

Parametric Contrastive Learning

Jiequan Cui¹ Zhisheng Zhong¹ Shu Liu² Bei Yu^{1,2} Jiaya Jia^{1,2}
¹The Chinese University of Hong Kong ²SmartMore

{jqcui, zszhong21, byu, leojia}@cse.cuhk.edu.hk, {liushuhust}@gmail.com

Abstract

In this paper, we propose *Parametric Contrastive Learning (PaCo)* to tackle long-tailed recognition. Based on theoretical analysis, we observe supervised contrastive loss tends to bias on high-frequency classes and thus increases the difficulty of imbalanced learning. We introduce a set of parametric class-wise learnable centers to rebalance from an optimization perspective. Further, we analyze our PaCo loss under a balanced setting. Our analysis demonstrates that PaCo can adaptively enhance the intensity of pushing samples of the same class close as more samples are pulled together with their corresponding centers and benefit hard example learning. Experiments on long-tailed CIFAR, ImageNet, Places, and iNaturalist 2018 manifest the new state-of-the-art for long-tailed recognition. On full ImageNet, models trained with PaCo loss surpass supervised contrastive learning across various ResNet backbones, e.g., our ResNet-200 achieves 81.8% top-1 accuracy. Our code is available at <https://github.com/dvlab-research/Parametric-Contrastive-Learning>.

1. Introduction

Convolutional neural networks (CNNs) have achieved great success in various tasks, including image classification [22, 43], object detection [31, 34] and semantic segmentation [55]. With neural network search [60, 33, 45, 13, 4], performance of CNNs further boosts. Impressive progress highly depends on large-scale and high-quality datasets, such as ImageNet [40], MS COCO [32] and Places [59]. When dealing with real-world applications, generally we face the long-tailed distribution problem – a few classes contain many instances, while most classes contain only a few instances. Learning in such an imbalanced setting is challenging as the low-frequency classes can be easily overwhelmed by high-frequency ones. Without considering this situation, CNNs will suffer from significant performance degradation.




Figure 1: Comparison with state-of-the-arts on ImageNet-LT [35]. Inference time is calculated with a batch of 64 images on Nvidia GeForce 2080Ti GPU. The same experimental setting is adopted for comparison. ResNet-50, ResNeXt-50, and ResNeXt101 are used for Balanced Softmax [38], Decouple [28], and PaCo. For RIDE [49], various numbers of expert with RIDEResNet and RIDEResNeXt are adopted. PaCo significantly outperforms recent SOTA. Detailed numbers for RIDE is in the supplementary file.

Contrastive learning [9, 21, 10, 19, 7] is a major research topic due to its success in self-supervised representation learning. Khosla *et al.* [29] extend non-parametric contrastive loss into non-parametric supervised contrastive loss by leveraging label information, which trains representation in the first stage and learns the linear classifier with the fixed backbone in the second stage. Though supervised contrastive learning works well in a balanced setting, for imbalanced datasets, our theoretical analysis shows that high-frequency classes will have a higher lower bound of loss and contribute much higher importance than low-frequency classes when equipping it in training.

This phenomenon leads to model bias on high-frequency classes and increases the difficulty of imbalanced learning. As shown in Fig. 2, when the model is trained with supervised contrastive loss on ImageNet-LT, the gradient norm varying from the most frequent class to the least one is rather steep. In particular, the gradient norm dramatically decreases for the top 200 most frequent classes.

Previous work [1, 25, 15, 20, 8, 41, 38, 28, 49, 14, 46, 14, 57] explored rebalancing in traditional supervised cross-




Figure 2: Rebalance in contrastive learning. We collect the average L2 norm of the gradient of weights in the last classifier layer on ImageNet-LT. Category indices are sorted by their image counts. The gradient norm varying from the most frequent class to the least one is steep for supervised contrastive learning [29]. In particular, the gradient norm dramatically decreases for the top 200 most frequent classes. Trained with PaCo, the gradient norm is better balanced.

entropy learning. In this paper, we tackle the above mentioned imbalance issue in supervised contrastive learning and make use of contrastive learning for long-tailed recognition. To our knowledge, it is the first attempt of using rebalance in contrastive learning.

To rebalance in supervised contrastive learning, we introduce a set of parametric class-wise learnable centers into supervised contrastive learning. We name our algorithm **Parametric Contrastive Learning (PaCo)** shown in Fig. 3. With such a simple and yet effective operation, we theoretically prove that the optimal values for the probability that two samples are a true positive pair (belonging to the same class), varying from the most frequent class to the least frequent class, are more balanced. Thus their lower bound of loss values are better organized. This phenomenon means the model takes more care of low-frequency classes, making the PaCo loss benefit imbalanced learning. Fig. 2 shows that, with our PaCo loss in training, gradient norm varying from the most frequent class to the least one are moderated better than supervised contrastive learning, which matches our analysis.

Further, we analyze the PaCo loss under a balanced setting. Our analysis demonstrates that with more samples clustered around their corresponding centers in training, the PaCo loss increases the intensity of pushing samples of the same class close, which benefits hard examples learning.

Finally, we conduct experiments on long-tailed version of CIFAR [16, 5], ImageNet [35], Places [35] and iNaturalist 2018 [48]. Experimental results show that we create a new record for long-tailed recognition. We also conduct experiments on full ImageNet [40] and CIFAR [30]. ResNet

models trained with PaCo also outperform the ones by supervised contrastive learning on such balanced datasets. Our key contributions are as follows.

- We identify the shortcoming of supervised contrastive learning under an imbalanced setting – it tends to bias towards high-frequency classes.
- We extend supervised contrastive loss to the PaCo loss, which is more friendly to imbalance learning, by introducing a set of parametric class-wise learnable centers.
- Equipped with the PaCo loss, we create new record across various benchmarks for long-tailed recognition. Moreover, experimental results on full ImageNet and CIFAR validate the effectiveness of PaCo under a balanced setting.

2. Related Work

Re-sampling/re-weighting. The most classical way to deal with long-tailed datasets is to over-sample low-frequency class images [41, 56, 2, 3] or under-sample high-frequency class images [20, 27, 2]. However, Oversampling can suffer from heavy over-fitting to low-frequency classes especially on small datasets. For under-sampling, discarding a large portion of high-frequency class data inevitably causes degradation of the generalization ability of CNNs. Re-weighting [23, 24, 51, 39, 42, 26] the loss functions is an alternative way to rebalance by either enlarging weights on more challenging and sparse classes or randomly ignoring gradients from high-frequency classes [44]. However, with large-scale data, re-weighting makes CNNs difficult to optimize during training [23, 24].

One/two-stage Methods. Since deferred re-weighting and re-sampling were proposed by Cao *et al.* [5], Kang *et al.* [28] and Zhou *et al.* [58] observed re-weighting or re-sampling strategies could benefit classifier learning while hurting representation learning. Kang *et al.* [28] proposed to decompose representation and classifier learning. It first trains the CNNs with uniform sampling, and then fine-tune the classifier with class-balanced sampling while keeping parameters of representation learning fixed. Zhou *et al.* [58] proposed one cumulative learning strategy, with which they bridge representation learning and classifier re-balancing.

The two-stage design is not for end-to-end frameworks. Tang *et al.* [46] analyzed the reason from the perspective of causal graph and concluded that the bad momentum causal effects played a vital role. Cui *et al.* [14] proposed residual learning mechanism to address this issue.

Non-parametric Contrastive Loss. Contrastive learning [9, 21, 10, 19, 7] is a framework that learns similar/dissimilar

representations from data that are organized into similar/dissimilar pairs. An effective contrastive loss function, called InfoNCE [47], is

$$\mathcal{L}_{q,k^+,\{k^-\}} = -\log \frac{\exp(q \cdot k^+ / \tau)}{\exp(q \cdot k^+ / \tau) + \sum_{k^-} \exp(q \cdot k^- / \tau)}, \quad (1)$$

where q is a query representation, k^+ is for the positive (similar) key sample, and $\{k^-\}$ denotes negative (dissimilar) key samples. τ is a temperature hyper-parameter. In the instance discrimination pretext task [53] for self-supervised learning, a query and a key form a positive pair if they are data-augmented versions of the same image. It forms a negative pair otherwise.

Traditional cross-entropy with linear fc layer weight w and true label y among n classes is expressed as

$$\mathcal{L}_{cross-entropy} = -\log \frac{\exp(q \cdot w_y)}{\sum_{i=1}^n \exp(q \cdot w_i)}. \quad (2)$$

Compared to it, InfoNCE does not get involved with parametric learnable parameters. To distinguish our proposed parametric contrastive learning from previous ones, we treat the InfoNCE as a non-parametric contrastive loss following [54].

Chen *et al.* [9] used self-supervised contrastive learning SimCLR to first match the performance of a supervised ResNet-50 with only a linear classifier trained on self-supervised representation on full ImageNet. He *et al.* [21] proposed MoCo and Chen *et al.* [10] extended MoCo to MoCo v2, with which small batch size training can also achieve competitive results on full ImageNet [40]. In addition, many other methods [19, 7] are also proposed to further boost performance.

3. Parametric Contrastive Learning

3.1. Supervised Contrastive Learning

Khosla *et al.* [29] extended the self-supervised contrastive loss with label information into supervised contrastive loss. Here we present it in the framework of MoCo [21, 10] as

$$\mathcal{L}_i = - \sum_{z_+ \in P(i)} \log \frac{\exp(z_+ \cdot T(x_i))}{\sum_{z_k \in A(i)} \exp(z_k \cdot T(x_i))}. \quad (3)$$

MoCo framework [21, 10] consists of two networks with the same structure, *i.e.*, *query network* and *key network*. The key network is driven by a momentum update with the query network in training. For each network, it usually contains one encoder CNN and one two-layer MLP transform.

During training, for one two-viewed image batch $B = (B_{v1}, B_{v2})$ and label y , B_{v1} and B_{v2} are fed into the query

Table 1: Top-1 accuracy (%) of supervised contrastive learning on ImageNet-LT with ResNet-50. Implementation details are in supplementary file. †represents that the model is trained with PaCo loss without center learning rebalance.

Method	Many	Medium	Few	All
Cross-Entropy	67.5	42.6	13.7	48.4
SupCon	53.4	2.9	0	22.0
PaCo (ours) †	69.6	45.8	16.0	51.0

network and key network respectively and we denote their outputs as Z_{v1} and Z_{v2} . Especially, Z_{v2} is used to update the momentum *queue*.

In Eq. (3), x_i is the representation for image X_i in B_{v1} obtained by the encoder of query network. The transform $T(\cdot)$ also belongs to the query network. We write

$$A(i) = \{z_k \in queue \cup Z_{v1} \cup Z_{v2}\} \setminus \{z_k \in Z_{v1} : k = i\},$$

$$P(i) = \{z_k \in A(i) : y_k = y_i\}.$$

In implementation, the loss is usually scaled by $\frac{1}{|P(i)|}$ and a temperature τ is applied like in Eq. (1). Different from self-supervised contrastive loss, which treats query and key as a positive pair if they are the data-augmented version of the same image, supervised contrastive loss treats them as one positive pair if they belong to the same class.

3.2. Theoretical Motivation

Analysis of Supervised Contrastive Learning. Khosla *et al.* [29] introduced supervised contrastive learning to encourage more compact representation. We observe that it is not directly applicable to long-tailed recognition. As shown in Table 1, the performance significantly decreases compared with traditional supervised cross-entropy. From an optimization point of view, supervised contrastive loss concentrates more on high-frequency classes than low-frequency ones, which is unfriendly for imbalanced learning.

Remark 1 (Optimal value for supervised contrastive learning). When supervised contrastive loss converges, the optimal value for the probability that two samples are a true positive pair with label y is $\frac{1}{K_y}$, where, $q(y)$ is the frequency of class y over the whole dataset, *queue* is the momentum *queue* in MoCo [21, 10] and $K_y \approx \text{length}(\text{queue}) \cdot q(y)$.

Interpretation. As indicated by Remark 1, high-frequency classes have a higher lower bound of loss value and contribute much higher importance than low-frequency classes in training. Thus the training process can be dominated by high-frequency classes. To handle this issue, we introduce a set of parametric class-wise learnable centers for rebalancing in contrastive learning.




Figure 3: Parametric contrastive learning (PaCo). We introduce a set of parametric class-wise learnable centers for rebalancing in contrastive learning. More analysis is in Section 3.3 for PaCo.

3.3. Rebalance in Contrastive Learning

As described in Fig. 3, we introduce a set of parametric class-wise learnable centers $\mathbf{C} = \{c_1, c_2, \dots, c_n\}$ into the original supervised contrastive learning, and call this new form **Parametric Contrastive Learning (PaCo)**. Correspondingly, the loss function is changed to

$$\mathcal{L}_i = \sum_{z_+ \in P(i) \cup \{c_y\}} -w(z_+) \log \frac{\exp(z_+ \cdot T(x_i))}{\sum_{z_k \in A(i) \cup \mathbf{C}} \exp(z_k \cdot T(x_i))}, \quad (4)$$

where

$$w(z_+) = \begin{cases} \alpha, & z_+ \in P(i) \\ 1.0, & z_+ \in \{c_y\} \end{cases}$$

and

$$z \cdot T(x_i) = \begin{cases} z \cdot \mathcal{G}(x_i), & z \in A(i) \\ z \cdot \mathcal{F}(x_i), & z \in \mathbf{C}. \end{cases}$$

Following Chen *et al.* [10], the transform $\mathcal{G}(\cdot)$ is a two-layer MLP while $\mathcal{F}(\cdot)$ is the identity mapping, *i.e.*, $\mathcal{F}(x) = x$. α is one hyper-parameter in $(0, 1)$. $P(i)$ and $A(i)$ are the same with supervised contrastive learning in Eq. (3). In implementation, the loss is scaled by $\frac{1}{\sum_{z_+ \in P(i) \cup \{c_y\}} w(z_+)}$ and a temperature τ is applied like that in Eq. (3).

Remark 2 (Optimal value for parametric contrastive learning) When parametric contrastive loss converges, the optimal value for the probability that two samples are a true positive pair with label y is $\frac{\alpha}{1 + \alpha \cdot K_y}$ and the optimal value for the probability that a sample is closest to its corresponding center c_y among \mathbf{C} is $\frac{1}{1 + \alpha \cdot K_y}$, where $q(y)$ is the frequency of class y over the whole dataset, *queue* is the momentum queue in MoCo [21, 10] and $K_y \approx \text{length}(\text{queue}) \cdot q(y)$.




Figure 4: Curve for \mathcal{L}_{extra} .

Interpretation. Suppose the most frequent class y_h has $K_{y_h} \approx q(y_h) \cdot \text{length}(\text{queue})$ and the least frequent class y_t has $K_{y_t} \approx q(y_t) \cdot \text{length}(\text{queue})$. As indicated by Remarks 2 and 1, the optimal value for the probability that two samples are a true positive pair varying from the most frequent class to the least one is rebalanced from $\frac{1}{K_{y_h}} \rightarrow \frac{1}{K_{y_t}}$ to $\frac{1}{\frac{1}{\alpha} + K_{y_h}} \rightarrow \frac{1}{\frac{1}{\alpha} + K_{y_t}}$. The smaller α is, the more uniform the optimal value from the most frequent class to the least one is, friendly to low-frequency classes learning.

However, when α decreases, the intensity of contrast among samples becomes weaker, the intensity of contrast between samples and centers is stronger. The whole loss becomes closer to supervised cross-entropy. To make good use of contrastive learning and rebalance at the same time, we observe that $\alpha=0.05$ is a reasonable choice.

3.4. PaCo under Balanced Setting

For balanced datasets, all classes have the same frequency, *i.e.*, $q^*=q(y_i)=q(y_j)$ and $K^*=K_{y_i}=K_{y_j}$ for any class y_i and class y_j . In this case, PaCo reduces to an improved version of multi-task with the weighted sum of supervised cross-entropy loss and supervised contrastive loss. The connection between PaCo and multi-task is

$$\text{ExpSum} = \sum_{c_k \in \mathbf{C}} \exp(c_k \cdot \mathcal{F}(x_i)) + \sum_{z_k \in A(i)} \exp(z_k \cdot \mathcal{G}(x_i)).$$

We also write the PaCo loss as

$$\begin{aligned}
\mathcal{L}_i &= \sum_{z_+ \in P(i) \cup \{c_y\}} -w(z_+) \log \frac{\exp(z_+ \cdot T(x_i))}{\sum_{z_k \in A(i) \cup \mathbf{C}} \exp(z_k \cdot T(x_i))} \\
&= -\log \frac{\exp(c_y \cdot \mathcal{F}(x_i))}{ExpSum} - \alpha \sum_{z_+ \in P(i)} \log \frac{\exp(z_+ \cdot \mathcal{G}(x_i))}{ExpSum} \\
&= \mathcal{L}_{sup} + \alpha \mathcal{L}_{supcon} - (\log P_{sup} + \alpha K^* \log P_{supcon}) \\
&= \mathcal{L}_{sup} + \alpha \mathcal{L}_{supcon} - (\log P_{sup} + \alpha K^* \log(1 - P_{sup})),
\end{aligned}$$

$$\text{where } \begin{cases} P_{sup} = \frac{\sum_{c_k \in \mathbf{C}} \exp(c_k \cdot \mathcal{F}(x_i))}{ExpSum}; \\ P_{supcon} = \frac{\sum_{z_k \in A(i)} \exp(z_k \cdot \mathcal{G}(x_i))}{ExpSum}. \end{cases} \quad (5)$$

Multi-task learning combines supervised cross-entropy loss and supervised contrastive loss with a fixed weighted scalar. When these two losses conflict, training can suffer from slow convergence or sub-optimization. Our PaCo contrarily adjusts the intensity of supervised cross-entropy loss and supervised contrastive loss in an adaptive way and potentially avoids conflict as analyzed in the following.

3.4.1 Analysis of PaCo under Balanced Setting

As indicated by Eq. (5), compared with multi-task, PaCo has an additional loss item of

$$\mathcal{L}_{extra} = -\log(P_{sup}) - \alpha K^* \log(1 - P_{sup}). \quad (6)$$

Here, we take full ImageNet as an example, *i.e.*, $q^* = 0.001$, $\text{length}(\text{queue}) = 8192$, $\alpha = 0.05$, and $\alpha K^* = 0.41$. Then the function curve for \mathcal{L}_{extra} is shown in Fig. 4. With P_{sup} increases from 0 to 1.0, the function value decreases until $P_{sup} = 0.71$ and then goes up, which implies \mathcal{L}_{extra} obtains the smallest loss value when $P_{sup} = 0.71$. Note that, when the whole PaCo loss in Eq. (4) achieves the optimal solution, $P_{sup} = 0.71$ still holds as demonstrated in Remark 2. With P_{sup} increases during training, we analyze how it affects the intensity of supervised contrastive loss and supervised cross-entropy loss in the following.

Adaptive Weighting between \mathcal{L}_{sup} and \mathcal{L}_{supcon} . Note that the optimal value is 0.035 for the probability that two samples are a true positive pair with label y as indicated by Remark 2. We suppose $p_l, p_h \in (0, 0.71)$ and $p_l < p_h$. To achieve the optimal value, when $P_{sup}=p$, the supervised contrastive loss value \mathcal{L}_{supcon} decreases as shown in Eq. (7).

Here P_{sup} increases from p_l to p_h , \mathcal{L}_{supcon} decreases to a much smaller loss value to achieve the optimal solution, which implies the need to make two different class samples

$$\begin{aligned}
\mathcal{L}_{supcon} &= - \sum_{z_+ \in P(i)} \log \frac{\exp(z_+ \cdot \mathcal{G}(x_i))}{\sum_{z_k \in A(i)} \exp(z_k \cdot \mathcal{G}(x_i))} \\
&= - \sum_{z_+ \in P(i)} \log \frac{\frac{\exp(z_+ \cdot \mathcal{G}(x_i))}{ExpSum}}{\frac{\sum_{z_k \in A(i)} \exp(z_k \cdot \mathcal{G}(x_i))}{ExpSum}} \\
&= - K^* \log \frac{0.035}{1-p}.
\end{aligned} \quad (7)$$

much more discriminative, *i.e.*, increasing inter-class margins. Thus the intensity of supervised contrastive loss enlarges.

An intuition is that as P_{sup} increases, more samples are pulled together with their corresponding centers. Along with stronger intensity of supervised contrastive loss at that time, it is more likely to push hard examples close to those samples that are already around right centers.

3.5. Center Learning Rebalance

PaCo balances the contrastive learning (for moderating contrast among samples). However the center learning also needs to be balanced, which has been explored in [1, 25, 15, 20, 8, 41, 38, 28, 49, 14, 46, 18, 57]. We incorporate Balanced Softmax [38] into the center learning. Then the PaCo loss is changed from Eq. (4) to

$$\mathcal{L}_i = \sum_{z_+ \in P(i) \cup \{c_y\}} -w(z_+) \log \frac{\psi(z_+, T(x_i))}{\sum_{z_k \in A(i) \cup \mathbf{C}} \psi(z_k, T(x_i))}, \quad (8)$$

where

$$\psi(z_k, T(x_i)) = \begin{cases} \exp(z_k \cdot \mathcal{G}(x_i)), & z_k \in A(i); \\ \exp(z_k \cdot \mathcal{F}(x_i)) \cdot q(y_k), & z_k \in \mathbf{C}. \end{cases}$$

We emphasize that Balanced Softmax is only a practical remedy for center learning rebalance. Theoretical analysis remains future work.

4. Experiments

4.1. Ablation Study

Data augmentation strategy for PaCo. Data augmentation is the key for success of contrastive learning as indicated by Chen *et al.* [9]. For PaCo, we also conduct ablation studies for different augmentation strategies. Several observations are intriguingly different from those of [37]. We experiment with the following ways of data augmentation.

- SimAugment: an augmentation policy [21, 10] that applies random flips and color jitters followed by Gaussian blur.
- RandAugment [12]: A two stage augmentation policy that uses random parameters in place of parameters

Table 2: Comparison with different augmentation strategies for PaCo on ImageNet-LT with ResNet-50.

Methods	Aug. strategy	Top-1 Accuracy
PaCo	Strategy (1)	55.0
PaCo	Strategy (2)	56.5
PaCo	Strategy (3)	57.0

tuned by AutoAugment. The random parameters do not need to be tuned and hence reduces the search space.

For the common *random resized crop* used along with the above two strategies, work of [37] explains that the optimal hyper-parameter for random resized crop is (0.2,1) in self-supervised contrastive learning. This setting is also adopted by work of [9, 21, 10, 19, 7].

However, in this paper, we observe severe performance degradation on ImageNet-LT with ResNet-50 (55.0% vs 52.2%) for PaCo when we change the hyper-parameter from (0.08,1) to (0.2, 1). This is because PaCo involves center learning while other self-supervised frameworks only apply non-parametric contrastive loss as described in Section 2. Note that the same phenomenon is also observed on traditional supervised learning with cross-entropy loss.

Another observation is on RandAugment [12]. The work of [50] demonstrated that directly applying strong data augmentation in MoCo [21, 10] does not work well. Here we observe a similar conclusion with RandAugment [12]. We experiment with 3 different augmentation strategies of (1) encoder and momentum encoder input both with SimAugment; (2) encoder and momentum encoder input both with RandAugment; and (3) encoder input using RandAugment and momentum encoder input with SimAugment. The experimental results are presented in Table 2. With strategy (3), PaCo achieves the best performance, showing center learning requires more aggressive data augmentation compared with contrastive learning among samples.

4.2. Baseline Implementation

Contrastive learning benefits from longer training compared with traditional supervised learning with cross-entropy as Chen *et al.* [9] concluded, which is also validated by previous work of [9, 21, 10, 19, 7]. We run PaCo with 400 epochs on CIFAR-LT, ImageNet-LT, iNaturalist 2018, full CIFAR, and full ImageNet except for Places-LT. With Places-LT, we follow previous work [35, 14] by loading the pre-trained model on ImageNet and finely tune 30 epochs on Places-LT. For fair comparison on ImageNet-LT and iNaturalist 2018, we reimplement baselines with the same training time and RandAugment [12]. Especially, for RIDE, based on model ensemble, we compare with it under comparable inference latency in Fig. 1.

4.3. Long-tailed Recognition

We follow the common evaluation protocol [35, 14, 28] in long-tailed recognition – that is, training models on the long-tailed source label distribution and evaluating their performance on the uniform target label distribution. We conduct experiments on long-tailed version of CIFAR-100 [16, 5], Places [35], ImageNet [35] and iNaturalist 2018 [48] datasets.

CIFAR-100-LT datasets. CIFAR-100 has 60,000 images – 50,000 for training and 10,000 for validation with 100 categories. For fair comparison, we use the long-tailed version of CIFAR datasets with the same setting as those used in [6, 58, 16]. They control the degrees of data imbalance with an imbalance factor β . $\beta = \frac{N_{max}}{N_{min}}$ where N_{max} and N_{min} are the numbers of training samples for the most and least frequent classes respectively. Following [58], we conduct experiments with imbalance factors 100, 50, and 10.

ImageNet-LT and Places-LT. ImageNet-LT and Places-LT were proposed in [35]. ImageNet-LT is a long-tailed version of ImageNet dataset [40] by sampling a subset following the Pareto distribution with power value $\alpha=6$. It contains 115.8K images from 1,000 categories, with class cardinality ranging from 5 to 1,280. Places-LT is a long-tailed version of the large-scale scene classification dataset Places [59]. It consists of 184.5K images from 365 categories with class cardinality ranging from 5 to 4,980.

iNaturalist 2018. The iNaturalist 2018 [48] is one species classification dataset, which is on a large scale and suffers from extremely imbalanced label distribution. It is composed of 437.5K images from 8,142 categories. In addition to the extreme imbalance, the iNaturalist 2018 dataset also confronts the fine-grained problem [52].

Implementation details. For image classification on ImageNet-LT, we used ResNet-50, ResNeXt-50-32x4d, and ResNeXt-101-32x4d as our backbones for experiments. For iNaturalist 2018, we conduct experiments with ResNet-50 and ResNet-152. All models were trained using SGD optimizer with momentum $\mu = 0.9$. Contrastive learning usually requires long training time to converge. MoCo [21, 10], BYOL [19] and SWAV [7] train 800 epochs for model convergence. Supervised contrastive learning [29] trains 350 epochs for feature learning and another 350 epochs for classifier learning.

Following MoCo [21, 10], when we train models with PaCo, the learning rate decays by a cosine scheduler from 0.02 to 0 with batch size 128 on 4 GPUs in 400 epochs. The temperature is set to 0.2. α is 0.05. For fair comparison, we reimplement baselines in the same setting for recent

Table 3: Long-tail recognition accuracy on ImageNet-LT for different backbone architectures. † denotes models trained with RandAugment [12] in 400 epochs. More comparisons with RIDE [49] are in Fig. 1.

Method	ResNet-50	ResNeXt-50	ResNeXt-101
CE(baseline)	41.6	44.4	44.8
Decouple-cRT	47.3	49.6	49.4
Decouple- τ -norm	46.7	49.4	49.6
De-confound-TDE	51.7	51.8	53.3
ResLT	-	52.9	54.1
MiSLAS	52.7	-	-
Decouple- τ -norm †	54.5	56.0	57.9
Balanced Softmax †	55.0	56.2	58.0
PaCo†	57.0	58.2	60.0

state-of-the-arts of Decouple [28], Balanced Softmax [38] and RIDE [49] as mentioned in Section 4.2.

For Places-LT, following previous setting [35, 14], we choose ResNet-152 as the backbone network, pre-train it on the full ImageNet-2012 dataset (provided by torchvision), and finely tune it for 30 epochs on Places-LT. Same as that on ImageNet-LT, the learning rate decays by a cosine scheduler from 0.02 to 0 with batch size 128. The temperature is set to 0.2. α is 0.05. For CIFAR-100-LT, we strictly follow the setting of [38] for fair comparison. A smaller temperature of 0.05 and $\alpha = 0.02$ are adopted for PaCo.

Comparison on ImageNet-LT. Table 3 shows extensive experimental results for comparison with recent SOTA methods. We observe that Balanced Softmax [38] still achieves comparable results with Decouple [28] across various backbones under such strong training setting on ImageNet-LT, consistent with what is claimed in the original paper. For RIDE that is based on model ensemble, we analyze the real inference speed by calculating inference time with a batch of 64 images on Nvidia GeForce 2080Ti GPU.

We observe RIDEResNet with 3 experts even has higher inference latency than a standard ResNeXt-50-32x4d (**15.3ms vs 13.1ms**); RIDEResNeXt with 3 experts yields higher inference latency than a standard ResNeXt-101-32x4d (**26ms vs 25ms**). This result is in accordance with the conclusion that network fragmentation reduces the degree of parallelism and thus decreases efficiency in [36, 13]. For fair comparison, we do not apply knowledge distillation tricks for all these methods. As shown in Fig. 1 and Table 3, under comparable inference latency, PaCo significantly surpasses these baselines.

Comparison on Places-LT. The experimental results on Places-LT are summarized in Table 4. Due to the architecture change of RIDE, it is not applicable to load the publicly

Table 4: Performance on Places-LT [35], starting from an ImageNet pre-trained ResNet-152 provided by torchvision. †denotes the model trained with RandAugment [12].

Method	Many	Medium	Few	All
CE(baseline)	45.7	27.3	8.2	30.2
OLTR	44.7	37.0	25.3	35.9
Decouple- τ -norm	37.8	40.7	31.8	37.9
Balanced Softmax	42.0	39.3	30.5	38.6
ResLT	39.8	43.6	31.4	39.8
MiSLAS	39.6	43.3	36.1	40.4
RIDE (2 experts)	-	-	-	-
PaCo	37.5	47.2	33.9	41.2
PaCo †	36.1	47.9	35.3	41.2

pre-trained model on full ImageNet, while PaCo is more flexible where the network architecture is the same as those of [35, 14]. Under fair training setting by finely tuning 30 epochs without RandAugment, PaCo surpasses SOTA Balanced Softmax by 2.6%. An interesting observation is that RandAugment has little effect on the Places-LT dataset. A similar phenomenon can be observed on the iNaturalist 2018 dataset. More evaluation numbers are in the supplementary file. They can be intuitively understood since RandAugment is designed for ImageNet classification, which inspires us to explore general augmentations across different domains.

Comparison on iNaturalist 2018. Table 5 lists experimental results on iNaturalist 2018. Under fair training setting, PaCo consistently surpasses recent SOTA methods of Decouple, Balanced Softmax and RIDE. Our method is 1.4% higher than Balanced Softmax. We also apply PaCo on large ResNet-152 architecture. And the performance boosts to **75.3%** top-1 accuracy. Note that we only transfer the hyper-parameters of PaCo on ImageNet-LT to iNaturalist 2018 without any change. Tuning hyper-parameters for PaCo will bring further improvement.

Comparison on CIFAR-100-LT. The experimental results on CIFAR-100-LT are listed in Table 6. For the CIFAR-100-LT dataset, we mainly compare with the SOTA method Balanced Softmax [38] with the same training setting where Cutout [17] and AutoAugment [11] are used in training. As shown in Table 6, PaCo consistently outperforms Balanced Softmax across different imbalance factors with such a strong setting. Specifically, PaCo surpasses Balanced Softmax by 1.2%, 1.8% and 1.2% under imbalance factor 100, 50 and 10 respectively, which testify the effectiveness of our PaCo method.

4.4. Full ImageNet and CIFAR Recognition

As analyzed in Section 3.4, for balanced datasets, PaCo reduces to an improved version of multi-task learning, which

Table 5: Top-1 accuracy over all classes on iNaturalist 2018 with ResNet-50. Knowledge distillation is not applied to all methods for fair comparison. We compare with RIDE ¹ under comparable inference latency. † denotes models trained with RandAugment [12] in 400 epochs.

Method	Top-1 Accuracy
CB-Focal	61.1
LDAM+DRW	68.0
Decouple- τ -norm	69.3
Decouple-LWS	69.5
BBN	69.6
ResLT	70.2
MiSLAS	71.6
RIDE (2 experts) †	69.5
Decouple- τ -norm †	71.5
Balanced Softmax †	71.8
PaCo †	73.2

Table 6: Top-1 accuracy on CIFAR-100-LT with different imbalance factors (†: models trained in same setting).

Dataset	CIFAR-100 LT		
	100	50	10
Focal Loss	38.4	44.3	55.8
LDAM+DRW	42.0	46.6	58.7
BBN	42.6	47.0	59.1
Causal Norm	44.1	50.3	59.6
ResLT	45.3	50.0	60.8
MiSLAS	47.0	52.3	63.2
Balanced Softmax †	50.8	54.2	63.0
PaCo †	52.0	56.0	64.2

adaptively adjusts the intensity of supervised cross-entropy loss and supervised contrastive loss. To verify the effectiveness of PaCo under this balanced setting, we conduct experiments on full ImageNet and full CIFAR. They are indicative to compare PaCo with supervised contrastive learning [29]. Note that, under full ImageNet and CIFAR, we remove the rebalance in center learning, *i.e.*, Balanced Softmax.

Full ImageNet. In the implementation, we transfer hyperparameters of PaCo on ImageNet-LT to full ImageNet without modification. SGD optimizer with momentum $\mu = 0.9$ is used. $\alpha=0.05$, temperature is 0.2 and queue size is 8,192. For multi-task training, the supervised contrastive loss is additional regularization and the loss weight is also set to 0.05. The same data augmentation strategy is applied as PaCo, which is discussed in Section 4.1.

The experimental results are summarized in Table 7. With SimAugment, our ResNet-50 model achieves 78.7% top-1 accuracy, which outperforms supervised contrastive learning model by 0.8%. Equipped with strong augmenta-

Table 7: Top-1 accuracy on full ImageNet with ResNets. “*” denotes supervised contrastive learning with additional operation of image warping before Gaussian blur.

Method	Model	augmentation	Top-1 Acc
Supcon	ResNet-50	SimAugment *	77.9
Supcon	ResNet-50	RandAugment	78.4
Supcon	ResNet-101	StackedRandAugment	80.2
multi-task	RandAugment	ResNet-50	78.1
PaCo	ResNet-50	SimAugment	78.7
PaCo	ResNet-50	RandAugment	79.3
PaCo	ResNet-101	StackedRandAugment	80.9
PaCo	ResNet-200	StackedRandAugment	81.8

Table 8: Top-1 accuracy on full CIFAR-100 (ResNet-50).

Method	dataset	Top-1 Accuracy
CE(baseline)	CIFAR-100	77.9
multi-task	CIFAR-100	78.0
Supcon	CIFAR-100	76.5
PaCo	CIFAR-100	79.1

tion, *i.e.*, RandAugment [12], the performance further improves to 79.3%. ResNet-101/200 trained with PaCo consistently surpass supervised contrastive learning.

Full CIFAR-100. For CIFAR implementation, we follow supervised contrastive learning and train ResNet-50 with only the SimAugment. Compared with full ImageNet, we adopt a smaller temperature of 0.07, $\alpha = 0.008$ and batch size 256 with learning rate 0.1. As shown in Table 8, on CIFAR-100, PaCo outperforms supervised contrastive learning by 2.6%, which validates the advantages of PaCo. Note that, following [13], we use a weight-decay of $5e - 4$.

5. Conclusion

In this paper, we have proposed Parametric Contrastive Learning (PaCo), which contains a set of parametric class-wise learnable centers to tackle the long-tailed recognition. It is based on the theoretical analysis of supervised contrastive learning. For balanced data, our analysis of PaCo demonstrates that it can adaptively enhance the intensity of pushing two samples of the same class close as more samples are pulled together with their corresponding centers, which can potentially benefit hard examples learning in training.

We conduct experiments on various benchmarks of CIFAR-LT, ImageNet-LT, Places-LT, and iNaturalist 2018. The experimental results show that we create a new state-of-the-art for long-tailed recognition. Further, experimental results on full ImageNet and CIFAR demonstrate that PaCo also benefits balanced datasets.

References

- [1] Mateusz Buda, Atsuto Maki, and Maciej A. Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 2018. 1, 5, 14
- [2] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 2018. 2
- [3] Jonathon Byrd and Zachary Lipton. What is the effect of importance weighting in deep learning? In *ICML*, 2019. 2
- [4] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. In *ICLR*, 2020. 1
- [5] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Aréchiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *NeurIPS*, 2019. 2, 6
- [6] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *NeurIPS*, 2019. 6
- [7] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *NeurIPS*, 2020. 1, 2, 3, 6
- [8] Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 2002. 1, 5, 14
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 1, 2, 3, 5, 6
- [10] Xinlei Chen, Haoqi Fan, Ross B. Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *CoRR*, 2020. 1, 2, 3, 4, 5, 6
- [11] Ekin D. Cubuk, Barret Zoph, Dandelion Mané, Vijay Vasudevan, and Quoc V. Le. Autoaugment: Learning augmentation strategies from data. In *CVPR*, 2019. 7
- [12] Ekin Dogus Cubuk, Barret Zoph, Jon Shlens, and Quoc Le. Randaugment: Practical automated data augmentation with a reduced search space. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *NeurIPS*, 2020. 5, 6, 7, 8
- [13] Jiequan Cui, Pengguang Chen, Ruiyu Li, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Fast and practical neural architecture search. In *ICCV*, 2019. 1, 7, 8
- [14] Jiequan Cui, Shu Liu, Zhuotao Tian, Zhisheng Zhong, and Jiaya Jia. Reslt: Residual learning for long-tailed recognition. *CoRR*, abs/2101.10633, 2021. 1, 2, 5, 6, 7, 14
- [15] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge J. Belongie. Class-balanced loss based on effective number of samples. In *CVPR*, 2019. 1, 5, 14
- [16] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *CVPR*, 2019. 2, 6, 14
- [17] Terrance Devries and Graham W. Taylor. Improved regularization of convolutional neural networks with cutout. *CoRR*, abs/1708.04552, 2017. 7
- [18] Rahul Duggal, Scott Freitas, Sunny Dhamnani, Duen Horng Jimeng Sun, et al. Elf: An early-exiting framework for long-tailed classification. *arXiv preprint arXiv:2006.11979*, 2020. 5, 14
- [19] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent - A new approach to self-supervised learning. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *NeurIPS*, 2020. 1, 2, 3, 6
- [20] Haibo He and Edwardo A Garcia. Learning from imbalanced data. *IEEE TKDE*, 2009. 1, 2, 5, 14
- [21] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 1, 2, 3, 4, 5, 6
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1
- [23] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *CVPR*, 2016. 2
- [24] Chen Huang, Yining Li, Change Loy Chen, and Xiaoou Tang. Deep imbalanced learning for face recognition and attribute prediction. *IEEE TPAMI*, 2019. 2
- [25] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *CVPR*, 2016. 1, 5, 14
- [26] Muhammad Abdullah Jamal, Matthew Brown, Ming-Hsuan Yang, Liqiang Wang, and Boqing Gong. Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. In *CVPR*, 2020. 2
- [27] Nathalie Japkowicz and Shaju Stephen. The class imbalance problem: A systematic study. *Intelligent Data Analysis*, 2002. 2
- [28] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. In *ICLR*, 2020. 1, 2, 5, 6, 7, 14
- [29] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *NeurIPS*, 2020. 1, 2, 3, 6, 8
- [30] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 2
- [31] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 1
- [32] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 1

- [33] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: differentiable architecture search. In *ICLR*. OpenReview.net, 2019. 1
- [34] Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, and Jiaya Jia. Path aggregation network for instance segmentation. In *CVPR*, 2018. 1
- [35] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X. Yu. Large-scale long-tailed recognition in an open world. In *CVPR*, 2019. 1, 2, 6, 7
- [36] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet V2: practical guidelines for efficient CNN architecture design. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *ECCV*, 2018. 7
- [37] Wanli Ouyang, Xiaogang Wang, Cong Zhang, and Xiaokang Yang. Factors in finetuning deep model for object detection with long-tail distribution. In *CVPR*, 2016. 5, 6
- [38] Jiawei Ren, Cunjun Yu, Shunan Sheng, Xiao Ma, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Balanced meta-softmax for long-tailed visual recognition. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *NeurIPS*, 2020. 1, 5, 7, 14
- [39] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *ICML*, 2018. 2
- [40] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, and Michael Bernstein. ImageNet large scale visual recognition challenge. *IJCV*, 2015. 1, 2, 3, 6
- [41] Li Shen, Zhouchen Lin, and Qingming Huang. Relay back-propagation for effective learning of deep convolutional neural networks. In *ECCV*, 2016. 1, 2, 5, 14
- [42] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. In *NeurIPS*, 2019. 2
- [43] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 1
- [44] Jingru Tan, Changbao Wang, Buyu Li, Quanquan Li, Wanli Ouyang, Changqing Yin, and Junjie Yan. Equalization loss for long-tailed object recognition. In *CVPR*, 2020. 2
- [45] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile. In *CVPR*, 2019. 1
- [46] Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. Long-tailed classification by keeping the good and removing the bad momentum causal effect. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *NeurIPS*, 2020. 1, 2, 5, 14
- [47] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *CoRR*, abs/1807.03748, 2018. 3
- [48] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The iNaturalist species classification and detection dataset. In *CVPR*, 2018. 2, 6
- [49] Xudong Wang, Long Lian, Zhongqi Miao, Ziwei Liu, and Stella X. Yu. Long-tailed recognition by routing diverse distribution-aware experts. *CoRR*, abs/2010.01809, 2020. 1, 5, 7, 14
- [50] Xiao Wang and Guo-Jun Qi. Contrastive learning with stronger augmentations, 2021. 6
- [51] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. In *NeurIPS*, 2017. 2
- [52] Xiu-Shen Wei, Peng Wang, Lingqiao Liu, Chunhua Shen, and Jianxin Wu. Piecewise classifier mappings: Learning fine-grained learners for novel categories with few examples. *IEEE TIP*, 2019. 6
- [53] Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 3
- [54] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 3
- [55] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017. 1
- [56] Q Zhong, C Li, Y Zhang, H Sun, S Yang, D Xie, and S Pu. Towards good practices for recognition & detection. In *CVPR workshops*, 2016. 2
- [57] Zhisheng Zhong, Jiequan Cui, Shu Liu, and Jiaya Jia. Improving calibration for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16489–16498, June 2021. 1, 5
- [58] Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. *arXiv preprint arXiv:1912.02413*, 2019. 2, 6
- [59] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE TPAMI*, 2018. 1, 6
- [60] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures for scalable image recognition. In *CVPR*, 2018. 1