previous ones in several aspects. Existing long-tailed recognition works using contrastive learning treat two kinds of positives equally. To the best of our knowledge, this is an early work revealing that treating two kinds of positives equally leads to a biased optimization across categories. We hence propose a decoupled supervised contrastive loss to pursue a balanced intra-category distance optimization. We further extend the contrastive learning by introducing patch-based self distillation to transfer knowledge between classes, mitigating the under-representation of the tailed classes and leading to a more effective optimization to inter-category distance. Different from other transfer




Figure 2: Average ratio of gradient L2 norm computed by pulling the anchor with two types of positives as in Eq. (6) on ImageNet-LT. ‘\*’ denotes the theoretical ratio. SCL treats two types of positives equally, and leads to the imbalanced optimization. Two types of positives denote the data augmentation and other images in the same category.

learning methods, PBSO leverages patch-based features to mine shared patterns between different classes and designs a self distillation procedure to transfer knowledge. The self distillation procedure does not rely on a large teacher model or multi-expert models (Li et al. 2022), making it efficient. Compared with patch-based contrastive learning methods that only mine similar patches from different views of an image, PBSO transfers knowledge between different images. Those differences and the promising performance in extensive experiments highlight the contribution of this work.

## Methodology

### Analysis of SCL

Given a training dataset  $\mathcal{D} = \{x_i, y_i\}_{i=1}^n$ , where  $x_i$  denotes an image and  $y_i \in \{1, \dots, K\}$  is its class label. Assuming that  $n^k$  denotes the cardinality of class  $k$  in  $\mathcal{D}$ , and indexes of classes are sorted by cardinality in decreasing order, *i.e.*, if  $a < b$ , then  $n^a \geq n^b$ . In long-tailed recognition, the training dataset is imbalanced, *i.e.*,  $n^1 \gg n^K$ , and the imbalance ratio is defined as  $n^1/n^K$ . For the image classification task, algorithm aims to learn a feature extraction backbone  $v_i = f_\theta(x_i)$ , and a linear classifier, which first maps an image  $x_i$  into a global feature map  $u_i$  and uses the global pooling to get a  $d$ -dim feature vector  $v_i$ . It hence classifies the feature vector to a  $K$ -dim classification score. Typically, the testing dataset  $\mathcal{T}$  is balanced.

Supervised Contrastive Learning (SCL) is commonly adopted to learn the feature extraction backbone. Given an anchor image  $x_i$ , defining  $z_i = g_\gamma(v_i)$  as the normalized feature extracted with the backbone and an extra projection head  $g_\gamma$  (He et al. 2020),  $z_i^+$  as the normalized feature of a positive sample of  $x_i$  generated by data augmentation. We use  $M$  to denote a set of sample features that can be acquired by the memory queue (He et al. 2020), and use  $P_i$  as the positive feature set of  $x_i$  drawn from  $M$  with  $P_i = \{z_t \in M : y_t = y_i\}$ .

SCL decreases the intra-category distance by pulling the anchor image and its positive samples together, meanwhile enlarges the inter-category distance through pushing images with different class labels apart, *i.e.*,

$$\mathcal{L}_{scl} = \frac{-1}{|P_i| + 1} \sum_{z_t \in \{z_i^+ \cup P_i\}} \log p(z_t | z_i), \quad (1)$$

where  $|P_i|$  is the cardinality of  $P_i$ . Using  $\tau$  to denote a pre-defined temperature parameter, the conditional probability  $p(z_t | z_i)$  is computed as,

$$p(z_t | z_i) = \frac{\exp(z_t \cdot z_i / \tau)}{\sum_{z_m \in \{z_i^+ \cup M\}} \exp(z_m \cdot z_i / \tau)}. \quad (2)$$

Eq. (1) can be formulated as a distribution alignment task,

$$\mathcal{L}_{align} = \sum_{z_t \in \{z_i^+ \cup M\}} -\hat{p}(z_t | z_i) \log p(z_t | z_i), \quad (3)$$

where  $\hat{p}(z_t | z_i)$  is the probability of the target distribution. For  $z_i^+$  and  $z_t \in P_i$ , SCL treats them equally as positive samples and sets their target probability as  $1/(|P_i| + 1)$ . For other images with different class labels in  $M$ , SCL treats them as negative samples and sets their target probability as 0.

For the feature  $z_i$  of an anchor image  $x_i$ , the gradient of SCL is,

$$\begin{aligned} \frac{\partial \mathcal{L}_{scl}}{\partial z_i} = \frac{1}{\tau} \left\{ \sum_{z_j \in N_i} z_j p(z_j | z_i) + z_i^+ \left( p(z_i^+ | z_i) - \frac{1}{|P_i| + 1} \right) \right. \\ \left. + \sum_{z_t \in P_i} z_t \left( p(z_t | z_i) - \frac{1}{|P_i| + 1} \right) \right\}, \end{aligned} \quad (4)$$

where  $N_i$  is the negative set of  $x_i$  containing features drawn from  $\{z_j \in M : y_j \neq y_i\}$ . SCL involves two types of positive samples  $z_i^+$  and  $z_t \in P_i$ . We compute gradients of pulling the anchor with two types of positive samples as,

$$\begin{aligned} \frac{\partial \mathcal{L}_{scl}}{\partial z_i} \Big|_{z_i^+} &= z_i^+ \left( p(z_i^+ | z_i) - \frac{1}{|P_i| + 1} \right), \\ \frac{\partial \mathcal{L}_{scl}}{\partial z_i} \Big|_{z_t} &= z_t \left( p(z_t | z_i) - \frac{1}{|P_i| + 1} \right), z_t \in P_i. \end{aligned} \quad (5)$$

At the beginning of the training, the ratio of gradient L2 norm of two kinds of positive samples is,

$$\frac{\left\| \frac{\partial \mathcal{L}_{scl}}{\partial z_i} \Big|_{z_i^+} \right\|_2}{\sum_{z_t \in P_i} \left\| \frac{\partial \mathcal{L}_{scl}}{\partial z_i} \Big|_{z_t} \right\|_2} \approx \frac{1}{|P_i|}. \quad (6)$$

When SCL converges, the optimal conditional probability of  $z_i^+$  is,

$$p(z_i^+ | z_i) = \frac{1}{|P_i| + 1}. \quad (7)$$

A detailed proof of above computations can be found in the Supplementary Material.

In SCL, the memory queue  $M$  is uniformly sampled from the training set, which leads to  $|P_i| \approx \frac{n^{y_i}}{n} |M|$ . In a balanced

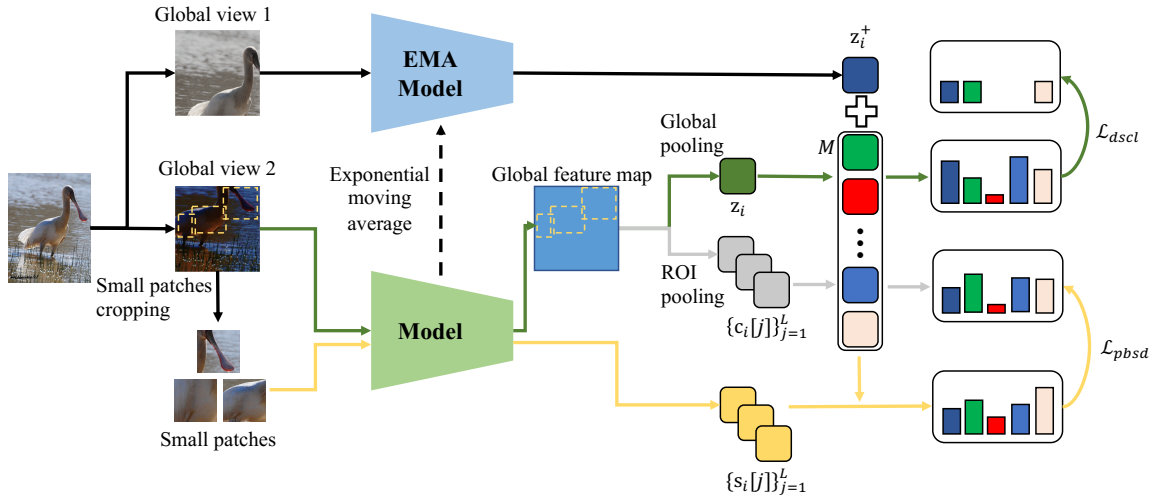


Figure 3: Illustration of the proposed method. Data augmentation is performed to get two global views of a training image. Then small patch is cropped from the global view. The backbone and Exponential Moving Average (EMA) backbone (He et al. 2020) are used to extract normalized features. These features are used to calculate the similarity distribution with memory queue  $M$ .  $\mathcal{L}_{dscl}$  optimizes the feature space by pulling the anchor image with its positive samples together and pushing the anchor image with its negative samples apart.  $\mathcal{L}_{pbsd}$  transfers knowledge through mimicking two similarity distributions.

dataset,  $n^1 \approx n^2 \approx \dots \approx n^K$ , resulting a balanced  $|P_i|$  across different categories. For a long-tail dataset with imbalanced  $|P_i|$ , SCL makes the head classes pay more attention to pulling the anchor  $z_i$  with features from  $P_i$  together as the gradient is dominated by the third term in Eq. (4).

As shown in Fig. 2, the ratio of gradient L2 norm of pulling two kinds of positive samples is unbalanced. When the training of SCL converges, the optimal value of  $p(z_i^+ | z_i)$  is also influenced by the  $|P_i|$  as shown in Eq. (7). The inconsistency of learned features across categories is illustrated in Fig. 1(a) and (b). This phenomena has also been validated by (Wei et al. 2020) that pulling  $z_i$  with  $z_i^+$  and samples from  $P_i$  leads to learning different representations, *i.e.*, appearance features for tail classes and semantic features for head classes, respectively.

Eq. (4) also indicates that, SCL equally pushes away all the negative samples to enlarges the inter-category distance. This strategy ignores the valuable similarity cues among different classes. To seek a better way to optimize intra and inter category distance, we propose the Decoupled Supervised Contrastive Loss (DSCL) to decouple two kinds of positive samples to prevent the biased optimization, as well as the Patch-based Self Distillation (PBSD) to leverage similarity cues among classes.

### Decoupled Supervised Contrastive Loss

DSCL is proposed to ensure a more balanced optimization to the intra-category distance across different categories. It decouples two kinds of positive samples and add different weights on them to make the gradient L2 norm ratio and the optimal value of  $p(z_i^+ | z_i)$  not influenced by the number of

samples in each category. We represent the DSCL as,

$$\mathcal{L}_{dscl} = \frac{-1}{|P_i| + 1} \sum_{z_t \in \{z_i^+ \cup P_i\}} \log \frac{\exp w_t(z_t \cdot z_i / \tau)}{\sum_{z_m \in \{z_i^+ \cup M\}} \exp(z_m \cdot z_i / \tau)}, \quad (8)$$

where,

$$w_t = \begin{cases} \alpha(|P_i| + 1), & z_t = z_i^+ \\ \frac{(1 - \alpha)(|P_i| + 1)}{|P_i|}, & z_t \in P_i \end{cases} \quad (9)$$

where  $\alpha \in [0, 1]$  is a pre-defined hyper-parameter. The proposed DSCL is a generalization of SCL in both balanced setting and imbalanced setting. If the dataset is balanced, DSCL is the same as SCL by setting  $\alpha = 1/(|P_i| + 1)$ .

We proceed to show why Eq. (8) leads to a more balanced optimization.

At the beginning of the training, the gradient L2 norm ratio of two kinds of positives is,

$$\frac{\left\| \frac{\partial \mathcal{L}_{dscl}}{\partial z_i} \Big|_{z_i^+} \right\|_2}{\sum_{z_t \in P_i} \left\| \frac{\partial \mathcal{L}_{dscl}}{\partial z_i} \Big|_{z_t} \right\|_2} \approx \frac{\alpha}{1 - \alpha}. \quad (10)$$

When DSCL converges, the optimal conditional probability of  $z_i^+$  is  $p(z_i^+ | z_i) = \alpha$ , where a detailed proof can be found in the Supplementary Material.

As shown in Eq. (10) and Fig. 2, the gradient ratio of two kinds of positive samples is not influenced by  $|P_i|$ . DSCL also ensures that the optimal value of  $p(z_i^+ | z_i)$  is not influenced by  $|P_i|$ , hence alleviates the inconsistent feature learning issue between head and tail classes.

### Patch-based Self Distillation

Visual patterns can be shared among different classes, *e.g.*, the visual pattern ‘wheel’ is shared by the class ‘truck’, ‘car’,

and ‘bus’. Features of many visual patterns in tail classes can be learned from head classes that share these visual patterns, hence reducing the difficult of representation learning in tail classes. SCL pushes two instances from different classes away in the feature space, even they share meaningful visual patterns. As shown in Fig. 4, we extract query patch features from yellow bounding boxes and retrieve the top-3 similar samples from the dataset. Retrieval results of SCL denoted by ‘w/o PBSB’ are not semantically related to the query patch, indicating that SCL is not effective in learning and leveraging patch-level semantic cues.

Inspired by the patch-based methods in fine-grained image recognition (Zhang et al. 2014; Quan et al. 2019; Sun et al. 2018), we introduce patch-based features to encode visual patterns. Given the global feature map  $u_i$  of an image  $x_i$  extracted by the backbone, we first randomly generate some patch boxes. The coordinates of those patch boxes denote  $\{B_i[j]\}_{j=1}^L$ , where  $L$  is the number of the patch boxes. We hence apply ROI pooling (He et al. 2017) according to the coordinates of those patch boxes and send pooled features into a projection head to get the normalized embedding features  $\{c_i[j]\}_{j=1}^L$  with

$$c_i[j] = g_\gamma(\text{ROI}(u_i, B_i[j])). \quad (11)$$

Similar to Eq. (2), the conditional probability is leveraged to calculate the similarity relationship between instances,

$$p(z_t|c_i^j) = \frac{\exp(z_t \cdot c_i[j]/\tau)}{\sum_{z_m \in \{z_i^+ \cup M\}} \exp(z_m \cdot c_i[j]/\tau)}. \quad (12)$$

If the image corresponding to  $z_t$  has shared visual patterns with the patch-based features,  $z_t$  and  $c_i[j]$  will have a high similarity. Therefore, Eq. (12) encodes the similarity cues between each pair of instances.

We use the similarity cues as the knowledge to supervise the training procedure. To maintain such knowledge, we also crop multi image patches from the image according to  $\{B_i[j]\}_{j=1}^L$ , and extract their feature embeddings  $\{s_i[j]\}_{j=1}^L$  with the backbone,

$$s_i[j] = g_\gamma(f_\theta(\text{Crop}(x_i, B_i[j]))). \quad (13)$$

PBSB enforces the feature embeddings of image patches to produce the same similarity distribution as the patch-based features via the following loss,

$$\mathcal{L}_{pbsd} = \frac{1}{L} \sum_{j=1}^L \sum_{z_t \in \{z_i^+ \cup M\}} -p(z_t|c_i[j]) \log p(z_t|s_i[j]), \quad (14)$$

note that  $p(z_t|c_i[j])$  is detached from the computation graph to block the gradient.

Local visual patterns of an object can be shared by different categories. We hence use patch-based features to represent visual patterns.  $p(z_t|c_i[j])$  is calculated to mine relationship of the shared patterns among images. Minimizing Eq. (14) maintains shared patterns to transfer knowledge and mitigate the under-representation of the tailed classes. The retrieval results shown in Fig. 4 indicate that our method effectively reinforces the learning of patch-level features and patch-to-image similarity, making it possible to mine shared visual

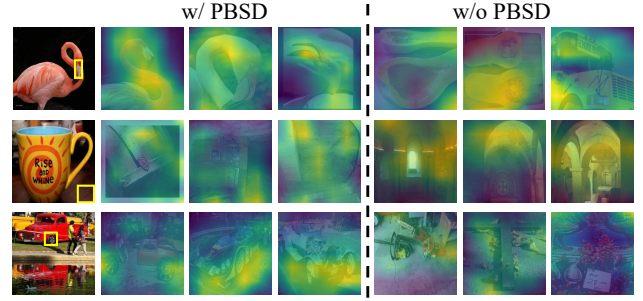


Figure 4: Patch-based image retrieval results (top 3 returned) on ImageNet-LT. Query patches are highlighted with yellow bounding boxes. The response map of query patch features on the retrieved images is also illustrated.

patterns of different classes. Experiments also validate that PBSB loss is important to the performance gain.

Multi-crop trick (Caron et al. 2020) is commonly used in self-supervised learning to generate more augmented samples of the anchor image. It introduces low resolution crops to reduce the computational complexity. Our motivation and loss design are different with the multi-crop strategy. PBSB is motivated to leverage shared patterns between head and tail classes to assist the learning of the tail classes. Patch-based features are obtained with ROI pooling to represent shared patterns. Eq. (14) performs self distillation to maintain shared patterns. We conduct an experiment by replacing PBSB with the multi-crop trick. As shown in Table 1, the performance on ImageNet-LT drops from 57.7% to 56.1%, indicating that PBSB is more effective than the multi-crop strategy.

## Training Pipeline

We illustrate our method in Fig. 3. To maintain a memory queue, we use a momentum updated model as in (He et al. 2020). The training is supervised by two losses, *i.e.*, the decoupled supervised contrastive loss and the patch-based self distillation loss. The overall training loss is denoted as,

$$\mathcal{L}_{overall} = \mathcal{L}_{dscl} + \lambda \mathcal{L}_{pbsd}, \quad (15)$$

where  $\lambda$  is the loss weight.

Our method focuses on the representation learning, and can be applied in different tasks by concatenating their losses. Following (Li et al. 2021; Kang et al. 2020), after the training of the backbone, we discard the learned projection head  $g_\gamma(\cdot)$  and train a linear classifier on top of the learned backbone using the standard cross-entropy loss with the class-balanced sampling strategy (Kang et al. 2019). The following section proceeds to present our evaluation to the proposed methods.

## Experiments

### Experimental Setup

**Datasets.** We use three popular datasets to evaluate the long-tailed recognition performance.

- **ImageNet-LT** (Liu et al. 2019) contains 115,846 training images of 1,000 classes sampled from the ImageNet1K (Russakovsky et al. 2015), with class cardinality ranging from 5 to 1,280.

Settings	Many	Medium	Few	Overall
Baseline	61.6	48.6	30.3	51.2
DSCL	63.4	50.0	31.4	52.6
+ PBSD	67.2	53.7	33.7	56.3
DSCL + PBSD*	67.2	53.9	33.7	56.2
DSCL + PBSD†	68.0	53.3	32.3	56.1
DSCL + PBSD	<b>68.5</b>	<b>55.2</b>	<b>35.4</b>	<b>57.7</b>

Table 1: Effectiveness of each component in our method on ImageNet-LT. SCL is used as baseline. \* denotes using features of the global view instead of patch-based features to calculate Eq. (14). † denotes using the multi-crop trick (Caron et al. 2020) instead of PBSD.

- **iNaturalist 2018** (Van Horn et al. 2018) is a real-world long-tailed dataset with 437,513 training images of 8,142 classes, with class cardinality ranging from 2 to 1,000.
- **Places-LT** (Liu et al. 2019) contains 62,500 training images of 365 classes sampled from the Places (Zhou et al. 2017), with class cardinality ranging from 5 to 4,980.

**Evaluation Metrics.** We follow the standard evaluation metrics that evaluate our models on the testing set and report the overall top-1 accuracy across all classes. To give a detailed analysis, we follow (Liu et al. 2019) that groups the classes into splits according to their number of images: Many ( $> 100$ ), Medium (20 - 100), and Few ( $< 20$ ).

**Implementation Details.** For a fair comparison, we follow the implementations of TSC (Li et al. 2021) and KCL (Kang et al. 2020) that train the backbone at the first stage and train the linear classifier with the frozen learned backbone at the second stage. We adopt ResNet-50 (He et al. 2016) as backbone for all experiments except that using ResNet-152 pre-trained on ImageNet1K for Places-LT. The  $\alpha$  in Eq. (9) is set as 0.1 and the loss weight  $\lambda$  in Eq. (15) is 1.5.

At the first stage, the basic framework is the same as MoCoV2 (Chen et al. 2020), the momentum value for the updating of EMA model is 0.999, the temperature  $\tau$  is 0.07, the size of memory queue  $M$  is 65536, and the output dimension of projection head is 128. The data augmentation is the same as MoCoV2 (Chen et al. 2020). Locations to get the patch-based features are sampled randomly from the global view with the scale of (0.05, 0.6). Image patches cropped from the global view are resized to 64. The number of patch-based feature  $L$  per anchor image is 5. SGD optimizer is used with a learning rate decays by cosine scheduler from 0.1 to 0 with batch size 256 on 2 Nvidia RTX 3090 in 200 epochs. For Places-LT, we only fine-tune the last block of the backbone for 30 epochs (Kang et al. 2019). At the second stage, the parameters are the same as (Li et al. 2021). The linear classifier is trained for 40 epochs with CE loss and class-balanced sampling (Kang et al. 2019) with batch size 2048 using SGD optimizer. The learning rate is initialized as 10, 30, 2.5 for ImageNet-LT, iNaturalist 2018, and Places-LT, respectively, and multiplied by 0.1 at epoch 20 and 30.

## Ablation Study

**Components analysis.** We analyze the effectiveness of each proposed component on ImageNet-LT in Table 1.

Settings	ResNet50	ResNeXt50
Baseline	51.2	51.8
DSCL	52.6	53.2
+ PBSD	56.3	57.7
DSCL + PBSD	57.7	58.7

Table 2: Ablation study of each component in our method on different backbones.




Figure 5: Evaluation of  $\alpha$  in Eq. (9), the number of patch-based features  $L$  per anchor image, and the loss weight  $\lambda$  on ImageNet-LT in (a), (b), and (c), respectively. Green dotted line in (a) denotes the baseline SCL.

SCL (Khosla et al. 2020) is used as baseline. Compared with the SCL baseline, DSCL improves the top-1 accuracy by 1.4%. This result is already better than the recent contrastive learning based method TSC (Li et al. 2021). Many methods for long-tailed classification could improve the performance of tail classes but sacrifice the head classes performance. Different from those works, PBSD improves the performance of both head and tail classes. The Table 1 clearly indicates that, the combination of DSCL and PBSD achieves the best performance. The introduction of patch-based features are important to PBSD. We conduct experiment by using features of global view to calculate Eq. (14). It decreases the overall accuracy by about 1.5%. In addition, our method is also more effective than the multi-crop trick, *i.e.*, it improves the overall accuracy by 1.6% over the multi-crop trick. In summarize, each component in our method is effective in boosting the performance.

**Components analysis on different backbones.** To validate that our method generalizes well on different backbones, we further conduct experiments using the ResNeXt50 (Xie et al. 2017) as backbone on ImageNet-LT. Results are summarized in Table 2, where our proposed components are also effective on ResNext50. Both DSCL and PBSD can bring performance improvement over the baseline. The combination of them achieves the best performance.

**The impact of  $\alpha$  in Eq. (9)** is investigated in Fig. 5 (a).  $\alpha$  determines the weight of pulling the anchor with its data augmented one.  $\alpha = 0$  means only pulling the anchor with other images from the same class. This setting decreases the accuracy from 57.7% to 56.8%, showing the importance of involving two kinds of positives. In addition, this setting still outperforms the SCL baseline, *i.e.*, denoted by the green dotted line in the figure. It indicates that preventing the biased features is important.  $\alpha = 1$  degenerates the loss into the self-supervised loss. The accuracy is only 39.8% because of the lack of label information. We set  $\alpha$  as 0.1, which gets the best performance. Setting  $\alpha = 0.1$  also gets competi-

Methods	Reference	ImageNet-LT				iNaturaList 2018				Places-LT
		Many	Medium	Few	Overall	Many	Medium	Few	Overall	Overall
CE	-	64.0	33.8	5.8	41.6	72.2	63.0	57.2	61.7	30.2
Balanced	NeurIPS20	61.1	47.5	27.6	50.1	-	-	-	-	38.7
cRT	ICLR20	58.8	44.4	26.1	47.3	69.0	66.0	63.2	65.2	36.7
DisAlign	CVPR21	59.9	49.9	31.8	51.3	-	-	-	67.8	39.3
BatchFormer	CVPR22	61.4	47.8	33.6	51.1	-	-	-	-	38.2
KCL	ICLR20	61.8	49.4	30.9	51.5	-	-	-	68.6	-
PaCo‡	ICCV21	59.7	53.2	38.1	53.6	66.3	70.8	70.6	70.2	41.2
TSC	CVPR22	63.5	49.7	30.4	52.4	72.6	70.6	67.8	69.7	-
BCL	CVPR22	-	-	-	56.0	-	-	-	71.8	-
Our*	This paper	67.2	54.8	38.7	57.4	-	-	-	-	-
Our	This paper	68.5	55.2	35.4	<b>57.7</b>	74.2	72.9	70.3	<b>72.0</b>	<b>42.4</b>
RIDE	ICLR21	-	-	-	55.4	70.9	72.4	73.1	72.6	40.3
NCL†	CVPR22	-	-	-	59.5	72.7	75.6	74.5	<b>74.9</b>	-
SADE	NeurIPS22	-	-	-	-	-	-	-	72.9	40.9
Our + RIDE	This paper	70.1	57.5	37.7	<b>59.7</b>	76.2	75.7	73.6	<b>74.9</b>	-

Table 3: Comparison with recent methods on ImageNet-LT, iNaturaList2018, and Places-LT. CE denotes training the model with the cross entropy loss. \* denotes the learning rate at the second stage of our method is initialized as 2.5. ‡ denotes the model is trained without RandAug (Cubuk et al. 2020) and with 200 epochs for fair comparison. † denotes the model is trained with RandAug and 400 epochs, which is a more expensive training setup than ours.

tive performance on different datasets as shown in following experiments.

**The impact of the number of patch-based features** per anchor image is shown in Fig. 5 (b). The model benefits from involving more patch-based features into training. The top-1 accuracy improves from 55.0% to 57.7% when increasing  $L$  from 1 to 5. We set  $L$  as 5 for a reasonable trade-off between training cost and accuracy.

**The impact of the loss weight**  $\lambda$  is shown in Fig. 5 (c). Because  $\lambda$  weights the influence of PBSB, the figure shows that PBSB is important. Setting  $\lambda$  from 1 to 2 gets similar performance. We set it as 1.5 for different datasets.

### Comparison with Recent Works

We compare our method with recent works on ImageNet-LT, iNaturaList 2018, and Places-LT. The compared methods include re-balancing methods (Ren et al. 2020), decoupling methods (Kang et al. 2019; Zhang et al. 2021), transfer learning based methods (Hou, Yu, and Tao 2022), methods that extend SCL (Kang et al. 2020; Li et al. 2021; Cui et al. 2021; Zhu et al. 2022), and ensemble-based methods (Li et al. 2022; Zhang et al. 2022; Wang et al. 2020). Experimental results are summarized in Table 3.

As shown in Table 3, directly using cross entropy loss leads to a poor performance on tail classes. Most long-tailed recognition methods improve the overall performance, but sacrifice the accuracy on ‘Many’ split. Compared with the re-balancing methods, decoupling methods adjust the classifier after the training, and achieve a better performance, showing the effectiveness of the two-stage training strategy. Compared with above works, transfer learning based methods get better performance on head classes. For instance, BatchFormer gets a higher accuracy on ‘Many’ split than DisAlign which has the same overall accuracy with it.

Our method achieves the best overall accuracy of 57.7% on ImageNet-LT. It also outperforms PaCo (Cui et al. 2021) that

uses stronger data augmentation and twice training epochs. To make a fair comparison, we train PaCo with the same data augmentation and training epochs as our method, which decreases its accuracy from 57.0% to 53.6%. We also found that the learning rate of the second stage linear classifier training can change the accuracy distribution on ‘Many’, ‘Medium’ and ‘Few’ splits, while maintaining the same overall accuracy. For instance, with a learning rate of 2.5 at the second training stage, the accuracy on ‘Few’ split increases from 35.4% to 38.7%, while the overall accuracy only decreases by about 0.3%. We hence note that, the overall accuracy could be more meaningful than individual accuracy on each split, which can be adjusted by hyperparameters.

Our method can also be combined with ensemble-based method to further boost its performance. Combined with RIDE, our method achieves 59.7% overall accuracy on ImageNet-LT, outperforming all those compared ensemble-based methods. Our method also achieves superior performance on iNaturaList 2018, where it gets comparable performance with NCL that is trained with stronger data augmentation and twice training epochs. With only a single model, our method achieves the best performance on Places-LT.

### Conclusion

To tackle the challenge of long-tailed recognition, this paper analyzed two issues in SCL and addressed them with DSCL and PBSB. The DSCL decouples two types of positives in SCL, and optimizes their relations toward different objectives to alleviate the influence of the imbalanced dataset. The PBSB leverages head classes to facilitate the representation learning in tail classes by exploring patch-level similarity relationship. Experiments on different benchmarks demonstrated the promising performance of our method, where it outperforms recent works using more expensive setups. Extending our method to long-tailed detection is considered as the future work.

## Acknowledgements

This work is supported in part by Natural Science Foundation of China under Grant No. U20B2052, 61936011, in part by the Okawa Foundation Research Award.

## References

- Byrd, J.; and Lipton, Z. 2019. What is the effect of importance weighting in deep learning? In *ICML*, 872–881. PMLR.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *NeurIPS*, 33: 9912–9924.
- Chen, X.; Fan, H.; Girshick, R.; and He, K. 2020. Improved baselines with momentum contrastive learning. arXiv:2003.04297.
- Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. V. 2020. Randaugment: Practical automated data augmentation with a reduced search space. In *CVPRW*, 702–703.
- Cui, J.; Zhong, Z.; Liu, S.; Yu, B.; and Jia, J. 2021. Parametric contrastive learning. In *ICCV*, 715–724.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 9729–9738.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *ICCV*, 2961–2969.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hou, Z.; Yu, B.; and Tao, D. 2022. BatchFormer: Learning to Explore Sample Relationships for Robust Representation Learning. arXiv:2203.01522.
- Japkowicz, N.; and Stephen, S. 2002. The class imbalance problem: A systematic study. *Intelligent data analysis*, 6(5): 429–449.
- Kang, B.; Li, Y.; Xie, S.; Yuan, Z.; and Feng, J. 2020. Exploring balanced feature spaces for representation learning. In *ICLR*.
- Kang, B.; Xie, S.; Rohrbach, M.; Yan, Z.; Gordo, A.; Feng, J.; and Kalantidis, Y. 2019. Decoupling representation and classifier for long-tailed recognition. arXiv:1910.09217.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. arXiv:2004.11362.
- Li, J.; Tan, Z.; Wan, J.; Lei, Z.; and Guo, G. 2022. Nested Collaborative Learning for Long-Tailed Visual Recognition. In *CVPR*, 6949–6958.
- Li, T.; Cao, P.; Yuan, Y.; Fan, L.; Yang, Y.; Feris, R.; Indyk, P.; and Katabi, D. 2021. Targeted Supervised Contrastive Learning for Long-Tailed Recognition. arXiv:2111.13998.
- 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.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *CVPR*, 3431–3440.
- Ma, Y.; Jiao, L.; Liu, F.; Yang, S.; Liu, X.; and Li, L. 2023. Curvature-Balanced Feature Manifold Learning for Long-Tailed Classification. In *CVPR*.
- Quan, R.; Dong, X.; Wu, Y.; Zhu, L.; and Yang, Y. 2019. Auto-reid: Searching for a part-aware convnet for person re-identification. In *ICCV*, 3750–3759.
- Ren, J.; Yu, C.; Ma, X.; Zhao, H.; Yi, S.; et al. 2020. Balanced meta-softmax for long-tailed visual recognition. *NeurIPS*, 33: 4175–4186.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *IJCV*, 115(3): 211–252.
- Sun, Y.; Zheng, L.; Yang, Y.; Tian, Q.; and Wang, S. 2018. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *ECCV*, 480–496.
- Van Horn, G.; Mac Aodha, O.; Song, Y.; Cui, Y.; Sun, C.; Shepard, A.; Adam, H.; Perona, P.; and Belongie, S. 2018. The inaturalist species classification and detection dataset. In *CVPR*, 8769–8778.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *NeurIPS*, 30.
- Wang, X.; Lian, L.; Miao, Z.; Liu, Z.; and Yu, S. X. 2020. Long-tailed recognition by routing diverse distribution-aware experts. arXiv:2010.01809.
- Wei, L.; Xie, L.; He, J.; Chang, J.; Zhang, X.; Zhou, W.; Li, H.; and Tian, Q. 2020. Can Semantic Labels Assist Self-Supervised Visual Representation Learning? arXiv:2011.08621.
- Xie, S.; Girshick, R.; Dollár, P.; Tu, Z.; and He, K. 2017. Aggregated residual transformations for deep neural networks. In *CVPR*, 1492–1500.
- Yun, S.; Lee, H.; Kim, J.; and Shin, J. 2022. Patch-level representation learning for self-supervised vision transformers. In *CVPR*, 8354–8363.
- Zhang, N.; Donahue, J.; Girshick, R.; and Darrell, T. 2014. Part-based R-CNNs for fine-grained category detection. In *ECCV*, 834–849. Springer.
- Zhang, S.; Li, Z.; Yan, S.; He, X.; and Sun, J. 2021. Distribution alignment: A unified framework for long-tail visual recognition. In *CVPR*, 2361–2370.
- Zhang, S.; Zhou, Q.; Wang, Z.; Wang, F.; and Yan, J. 2023. Patch-level Contrastive Learning via Positional Query for Visual Pre-training. In *ICML*, 41990–41999. PMLR.
- Zhang, Y.; Hooi, B.; Hong, L.; and Feng, J. 2022. Self-supervised aggregation of diverse experts for test-agnostic long-tailed recognition. *NeurIPS*, 35: 34077–34090.
- Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; and Torralba, A. 2017. Places: A 10 million image database for scene recognition. *TPAMI*, 40(6): 1452–1464.
- Zhu, J.; Wang, Z.; Chen, J.; Chen, Y.-P. P.; and Jiang, Y.-G. 2022. Balanced contrastive learning for long-tailed visual recognition. In *CVPR*, 6908–6917.