
Contrastive Learning with Boosted Memorization

Zhihan Zhou¹ Jiangchao Yao¹ Yanfeng Wang^{1,2} Bo Han³ Ya Zhang^{1,2}

Abstract

Self-supervised learning has achieved a great success in the representation learning of visual and textual data. However, the current methods are mainly validated on the well-curated datasets, which do not exhibit the real-world long-tailed distribution. Recent attempts to consider self-supervised long-tailed learning are made by re-balancing in the *loss perspective* or the *model perspective*, resembling the paradigms in the supervised long-tailed learning. Nevertheless, without the aid of labels, these explorations have not shown the expected significant promise due to the limitation in tail sample discovery or the heuristic structure design. Different from previous works, we explore this direction from an alternative perspective, *i.e.*, the *data perspective*, and propose a novel Boosted Contrastive Learning (BCL) method. Specifically, BCL leverages the memorization effect of deep neural networks to automatically drive the information discrepancy of the sample views in contrastive learning, which is more efficient to enhance the long-tailed learning in the label-unaware context. Extensive experiments on a range of benchmark datasets demonstrate the effectiveness of BCL over several state-of-the-art methods. Our code is available at <https://github.com/MediaBrain-SJTU/BCL>.

1. Introduction

Self-supervised learning (Doersch et al., 2015; Wang & Gupta, 2015) that learns the robust representation for downstream tasks have achieved a significant success in the area of computer vision (Chen et al., 2020a; He et al., 2020) and natural language processing (Lan et al., 2019; Brown et al.,

2020). Nevertheless, previous studies are mainly conducted on the well-curated datasets like ImageNet (Deng et al., 2009), which is usually balanced among categories. In comparison, the real-world natural sources usually follow a long-tailed, even heavy-tailed distributions (Reed, 2001) that is challenging to learn for the current machine learning methods. Specially, recent attempts (Jiang et al., 2021b) have shown that self-supervised learning under long-tailed distribution still requires more explorations to achieve the satisfying performance (Van Horn et al., 2018).

Existing works for self-supervised long-tailed learning are mainly from the *loss perspective* or the *model perspective*. The former relies on the loss reweighting, *e.g.*, the focal loss in hard example mining (Lin et al., 2017) or SAM by means of the sharpness of the loss surface (Liu et al., 2021), to draw more attention on tail samples during training. However, the effectiveness of these methods is sensitive to and limited by the accuracy of the tail sample discovery. The latter mainly resorts to the specific model design like the divide-and-contrast ensemble (Tian et al., 2021) or self-damaged-contrast via pruning (Jiang et al., 2021b) to make the model better capture the semantics of the tail samples. These designs require the empirical heuristic and are usually black-box to understand the potential working dynamics for the further improvement (Zhang et al., 2021).

In this paper, we propose to study the self-supervised long-tailed learning in the *data perspective*. Our framework is motivated by the memorization effect (Zhang et al., 2017; Arpit et al., 2017; Feldman, 2020) of deep neural networks on data, where the easy patterns are usually memorized prior to the hard patterns. As shown in the left panel of Figure 1, the memorization effect still holds under long-tailed datasets, where the loss and accuracy of the tail samples consistently fall behind those of head samples. This inspires us to approximately distinguish the head and tail samples by analyzing the memorization effect. Another important motivation is except the loss reweighting or model re-design, the data augmentation is very effective in self-supervised long-tailed learning to achieve the improvement by introducing the information discrepancy of two views (Tian et al., 2020). As illustrated in the right panel of Figure 1, we can see that the heavier augmentation consistently boosts the performance of the treatment tail samples. Besides, the data augmentation does not directly modify the loss or the model

¹Cooperative Medianet Innovation Center, Shanghai Jiao Tong University ²Shanghai AI Laboratory ³Department of Computer Science, Hong Kong Baptist University. Correspondence to: Jiangchao Yao <Sunarker@sjtu.edu.cn>, Yanfeng Wang <wangyanfeng@sjtu.edu.cn>.




Figure 1. (Left) Test accuracy and loss of head and tail classes during the training stage on CIFAT-100-LT. (Right) Test accuracy of tail classes when deploying different strength of RandAugment on tail classes on CIFAT-100-LT. k is a hyper-parameter controlling the amount of augmentations used in RandAugment.

structure and thus is more robust to the noisy tail discovery.

On basis of the aforementioned observations in Figure 1, we introduce a novel Boosted Contrastive Learning method from the data perspective. Concretely, we propose a momentum loss to capture the clues from the memorization effect of DNNs to anchor the mostly possible tail samples. Then, the momentum loss is used to drive an instance-wise augmentation by constructing different information discrepancy for head and tail samples. In an end-to-end manner, BCL maintains the learning of head samples, meanwhile enhances the learning of hard-to-memorize tail samples.

Main Contributions

- Different from previous works in the loss and model perspectives, we are the first to explore self-supervised long-tailed learning from the data perspective, which leverages the DNN memorization effect on data and the augmentation efficiency in self-supervised learning.
- We propose a Boosted Contrastive Learning method, which builds a momentum loss to capture clues from the memorization effect and drive the instance-wise augmentation to dynamically maintain the learning of head samples and enhance the learning of tail samples.
- The proposed BCL is orthogonal to the current self-supervised methods on long-tailed data. Extensive experiments on a range of benchmark datasets demonstrate the superior performance of BCL.

2. Related Works

Supervised Long-tailed Learning. Recent works (Yang & Xu, 2020; Kang et al., 2020) start to boost the long-tailed recognition via the lens of representation learning (Zheng et al., 2019). Kang et al. (2019) proposed to disentangle representation and classification learning in a two-stage

training scheme and empirically observed that the instance-balanced sampling performs best for the first stage, which attracts more attention to representation learning in long-tailed recognition. Yang & Xu (2020) theoretically investigated the necessity of the label information for long-tailed data and showed the promise of self-supervised pre-training stage on long-tailed recognition. Motivated by these findings, Kang et al. (2020) first leveraged supervised contrastive learning paradigm for long-tailed recognition and claimed that the learned feature space is more balanced compared with the supervised learning. Cui et al. (2021) theoretically showed that supervised contrastive learning still suffers from the bias from the head classes under imbalanced data. They proposed a parametric class-wise learnable center to rebalance the contrastive loss across different class cardinality. The concurrent work (Li et al., 2021) proposed a uniform class center assignment strategy to force a balanced feature space.

Self-supervised Long-tailed Learning. In self-supervised learning area, several works (Chen et al., 2020a; He et al., 2020; Chen & He, 2021) mainly target to the curated and balanced dataset and naturally build the uniformity assumption. For example, Wang & Isola (2020) concluded that one key property of contrastive learning is to learn a uniform feature space by information maximization. Caron et al. (2020) assumed that all the samples are distributed uniformly at the prototype level and operated the fast Sinkhorn-Knopp algorithm (Cuturi, 2013) for the uniform online clustering. However, it may cause performance degeneration to model the real-world distribution in a uniform way as the practical data generally follows a skewed distribution (Reed, 2001).

There exists a few attempts (Liu et al., 2021; Jiang et al., 2021b; Zheng et al., 2021) towards self-supervised long-tailed learning, which can be divide into two categories: *loss-based* or *model-based* methods. A classical solution in the first category, *i.e.*, the focal loss (Lin et al., 2017), relies on the individual sample difficulty to rebalance the learning. Recently, Liu et al. (2021) proposed a sharpness regularization on loss surface to enhance model generalization. From the model perspective, Jiang et al. (2021b) assumed tail samples to be easily forgotten and designed a asymmetric network with a pruned branch to identify the tail classes. An alternative (Tian et al., 2021) targeted at the uncurated data faces the similar challenges in long-tailed recognition. They proposed a multi-expert framework to extract the more fine-grained features in the separated clusters. Different from these works, we explores the benefit of the data perspective for the self-supervised long-tailed representation learning.

Memorization Effect. The definition on the memorization effect of DNNs can trace back to the generalization study on noisy data (Zhang et al., 2017; Arpit et al., 2017). These findings shed lights on a stream of loss-aware studies towards noisy representation learning (Jiang et al., 2018;

Ren et al., 2018; Han et al., 2018). Specifically, they regard the small-loss samples as clean samples and then employ the sample selection or loss reweighting. For example, co-teaching (Han et al., 2018; Yu et al., 2019) selects the small-loss samples and discards high-loss samples in the training stage. Meanwhile, Ren et al. (2018) proposed a meta-learning framework to assign different weights to the training samples according to the loss value.

Recently, Feldman (2020) extended the memorization effect of deep neural networks towards the long-tailed samples. They concluded that the memorization of DNNs is necessary for the rare and atypical instances and proposed a memorization measurement. Specifically, the memorization score are defined as the drop in the prediction accuracy for each sample in the training dataset when removing the respective sample. However, the computational cost of estimating this memorization score is expensive. The subsequent work (Jiang et al., 2021c) explored some more efficient proxies to alternate the hold-out estimator. In particular, a learning speed based proxy have shown the positive correlation with the memorization score, which is in consistency with the observation of the memorization effect in (Feldman, 2020). Different from these explorations that require labels available, our methods conversely focus on the annotation-free long-tailed sample discovery.

3. Methods

3.1. Preliminary

In this section, we give the basic notations of contrastive learning that our method builds on. Generally, the classical contrastive learning (Chen et al., 2020a), termed as SimCLR, is defined as follows,

$$\mathcal{L}_{CL} = \frac{1}{N} \sum_{i=1}^N -\log \frac{\exp\left(\frac{f(x_i)^\top f(x_i^+)}{\tau}\right)}{\sum_{x'_i \in X^- \cup \{x_i^+\}} \exp\left(\frac{f(x_i)^\top f(x'_i)}{\tau}\right)} \quad (1)$$

where (x_i, x_i^+) is the positive sample sample pair and X^- is the negative sample set of x , τ is the temperature and $f(\cdot)$ is the encoder function. In practical, x_i and x_i^+ are two views of one example, while $x'_i \in X^-$ is the view of other samples. Contrastive learning is to learn a representation that is invariant to itself in the small perturbation but keeps the variance among different samples.

3.2. Motivation

Deep supervised long-tailed learning has made great progresses in the last ten years (Zhang et al., 2021) to handle the real-world data distributions. Nevertheless, previous works mainly focus on the supervised learning case, namely the labels of natural sources must be available, while only few works (Jiang et al., 2021b; Liu et al., 2021) pay at-

tention to the study of such a skew distribution under the self-supervised learning scenario. Compared to the supervised learning, long-tailed learning without labels is more practical and important, since in a range of cases, *e.g.*, the large-scale datasets, it is expensive to collect the annotation of each sample. Concomitantly, this task is more challenging, since most of previous works build on top of the explicit label partition of head and tail samples.

Without labels, previous self-supervised learning study in this direction leverages the implicit balancing from the loss perspective (Lin et al., 2017; Liu et al., 2021) or the model perspective (Jiang et al., 2021b) to enhance the learning law on tail samples. Different from these works, BCL explicitly trace the memorization effect via a learning speed scope based on theoretical and empirical findings (Feldman, 2020; Jiang et al., 2021c) in the context of supervised image classification. The definition (Feldman, 2020) that describes how models memorize the patterns of the individual sample during the training is given as follows:

$$\text{mem}(\mathcal{A}, S, i) := \Pr_{h \leftarrow \mathcal{A}(S)} [h(x_i) = y_i] - \Pr_{h \leftarrow \mathcal{A}(S \setminus i)} [h(x_i) = y_i] \quad (2)$$

where \mathcal{A} denotes the training algorithm and $S \setminus i$ denotes removing the sample point (x_i, y_i) from the data collection S . Unfortunately, the hold-out retraining metric is computationally expensive and only limited to the supervised learning. Inspired by the learning speed proxy explored in the subsequent work (Jiang et al., 2021c), we first extend the memorization estimation to the self-supervised learning task. Specifically, we propose the momentum loss to characterize the learning speed of individual sample, which is used to reflect the memorization effect. Merits of the proposed historical statistic are two-fold: computationally efficient and robust to the randomness issue without the explicit label calibration in contrastive loss (Chen et al., 2020a).

Besides, we boost the performance of contrastive learning on tail samples from the data perspective, *i.e.*, construct the heavier information discrepancy between two views of the sample instead of the previous loss reweighting (Lin et al., 2017; Liu et al., 2021) or the model pruning (Jiang et al., 2021b). According to the ‘‘InfoMin Principle’’ (Tian et al., 2020), a good set of views are those that share the minimal information necessary to perform well at the downstream task. In this spirit, BCL dynamically constructs the information discrepancy between views to boost representation learning based on the memorization effect. Specifically, BCL constructs the stronger information discrepancy between views to emphasize the importance of tail samples, while maintains the relative high correlation between views for head samples to avoid fitting to task-irrelevant noise. This allows our model to capture more task-relevant information from samples in long-tailed distribution.




Figure 2. The illustration of Boosted Contrastive Learning. We trace the historical losses of each sample to find the clues about the memorization effect of DNNs, which then drives the augmentation strength to enhance the learning law on the tail samples. The “head” and “tail” indicators about the cat image and the tiger image are the exemplars and actually unknown during the training.

3.3. Boosted Contrastive Learning

In this section, we will present the formulation of the proposed Boosted Contrastive Learning, which leverages a momentum loss proxy to control the augmentation to affect the memorization effect of DNNs. Specifically, as tail samples tend to be learned slowly, they will be assigned with higher intensities of augmentation. Then, the model is driven to extract more information from the augmented views of tail samples for the better generalization.

Concretely, given a training sample x_i on the long-tailed dataset, we denote its contrastive loss as \mathcal{L}_i and $\{\mathcal{L}_{i,0}, \dots, \mathcal{L}_{i,t}, \dots, \mathcal{L}_{i,T}\}$ traces a sequence of the loss values \mathcal{L}_i in T epochs. We then define the following moving-average momentum loss,

$$\mathcal{L}_{i,0}^m = \mathcal{L}_{i,0}, \quad \mathcal{L}_{i,t}^m = \beta \mathcal{L}_{i,t-1}^m + (1 - \beta) \mathcal{L}_{i,t}$$

where β is a hyper-parameter to control the degree smoothed by the historical losses. After the training in the t -th epoch through the above moving-average, we could acquire a set of the momentum losses for each sample as $\{\mathcal{L}_{0,t}^m, \dots, \mathcal{L}_{i,t}^m, \dots, \mathcal{L}_{N,t}^m\}$, where N is the number of training samples in the dataset. Finally, we define the following normalization on the momentum losses,

$$\mathbf{M}_{i,t} = \frac{1}{2} \left(\frac{\mathcal{L}_{i,t}^m - \bar{\mathcal{L}}_t^m}{\max \{|\mathcal{L}_{i,t}^m - \bar{\mathcal{L}}_t^m|\}_{i=0, \dots, N}} + 1 \right) \quad (3)$$

where $\bar{\mathcal{L}}_t^m$ is the average momentum loss at the t -th training epoch. By Eq. (3), \mathbf{M}_i is normalized to $[0, 1]$ with the average value of 0.5, which reflects the intensity of the memorization effect. To boost the contrastive learning, we use \mathbf{M}_i as an indicator controlling the occurrence and strength of the augmentation. Specifically, we randomly selects k types of augmentations from RandAugment (Cubuk et al., 2020) and apply each augmentation with probability \mathbf{M}_i and strength

$[0, \mathbf{M}_i]$, respectively. For clarity, we assume augmentations defined in RandAugment as $\mathcal{A} = (A_1, \dots, A_j, \dots, A_K)$, where K denotes the amount of augmentations. In each step, only k augmentations are applied (Cubuk et al., 2020). We formulate the memorization-boosted augmentation $\Psi(x_i)$:

$$\Psi(x_i; \mathcal{A}, \mathbf{M}_i) = a_1(x_i) \circ \dots \circ a_k(x_i),$$

$$a_j(x_i) = \begin{cases} A_j(x_i; \mathbf{M}_i \zeta) & u \sim \mathcal{U}(0, 1) \text{ \& } u < \mathbf{M}_i \\ x_i & \text{otherwise} \end{cases} \quad (4)$$

where ζ is sampled from the uniform distribution $\mathcal{U}(0, 1)$ and $a_j(x_i)$ means we decide to keep x_i unchanged or augment x_i by $A_j(x_i; \mathbf{M}_i \zeta)$ based on whether u is greater than \mathbf{M}_i . $A_j(x_i; \mathbf{M}_i \zeta)$ represents applying the j -th augmentation to x_i with the strength $\mathbf{M}_i \zeta$, and \circ is the function composition¹ operator, namely, sequentially applying the selected k augmentations in \mathcal{A} . For simplicity, we use $\Psi(x_i)$ to represent $\Psi(x_i; \mathcal{A}, \mathbf{M}_i)$ in this paper. Our boosted contrastive learning loss are formulated as follows.

$$\mathcal{L}_{\text{BCL}} = \frac{1}{N} \sum_{i=1}^N -\log \frac{\exp \left(\frac{f(\Psi(x_i))^T f(\Psi(x_i^+))}{\tau} \right)}{\sum_{x'_i \in X'} \exp \left(\frac{f(\Psi(x_i))^T f(\Psi(x'_i))}{\tau} \right)} \quad (5)$$

where X' represents $X^- \cup \{x_i^+\}$ as Eq. (1). Intuitively, at a high level, BCL can be understood as a curriculum learning method that adaptively assigns the appropriate augmentation strength for the individual sample according to the feedback from the memorization clues. Let θ denotes the model parameters and we have the following procedure

$$\theta = \arg \min_{\theta} \mathcal{L}_{\text{BCL}}(X, \Psi, \theta),$$

$$\Psi = \Psi(x; \mathcal{A}, \mathbf{M}), \quad \mathbf{M} = \text{Normalize}(\mathcal{L}_{\text{BCL}}^m).$$

¹https://en.wikipedia.org/wiki/Function_composition

Algorithm 1 Boosted Contrastive Learning (BCL)

Input: dataset \mathcal{X} , the epoch number T , the weighting factor β , the number k used in RandAugment, the whole augmentation set \mathcal{A} (K augmentation types)

Output: pretrained model parameter θ_T

Initialize: model parameter θ_0

```

1: if  $t = 0$  then
2:   Train model  $\theta_0$  with Eq. (1) and initialize  $\mathcal{L}_0^m, \mathbf{M}_0$ .
3: end if
4: for  $t = 1, \dots, T - 1$  do
5:   for  $x$  in  $\mathcal{X}$  do
6:     Select  $k$  augmentations from the augmentation set
        $\mathcal{A}$  and construct augmented views  $\Psi_t(x)$  accord-
       ing to  $\mathbf{M}_{t-1}$  with Eq. (4).
7:   end for
8:   Train model  $\theta_t$  with Eq. (5) or Eq. (6) and obtain  $\mathcal{L}_t$ ;
9:   Obtain  $\mathcal{L}_t^m = \beta \mathcal{L}_{t-1}^m + (1 - \beta) \mathcal{L}_t$  with stored  $\mathcal{L}_{t-1}^m$ ;
10:  Update  $\mathbf{M}_t \leftarrow \frac{1}{2} \left( \frac{\mathcal{L}_{i,t}^m - \tilde{\mathcal{L}}_t^m}{\max\{|\mathcal{L}_{i,t}^m - \tilde{\mathcal{L}}_t^m|\}_{i=0,\dots,N}} + 1 \right)$ ;
11: end for

```

In this way, BCL continually depends on Ψ to highlight the training samples to which DNNs show the poor memorization effect until its momentum loss $\mathcal{L}_{\text{BCL}}^m$ degrades. By iteratively optimizing the model and building the memorization-boosted information discrepancy, we adaptively motivate model to learn “residual” information contained in tail samples. Note that, the form of \mathcal{L}_{BCL} can be flexibly replaced by the extensions from more self-supervised methods. In this paper, we mainly investigate two BCL types, *i.e.*, BCL-I (Identity) and BCL-D (Damaging). Specifically, BCL-I means the plain BCL in Eq. (5), while BCL-D is built on SDCLR and is formulated by the following equation,

$$\mathcal{L}_{\text{BCL-D}} = \frac{1}{N} \sum_{i=1}^N -\log \frac{\exp\left(\frac{f(\Psi(x_i))^\top g(\Psi(x_i^+))}{\tau}\right)}{\sum_{x'_i \in \mathcal{X}'} \exp\left(\frac{f(\Psi(x_i))^\top g(\Psi(x'_i))}{\tau}\right)} \quad (6)$$

where g is the pruned version of f as detailed in SDCLR (Jiang et al., 2021b). We illustrate BCL in Figure 2 and summarize the complete procedure in Algorithm 1.

3.4. More Discussions on BCL

Complexity. The additional storage in BCL compared with that in the standard contrastive learning methods is the momentum loss. In Eq. (3), we only need to save a scalar $\mathcal{L}_{i,t-1}^m$ of the previous epoch for each sample. Therefore, its storage cost is as cheap as that of one label in the float type.

Compatibility. BCL does not require the specific model structures and thus it is compatible with many self-supervised learning methods in the recent years (Chen et al.,

2020a; He et al., 2020; Grill et al., 2020; Ermolov et al., 2021; Chen & He, 2021). Besides, it can be potentially adapted to enhance the representation learning under the supervised long-tailed learning setting in the form of pre-training or regularization for the representation learning of head and tail samples.

Relation to loss re-weighting. Loss re-weighting is an explicit way to enhance the learning of the specific samples by enlarging the importance of their losses. Previous attempts like Focal loss (Lin et al., 2017) and SAM (Liu et al., 2021) belong to this case. In comparison, BCL does not directly modify the loss, but captures the memorization clues to drive the construction of information discrepancy for the implicit re-weighting. In the following section, we will show that this actually is a more efficient way to bootstrap the long-tailed representation learning without label annotations.

4. Experiments

4.1. Datasets and Baselines

We conduct extensive experiments on three benchmark long-tailed datasets: CIFAR-100-LT (Cao et al., 2019), ImageNet-LT (Liu et al., 2019) and Places-LT (Liu et al., 2019).

CIFAR-100-LT: The original CIFAR-100 is a small-scale dataset composed of 32×32 images from 100 classes. For the long-tailed version, we use the same sampled subsets of CIFAR-100 as in (Jiang et al., 2021b). The imbalance factor is defined by the number of the most frequent classes divided by the least frequent classes. Following (Jiang et al., 2021b), we set the imbalance factor as 100 and conduct experiments on five long-tailed splits to avoid randomness.

ImageNet-LT: ImageNet-LT (Liu et al., 2019) is a long-tailed version of ImageNet, which is down-sampled according to the Pareto distribution with the power value $\alpha = 6$. It contains 115.8K images of 1000 categories, ranging from 1,280 to 5 in terms of the class cardinality.

Places-LT: Places (Zhou et al., 2017) is a large-scale scene-centric dataset and Places-LT is a long-tailed subset of Places following the Pareto distribution (Liu et al., 2019). It contains 62,500 images in total from 365 categories, ranging from 4,980 to 5 under the class cardinality.

Baselines: To demonstrate the effectiveness of our method on benchmark datasets, we compare to many self-supervised methods related under long-tailed representation learning, including: (1) *contrastive learning baseline*: SimCLR (Chen et al., 2020a), (2) *hard example mining*: Focal loss (Lin et al., 2017), (3) *model ensemble*: DnC (Tian et al., 2021), (4) *model damaging*: SDCLR (Jiang et al., 2021b). As mentioned before, BCL can be combined with any self-supervised learning architectures. Here, we term its

Table 1. Fine-grained analysis for various methods pre-trained on CIFAR-100-LT, ImageNet-LT and Places-LT. Many/Medium/Few corresponds to three partitions on the long-tailed data. Std is the standard deviation of the accuracies among Many/Medium/Few groups.

| Methods | CIFAR-100-LT | | | | ImageNet-LT | | | | Places-LT | | | |
|---------|--------------|--------------|--------------|-------------|--------------|--------------|--------------|-------------|--------------|--------------|--------------|-------------|
| | Many | Medium | Few | Std | Many | Medium | Few | Std | Many | Medium | Few | Std |
| SimCLR | 48.70 | 46.81 | 44.02 | 2.36 | 41.16 | 32.91 | 31.76 | 5.13 | 31.12 | 33.85 | 35.62 | 2.27 |
| Focal | 48.46 | 46.73 | 44.12 | 2.18 | 40.55 | 32.91 | 31.29 | 4.95 | 30.18 | 31.56 | 33.32 | 1.57 |
| DnC | <u>54.00</u> | 46.68 | 45.65 | 4.55 | 29.54 | 19.62 | 18.38 | 6.12 | 28.20 | 28.07 | 28.46 | 0.20 |
| SDCLR | <u>51.22</u> | <u>49.22</u> | <u>45.85</u> | 2.71 | <u>41.24</u> | <u>33.62</u> | <u>32.15</u> | 4.88 | <u>32.08</u> | <u>35.08</u> | <u>35.94</u> | <u>2.03</u> |
| BCL-I | 50.45 | <u>48.23</u> | <u>45.97</u> | <u>2.24</u> | 42.53 | 35.66 | <u>33.93</u> | 4.54 | <u>32.27</u> | <u>34.96</u> | 38.03 | 2.88 |
| BCL-D | 53.98 | 51.97 | 49.52 | <u>2.23</u> | <u>41.92</u> | <u>35.29</u> | 34.07 | 4.22 | 32.34 | 35.44 | <u>37.75</u> | 2.71 |

combination with SimCLR as BCL-I and its combination with SDCLR as BCL-D, respectively.

4.2. Implementation Details

For all experiments, we use the SGD optimizer and the cosine annealing schedule. Similar to the backbone architecture and projection head proposed in (Chen et al., 2020a), we use ResNet-18 (He et al., 2016) as the backbone for experiments on CIFAR-100-LT and ResNet-50 on ImageNet-LT and Places-LT. The smoothing factor β in the momentum loss Eq. (3) is set as 0.97. Besides, we set $k = 1$ for BCL-I and $k = 2$ for BCL-D in the RandAugment. The whole augmentation set \mathcal{A} is aligned with RandAugment where $K = 16$. For the other pre-training settings, we follow (Jiang et al., 2021b) and during evaluation, we leverage the way in (Ermolov et al., 2021). Specifically, we train the classifier for 500 epochs and employ the learning rate decaying from 10^{-2} to 10^{-6} . We use the Adam optimizer with the weight decay 5×10^{-6} .

We follow (Ermolov et al., 2021) to conduct *linear probing evaluation*, where a linear classifier is trained on top of the frozen pretrained backbone and the test accuracy is calculated to measure the representation quality. To eliminate the effect of long-tailed distribution in the fine-tuning stage, the classifier is trained on a balanced dataset. Specifically, we report the *few-shot* performance of the classifier on basis of the pretrained representation. In the default case, we conduct 100-shot evaluation on CIFAR-100-LT, ImageNet-LT and Places-LT for performance evaluation. Meanwhile, we also implement the full-shot, 100-shot and 50-shot evaluation for ablation study on CIFAR-100-LT.

To visualize the fine-grained performance under the long-tailed setting, we divide each dataset to three partitions (*Many-Medium-Few*). Following (Jiang et al., 2021b) on CIFAR-100-LT, the resulted partitions are Many (34 classes, 500 to 106 samples in the cardinal classes), Medium (33 classes, 105 to 20 samples in the cardinal classes) and Few (33 classes, 19 to 5 samples in the cardinal classes), respectively. As for the large-scale datasets ImageNet-LT and Places-LT, we follow (Liu et al., 2019) to divide each dataset

Table 2. The overall performance of various methods pre-trained on CIFAR-100-LT, ImageNet-LT and Places-LT with 100-shot.

| Methods | CIFAR-100-LT | ImageNet-LT | Places-LT |
|---------|--------------|--------------|--------------|
| SimCLR | 46.53 | 35.93 | 33.22 |
| Focal | 46.46 | 35.63 | 31.41 |
| DnC | <u>48.53</u> | 23.27 | 28.19 |
| SDCLR | <u>48.79</u> | <u>36.35</u> | <u>34.17</u> |
| BCL-I | 48.24 | 38.07 | <u>34.59</u> |
| BCL-D | 51.84 | <u>37.68</u> | 34.78 |

into Many (over 100 samples), Medium (100 to 20 samples) and Few (under 20 samples). The average accuracy and the standard deviation are computed among three groups.

4.3. Performance Evaluation

Overall performance. In Table 2, we summarize the performance of different methods on three long-tailed datasets. According to the results, BCL-I and BCL-D significantly improve the few-shot performance by 1.71% and 3.05% on SimCLR and SDCLR on CIFAR-100-LT. On large-scale datasets ImageNet-LT and Places-LT, compared with SimCLR, SDCLR only improves the few-shot accuracy by 0.42% and 0.95%. However, our methods maintain a consistent gain over other self-supervised methods and specifically, BCL-I achieves comparable performance with BCL-D and outperforms SDCLR by 1.72% on ImageNet-LT.

Fine-grained analysis. In Table 1, we visualize the merit of BCL in the fine-grained perspective. According to the results on CIFAR-100-LT, ImageNet-LT and Places-LT, we can see that BCL achieves the new state-of-the-art performance on each partition across different benchmark datasets. For example, compared with SDCLR, BCL-D improves Many, Medium and Few accuracy by 2.77%, 2.75% and 3.87% on CIFAR-100-LT, respectively. We also apply standard deviation(Std) of average accuracy on each partition to measure the representation balancedness. As shown in Table 1, we see that our methods reduce Std by a considerable margin of 0.4–0.7 on CIFAR-100-LT and ImageNet-LT.




Figure 3. Long-tailed sample discovery with our momentum loss (ML) and the conventional contrastive loss (CL) under different training epochs on CIFAR-100-LT. ϕ means the proportion of head or tail classes in the top 10% large-loss samples of the dataset.

Note that, the results on Places-LT differ from the former datasets as the performance of three groups shows a reverse trend on the long-tailed distribution. Nevertheless, an interesting observation is that BCL-I still significantly improves Few accuracy by 2.09% while maintain at Many(0.19%) compared with SDCLR. The results confirm that BCL can boost the performance on tail classes and potentially handle the more complicated real-world data distribution.

Long-tailed sample discovery. We use ground-truth labels to validate the tail detection of the momentum loss mechanism in Eq. (3). First, we pre-train SimCLR and store the loss value of the training sample in each epoch. We then calculate the momentum loss and choose the training samples that have top-10% highest loss. To mitigate the effect of the group size, we apply the correlation metric in (Jiang et al., 2021a) and divide the train dataset into head (Major) and tail (Medium, Few). Specifically, the metric is defined as:

$$\phi = \frac{\mathcal{G} \cap \mathcal{X}_l}{\mathcal{G} \cap \mathcal{X}}, \quad \mathcal{X}_l = \arg \max_{\mathcal{X}': |\mathcal{X}'| \geq r|\mathcal{X}|} \mathcal{L}(\mathcal{X}')$$

where \mathcal{G} denotes the target group, \mathcal{X}_l represents the subset of large-loss samples and r represents the threshold ratio. We set $r = 0.1$ and compare the proposed \mathcal{L}_{CL}^m with \mathcal{L}_{CL} .

As shown in Figure 3, we can find that more tail samples are extracted by the proposed momentum loss compared those by the standard contrastive loss. Meanwhile, we find that momentum loss serves as a reliable tail detector only except the early stage of training process in Figure 3. As the momentum loss is built on the historical information, a long-term observation will yield a more stable estimation.

4.4. On Transferability for Downstream Tasks

Downstream supervised long-tailed classification. Self-supervised pre-training is proved to be useful for learning

Table 3. The classification accuracy of supervised learning with self-supervised pre-training on CIFAR-100-LT and ImageNet-LT.

| Dataset | CE | CE with the following model initialization | | | | | |
|--------------|------|---|-------|------|-------------|-------------|-------------|
| | | CL | Focal | DnC | SDCLR | BCL-I | BCL-D |
| CIFAR-100-LT | 41.7 | 44.4 | 44.4 | 44.4 | <u>44.6</u> | 45.1 | 45.4 |
| ImageNet-LT | 41.6 | 45.5 | 45.4 | 42.2 | <u>45.9</u> | 46.9 | 46.4 |
| Dataset | cRT | cRT with the following model initialization | | | | | |
| | | CL | Focal | DnC | SDCLR | BCL-I | BCL-D |
| CIFAR-100-LT | 44.1 | 48.9 | 48.7 | 48.6 | <u>49.8</u> | 49.9 | 50.0 |
| ImageNet-LT | 46.7 | <u>47.5</u> | 47.3 | 43.5 | 47.3 | 48.4 | 48.1 |
| Dataset | LA | LA with the following model initialization | | | | | |
| | | CL | Focal | DnC | SDCLR | BCL-I | BCL-D |
| CIFAR-100-LT | 45.7 | 50.1 | 49.5 | 49.7 | <u>50.4</u> | 50.8 | 50.5 |
| ImageNet-LT | 47.4 | <u>48.6</u> | 48.4 | 45.6 | 48.2 | 49.7 | 49.1 |

more generalizable representations by label-agnostic model initialization (Yang & Xu, 2020). All the self-supervised long-tailed baselines can be regarded as the pre-training methods that are compatible with supervised algorithms. In order to validate the effectiveness of BCL on downstream supervised long-tailed tasks, we use the pre-trained self-supervised models to initialize the supervised model backbone and then finetune all parameters. Specifically, we evaluate and compare 3 representative long-tailed methods: Cross Entropy, cRT (Kang et al., 2019) and Logit Adjustment (Menon et al., 2021) with 6 self-supervised initialization methods on CIFAR-100-LT and ImageNet-LT. The results of the finetuning experiment are summarized in Table 3, showing that initialization with self-supervised models always helps improve over the standard baseline, and BCL outperforms all other self-supervised pre-training methods. This indicates the potential merits of BCL to further boost the supervised long-tailed representation learning.

Downstream fine-grained classification. In order to validate the representation transferability of our memorization-boosted augmentation, we conduct experiments on various downstream fine-grained datasets: Caltech-UCSD Birds (CUB200) (Wah et al., 2011), Stanford Cars (Krause et al., 2013), Aircrafts (Maji et al., 2013), Stanford Dogs (Khosla et al., 2011), NABirds (Van Horn et al., 2015). The training and testing images of these datasets roughly range from 10k to 50k. Meanwhile, these datasets include five distinct categories, from birds to cars, where the intrinsic property of the data distribution varies. We first pre-train the model on ImageNet-LT and then conduct the linear probing evaluation on these target datasets individually.

In Table 4, we present the transfer results on various downstream tasks. According to the table, we can see that our methods consistently surpass other methods with a considerable margin in all cases. Specifically, our methods significantly improve the best Top-1 accuracy by 3.80%, 3.90% and 1.92% on Stanford Cars, Aircrafts and Dogs, and by 0.38% and 0.90% on the other two bird datasets,

Table 4. The linear probing performance of all methods on CUB, Cars, Aircrafts, Dogs and NABirds. We pretrain the backbone ResNet-50 on ImageNet-LT under different methods, and then transfer to these datasets for the linear probing evaluation. The top-1 and top-5 accuracies are reported by computing the highest and top-5 highest predictions to match the ground-truth labels.

| Methods | CUB | | Cars | | Aircrafts | | Dogs | | NABirds | | All | |
|---------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | Top-1 | Top-5 | Top-1 | Top-5 | Top-1 | Top-5 | Top-1 | Top-5 | Top-1 | Top-5 | Top-1 | Top-5 |
| SimCLR | <u>29.62</u> | <u>57.35</u> | 21.45 | 44.93 | 30.48 | 57.01 | 46.67 | 79.22 | <u>16.52</u> | <u>37.61</u> | 28.95 | 55.22 |
| Focal | <u>29.08</u> | 56.89 | 21.40 | 44.35 | 30.99 | 57.64 | 46.59 | 78.14 | 16.31 | 36.97 | 28.87 | 54.80 |
| DnC | <u>16.97</u> | 40.90 | 8.15 | 23.79 | 13.71 | 33.18 | 29.83 | 61.92 | 8.44 | 22.75 | 15.42 | 36.51 |
| SDCLR | 28.98 | 57.27 | <u>22.10</u> | <u>46.13</u> | <u>31.05</u> | <u>58.18</u> | <u>46.69</u> | <u>78.82</u> | 16.17 | 37.10 | <u>29.00</u> | <u>55.50</u> |
| BCL-I | 30.00 | 58.08 | <u>23.67</u> | <u>49.16</u> | <u>32.37</u> | <u>60.31</u> | 48.61 | 79.99 | 17.42 | 38.96 | <u>30.41</u> | <u>57.30</u> |
| BCL-D | 28.79 | <u>57.37</u> | 25.90 | 51.34 | 34.95 | 62.77 | <u>47.49</u> | <u>78.86</u> | <u>16.41</u> | <u>37.24</u> | 30.71 | 57.51 |

CUB and NABirds. Overall, BCL-D on average improves Top-1 and Top-5 accuracy by 1.71% and 2.01% on five target datasets. This confirms our intuition that there is discarded transferable information for tail samples, which is effectively extracted by BCL. Tracing out distinct mutual information for head or tail samples, BCL encourages to learn more generalizable and robust representation on the long-tailed dataset compared to the baselines from the loss and model perspectives.

4.5. Ablation Study

On augmentation components. In Table 5, we conduct various experiments to investigate the effect of individual augmentation components in \mathcal{A} of BCL. Specifically, we set the augmentation number $k = 1$ and additionally add each component to the sampled subset of augmentations. In this way, the monitored component dominants to construct the information discrepancy in the training stage. We then evaluate the effect of each component by computing the difference on the linear probing accuracy compared with *Identity* augmentation (i.e. $k = 1$) on CIFAR-100-LT. As shown in Table 5, we can see that the geometric-related augmentations are more helpful for representation learning. In particular, *ShearX*, *ShearY* and *Cutout* significantly improve the linear probing accuracy by 0.69%, 0.90% and 0.82%, respectively. However, some color-related augmentations lead to the degeneration of the linear probing accuracy except *Posterize*, *Sharpness* and *Brightness*. Intuitively, the color distortion augmentations in standard setting might be enough for contrastive learning methods, while some geometric-related semantics can further be captured by BCL.

Augmentation w vs. w/o the memorization guidance. To study the importance of the memorization guidance for the augmentation, we compare to RandAugment combined with SimCLR and SDCLR. For fair comparison, we fix the strength of augmentation in RandAugment. Note that, non-BCL means adopting strong and uniform augmentation to all samples in the dataset. Therefore, the performance bias from the augmentation is decoupled in these experiments.

Table 5. Improvement of linear probing performance on additionally adopting each component relative to that with *Identity* augmentation for BCL on CIFAR-100-LT. $\Delta(\%)$ means the relative gain.

| Component | $\Delta(\%)$ | Component | $\Delta(\%)$ |
|------------|--------------|--------------|--------------|
| Identity | 0.00 | Equalize | -1.28 |
| ShearX | 0.69 | Solarize | -2.69 |
| ShearY | 0.90 | Posterize | 0.59 |
| TranslateX | 0.44 | Contrast | -0.2 |
| TranslateY | 0.37 | Color | -0.53 |
| Rotate | 0.13 | Brightness | -0.08 |
| Cutout | 0.82 | Sharpness | -0.06 |
| Invert | -4.38 | AutoContrast | -0.96 |

As shown in the left panel of Figure 4, we can see that BCL consistently outperforms Non-BCL on linear probing evaluation under different shot on CIFAR-100-LT. Specifically, BCL-I and BCL-D improve full-shot performance by 1.45% and 0.9%, compared with Non-BCL-I and Non-BCL-D. The results confirm the effectiveness of the tailness detection mechanism and memorization-boosted design in BCL.

Other contrastive learning backbones. We extend our BCL to another two representative contrastive learning methods: MoCoV2 (He et al., 2020; Chen et al., 2020b), SimSiam (Chen & He, 2021), as shown in the middle panel of Figure 4. From the results, we can see that BCL maintains a consistent gain over MoCoV2 and SimSiam. The improvements show that BCL is orthogonal to the current self-supervised learning methods in long-tailed scenarios.

Performance on balanced datasets. Following (Jiang et al., 2021b), we also validate our BCL on balanced subsets of CIFAR-100 to explore whether BCL can benefit “implicit imbalancedness” *w.r.t.*, atypical samples or sampling bias on balanced data. Results are shown in the right panel of Figure 4. Similarly, BCL boosts the linear probing performance by 1.56%, 1.77% and 1.64% under different evaluations.

Impact of β in the momentum loss Eq. (3) In the left panel




Figure 4. (Left) Linear probing evaluation under different shots for BCL and Non-BCL (without the memorization guidance) pre-trained on CIFAR-100-LT. (Middle) Linear probing evaluation under different shots for MoCoV2 and SimSiam pre-trained on CIFAR-100-LT. (Right) Linear probing evaluation under different shots for BCL pre-trained on CIFAR-100, compared with SimCLR and SDCLR.

of Figure 5, we conduct several experiments with different β value to validate the stability of BCL. We compare different β in a high range (0.85-0.99) as the longer observations of the memorization effect are preferred to construct a reliable tail discovery. From the curve, we can see that BCL is mostly promising as the performance fluctuates a little.

Different augmentation number k . In the right panel of Figure 5, we validate BCL by training with different numbers of augmentations sampled from RandAugment. We can see that BCL achieves the appealing results with $k = 1, 2$ but degenerates at settings with the higher augmentation number k . Specifically, our method achieves 54.90% and 54.68% when adopting $k = 1, 2$ for the RandAugment, and 52.95%, 52.29%, 51.68% for $k = 3, 4, 5$, respectively. The performance difference reaches 3.22% between $k = 1$ and $k = 5$. We trace several augmented views and find that they are extremely distorted with limited information available when adopting $k = 5$ for RandAugment. We conjecture that too strong augmentation may lead to too much information loss and it becomes hard for BCL to encode the important details to the representation. On the other hand, a smaller k is also preferred due to the small computational cost.

5. Conclusion

In this paper, we propose a novel Boosted Contrastive Learning (BCL) method for the representation learning under the long-tailed data distribution. It leverages the clues of memorization effect in the historical training losses to automatically construct the information discrepancy for head and tail samples, which then drives contrastive learning to pay more attention to the tail samples. Different from previous methods that builds in the perspective of the loss or the model, BCL is essentially from the data perspective and orthogonal to the early explorations. Through extensive experiments, we demonstrate the effectiveness of BCL under different settings. In the future, we will extend BCL to more challenging long-tailed data like iNaturalist and explore the properties of the tail samples in more practical scenarios.




Figure 5. (Left) Linear probing performance under different β for BCL-I on CIFAR-100-LT. (Right) Linear probing performance with different k for BCL-I on CIFAR-100-LT.

Acknowledgement

This work is partially supported by the National Key R&D Program of China (No. 2019YFB1804304), 111 plan (No. BP0719010), STCSM (No. 18DZ2270700, No. 21DZ1100100), and State Key Laboratory of UHD Video and Audio Production and Presentation. BH was supported by the RGC Early Career Scheme No. 22200720, NSFC Young Scientists Fund No. 62006202, and Guangdong Basic and Applied Basic Research Foundation No. 2022A1515011652.

References

Arpit, D., Jastrzebski, S., Ballas, N., Krueger, D., Bengio, E., Kanwal, M. S., Maharaj, T., Fischer, A., Courville, A., Bengio, Y., et al. A closer look at memorization in deep networks. In *International Conference on Machine Learning*, pp. 233–242. PMLR, 2017.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33: 1877–1901, 2020.

- Cao, K., Wei, C., Gaidon, A., Arechiga, N., and Ma, T. Learning imbalanced datasets with label-distribution-aware margin loss. *Advances in Neural Information Processing Systems*, 32:1567–1578, 2019.
- Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., and Joulin, A. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33:9912–9924, 2020.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, pp. 1597–1607. PMLR, 2020a.
- Chen, X. and He, K. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15750–15758, 2021.
- Chen, X., Fan, H., Girshick, R., and He, K. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020b.
- Cubuk, E. D., Zoph, B., Shlens, J., and Le, Q. V. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 702–703, 2020.
- Cui, J., Zhong, Z., Liu, S., Yu, B., and Jia, J. Parametric contrastive learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 715–724, 2021.
- Cuturi, M. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in Neural Information Processing Systems*, 26:2292–2300, 2013.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255. Ieee, 2009.
- Doersch, C., Gupta, A., and Efros, A. A. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1422–1430, 2015.
- Ermolov, A., Siarohin, A., Sangineto, E., and Sebe, N. Whitening for self-supervised representation learning. In *International Conference on Machine Learning*, pp. 3015–3024. PMLR, 2021.
- Feldman, V. Does learning require memorization? a short tale about a long tail. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 954–959, 2020.
- Grill, J.-B., Strub, F., Alché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., Piot, B., kavukcuoglu, k., Munos, R., and Valko, M. Bootstrap your own latent - a new approach to self-supervised learning. In *Advances in Neural Information Processing Systems*, volume 33, pp. 21271–21284, 2020.
- Han, B., Yao, Q., Yu, X., Niu, G., Xu, M., Hu, W., Tsang, I. W., and Sugiyama, M. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in Neural Information Processing Systems*, 2018.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9729–9738, 2020.
- Jiang, L., Zhou, Z., Leung, T., Li, L.-J., and Fei-Fei, L. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International Conference on Machine Learning*, pp. 2304–2313. PMLR, 2018.
- Jiang, Z., Chen, T., Chen, T., and Wang, Z. Improving contrastive learning on imbalanced data via open-world sampling. *Advances in Neural Information Processing Systems*, 34, 2021a.
- Jiang, Z., Chen, T., Mortazavi, B. J., and Wang, Z. Self-damaging contrastive learning. In *International Conference on Machine Learning*, pp. 4927–4939. PMLR, 2021b.
- Jiang, Z., Zhang, C., Talwar, K., and Mozer, M. C. Characterizing structural regularities of labeled data in over-parameterized models. In *International Conference on Machine Learning*, volume 139, pp. 5034–5044. PMLR, 2021c.
- Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., and Kalantidis, Y. Decoupling representation and classifier for long-tailed recognition. In *International Conference on Learning Representations*, 2019.
- Kang, B., Li, Y., Xie, S., Yuan, Z., and Feng, J. Exploring balanced feature spaces for representation learning. In *International Conference on Learning Representations*, 2020.

- Khosla, A., Jayadevaprakash, N., Yao, B., and Li, F.-F. Novel dataset for fine-grained image categorization: Stanford dogs. In *Proc. CVPR Workshop on Fine-Grained Visual Categorization (FGVC)*, volume 2. Citeseer, 2011.
- Krause, J., Stark, M., Deng, J., and Fei-Fei, L. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE International Conference on Computer Vision workshops*, pp. 554–561, 2013.
- Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., and Soricut, R. Albert: A lite bert for self-supervised learning of language representations. *International Conference on Learning Representation*, 2019.
- Li, T., Cao, P., Yuan, Y., Fan, L., Yang, Y., Feris, R., Indyk, P., and Katabi, D. Targeted supervised contrastive learning for long-tailed recognition. *arXiv preprint arXiv:2111.13998*, 2021.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollár, P. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2980–2988, 2017.
- Liu, H., HaoChen, J. Z., Gaidon, A., and Ma, T. Self-supervised learning is more robust to dataset imbalance. *arXiv preprint arXiv:2110.05025*, 2021.
- Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., and Yu, S. X. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2537–2546, 2019.
- Maji, S., Rahtu, E., Kannala, J., Blaschko, M., and Vedaldi, A. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.
- Menon, A. K., Jayasumana, S., Jain, H., Veit, A., Kumar, S., and Rawat, A. S. Long-tail learning via logit adjustment. In *International Conference on Learning Representations*, 2021.
- Reed, W. J. The pareto, zipf and other power laws. *Economics letters*, 74(1):15–19, 2001.
- Ren, M., Zeng, W., Yang, B., and Urtasun, R. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning*, pp. 4334–4343. PMLR, 2018.
- Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., and Isola, P. What makes for good views for contrastive learning? In *Advances in Neural Information Processing Systems*, pp. 6827–6839, 2020.
- Tian, Y., Henaff, O. J., and van den Oord, A. Divide and contrast: Self-supervised learning from uncurated data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10063–10074, 2021.
- Van Horn, G., Branson, S., Farrell, R., Haber, S., Barry, J., Ipeirotis, P., Perona, P., and Belongie, S. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 595–604, 2015.
- Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., and Belongie, S. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8769–8778, 2018.
- Wah, C., Branson, S., Welinder, P., Perona, P., and Belongie, S. The caltech-ucsd birds-200-2011 dataset. 2011.
- Wang, T. and Isola, P. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pp. 9929–9939. PMLR, 2020.
- Wang, X. and Gupta, A. Unsupervised learning of visual representations using videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2794–2802, 2015.
- Yang, Y. and Xu, Z. Rethinking the value of labels for improving class-imbalanced learning. In *NeurIPS*, 2020.
- Yu, X., Han, B., Yao, J., Niu, G., Tsang, I., and Sugiyama, M. How does disagreement help generalization against label corruption? In *International Conference on Machine Learning*, pp. 7164–7173. PMLR, 2019.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.
- Zhang, Y., Kang, B., Hooi, B., Yan, S., and Feng, J. Deep long-tailed learning: A survey. *arXiv preprint arXiv:2110.04596*, 2021.
- Zheng, H., Yao, J., Zhang, Y., Tsang, I. W., and Wang, J. Understanding vaes in fisher-shannon plane. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 5917–5924, 2019.
- Zheng, H., Chen, X., Yao, J., Yang, H., Li, C., Zhang, Y., Zhang, H., Tsang, I., Zhou, J., and Zhou, M. Contrastive attraction and contrastive repulsion for representation learning. *arXiv preprint arXiv:2105.03746*, 2021.
- Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., and Torralba, A. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1452–1464, 2017.