Distribution Alignment: A Unified Framework for Long-tail Visual Recognition

Songyang Zhang[1,3,5,*] Zeming Li[2] Shipeng Yan[1] Xuming He[1,4] Jian Sun[2]
[1]ShanghaiTech University [2]Megvii Technology [3] University of Chinese Academy of Sciences
[4]Shanghai Engineering Research Center of Intelligent Vision and Imaging
[5]Shanghai Institute of Microsystem and Information Technology, Chinese Academy of Sciences
{zhangsy1, yanshp, hexm}@shanghaitech.edu.cn, {lizeming,sunjian}@megvii.com

Abstract

Despite the recent success of deep neural networks, it remains challenging to effectively model the long-tail class distribution in visual recognition tasks. To address this problem, we first investigate the performance bottleneck of the two-stage learning framework via ablative study. Motivated by our discovery, we propose a unified distribution alignment strategy for long-tail visual recognition. Specifically, we develop an adaptive calibration function that enables us to adjust the classification scores for each data point. We then introduce a generalized re-weight method in the two-stage learning to balance the class prior, which provides a flexible and unified solution to diverse scenarios in visual recognition tasks. We validate our method by extensive experiments on four tasks, including image classification, semantic segmentation, object detection, and instance segmentation. Our approach achieves the state-of-the-art results across all four recognition tasks with a simple and unified framework.

1. Introduction

While deep convolutional networks have achieved great successes in many vision tasks, it usually requires a large number of training examples for each visual category. More importantly, prior research mostly focuses on learning from a *balanced dataset* [22], where different object classes are approximately evenly distributed. However, for large-scale vision recognition tasks, partially due to the non-uniform distribution of natural object classes and varying annotation costs, we typically learn from datasets with a *long-tail* class label distribution. In such scenarios, the number of training instances per class varies significantly, from as few as one example for tail classes to hundreds or thousands for head classes [48, 26, 14, 51, 49, 36].

Figure 1: Per-class performance of the two-stage learning *baseline* and our *empirical classification bound* on ImageNet-LT val split. Two methods share the same representation while our bound setting retrains the classifier head with the balanced full dataset.

The intrinsic long-tail property of our visual data introduces a multitude of challenges for recognition in the wild [1], as a deep network model has to simultaneously cope with imbalanced annotations among the head and medium-sized classes, and few-shot learning in the tail classes. A naively learned model would be largely dominated by those few head classes while its performance is much degraded for many other tail classes.

Early works on re-balancing data distribution focus on learning one-stage models, which achieve limited successes due to lack of principled design in their strategies [2, 33, 3, 9, 26, 41]. More recent efforts aim to improve the long-tail prediction by decoupling the representation learning and classifier head learning [19, 28, 35, 38, 23]. However, such a two-stage strategy typically relies on heuristic design to adjust the decision boundary of the initially learned classifier head, which often requires tedious hyper-parameter tuning in practice. This severely limits its capacity to resolve the mismatch between *imbalanced training data distribution* and *balanced evaluation metrics*.

In this work, we first perform an ablative analysis on the two-stage learning strategy to shed light on its performance bottleneck. Specifically, our study estimates an 'ideal' clas-

*This work was done when Songyang Zhang was a research intern at Megvii Technology. This work was supported by Shanghai NSF Grant (No. 18ZR1425100), National Key R&D Program of China (No. 2017YFA0700800), and Beijing Academy of Artificial Intelligence (BAAI). Code is available: https://github.com/Megvii-BaseDetection/DisAlign sification accuracy using a balanced dataset to retrain the classifier head while keeping the first-stage representation fixed. Interestingly, as shown in Fig. 1, we find a substantial gap between this ideal performance and the baseline network, which indicates that the first-stage learning with unbalanced data provides a good representation, but there is a large room for improvement in the second stage due to the *biased decision boundary* (See Sec. 3.1 for details).

Based on those findings, we propose a simple and yet effective two-stage learning scheme for long-tail visual recognition problems. Our approach focuses on improving the second-stage training of the classifier after learning a feature representation in a standard manner. To this end, we develop a *unified* distribution alignment strategy to calibrate the classifier output via matching it to a reference distribution of classes that favors the balanced prediction. Such an alignment strategy enables us to exploit the class prior and data input in a principled manner for learning class decision boundary, which eliminates the needs for tedious hyper-parameter tuning and can be easily applied to various visual recognition tasks.

Specifically, we develop a light-weight distribution alignment module for calibrating classification scores, which consists of two main components. In the first component, we introduce an adaptive calibration function that equips the class scores with an input-dependent, learnable magnitude and margin. This allows us to achieve a flexible and confidence-aware distribution alignment for each data point. Our second component explicitly incorporates a balanced class prior by employing a generalized re-weight design for the reference class distribution, which provides a unified strategy to cope with diverse scenarios of label imbalance in different visual recognition tasks.

We extensively validate our model on four typical visual recognition tasks, including image classification on three benchmarks (ImageNet-LT [26], iNaturalist [36] and Places365-LT [26]), semantic segmentation on ADE20k dataset [49], object detection and instance segmentation on LVIS dataset [14]. The empirical results and ablative study show our method consistently outperforms the state-of-the-art approaches on all the benchmarks. To summarize, the main contributions of our works are three-folds:

- We conduct an empirical study to investigate the performance bottleneck of long-tail recognition and reveal a critical gap caused by biased decision boundary.
- We develop a simple and effective distribution alignment strategy with a generalized re-weight method, which can be easily optimized for various long-tail recognition tasks without whistles and bells.
- Our models outperform previous work with a large margin and achieve state-of-the-art performance on long-tail image classification, semantic segmentation, object detection, and instance segmentation.

2. Related Works

One-stage Imbalance Learning To alleviate the adverse effect of the long-tail class distribution in visual recognition, prior work have extensively studied the one-stage methods, which either leverage the re-balancing ideas or explore knowledge transfer from head categories. The basic idea of resample-based methods is to over-sample the minority categories [4, 15] or to under-sample the frequent categories in the training process [10, 2]. Class-aware sampling [33] proposes to choose samples of each category with equal probabilities, which is widely used in vision tasks [26, 11]. Repeat factor sampling [27] is a smoothed sampling method conducting repeated sampling for tail categories, which demonstrates its efficacy in instance segmentation [14]. In addition, [37] proposes to increase the sampling rate for categories with low performance after each training epoch and balances the feature learning for under-privileged categories.

An alternative strategy is to re-weight the loss function in training. Class-level methods typically re-weight the standard loss with category-specific coefficients correlated with the sample distributions [18, 9, 3, 21, 20, 34]. Sample-level methods [24, 30] try to introduce a more fine-grained control of loss for imbalanced learning. Other work aim to enhance the representation or classifier head of tail categories by transferring knowledge from the head classes [41, 40, 26, 46, 8, 43, 42]. Nevertheless, these methods require designing a task specific network module or structure, which is usually non-trivial to be generalized to different vision tasks.

Two-stage Imbalance Learning More recent efforts aims to improve the long-tail prediction by decoupling the learning of representation and classifier head. Decouple [19] proposes an instance-balanced sampling scheme, which generates more generalizable representations and achieves strong performance after properly re-balancing the classifier heads. The similar idea is adopted in [38, 39, 23], which develop effective strategies for long-tail object detection tasks. [28, 35] improve the two-stage ideas by introducing a post-process to adjust the prediction score. However, such a two-stage strategy typically relies on heuristic design in order to adjust the decision boundary of initially learned classifiers and requires tedious hyper-parameter tuning in practice.

Visual Recognition Tasks Visual recognition community has witnessed significant progress with deep convolutional networks in recent years. In this study, we focus on four types of visual tasks, including image classification, object detection, semantic and instance segmentation, which have been actively studied in a large amount of prior work. For object detection, we consider the typical deep network architecture used in the R-CNN series method [13, 12, 31], which detects objects based on the region proposals. For instance segmentation, we take the Mask R-CNN [16] as our example, which extends the Faster R-CNN[31] by adding a branch for predicting the object masks in parallel with the existing branch for bounding box recognition. For the pixel-wise task, semantic segmentation, we use the FCN-based methods [32] and the widely-adopted encoder-decoder structures [7, 5, 6]. Despite those specific choices, we note that our strategy can be easily extended to other types of deep network methods for those visual recognition tasks.

3. Our Approach

Our goal is to address the problem of large-scale long-tail visual recognition, which typically has a large number of classes and severe class imbalance in its training data. To this end, we adopt a two-stage learning framework that first learns a feature representation and a classifier head from the unbalanced data, followed by a calibration stage that adjusts the classification scores. Inspired by our ablative study on existing two-stage methods, we propose a principled calibration method that aligns the model prediction with a reference class distribution favoring the balanced evaluation metrics. Our distribution alignment strategy is simple and yet effective, enabling us to tackle different types of large-scale long-tail visual recognition tasks in a unified framework.

Below we start with a brief introduction to the long-tail classification and an empirical study of two-stage methods in Sec.3.1. We then describe our proposed distribution alignment strategy in Sec.3.2. Finally, we present a comparison with previous methods from the distribution match perspective in Sec.3.3.

3.1. Problem Setting and Empirical Study

We now introduce the problem setting of long-tail classification and review the two-stage learning framework for deep networks. Subsequently, we perform an empirical ablative study on a large-scale image classification task, which motivates our proposed approach.

Problem Definition The task of long-tail recognition aims to learn a classification model from a training dataset with long-tail class distribution. Formally, we denote the input as $\mathbf{I}$, and the target label space as $\mathcal{C} = \{c_1, \cdots, c_K\}$, where K is the number of classes. The classification model $\mathcal{M}$ defines a mapping from the input to the label space: $y = \mathcal{M}(\mathbf{I}; \Theta)$, where $y \in \mathcal{C}$ and Θ are its parameters. Our goal is to learn the model parameter from an imbalanced training dataset $\mathcal{D}_{tr}$ so that $\mathcal{M}$ achieves optimal performance on an evaluation dataset $\mathcal{D}_{eval}$ with respect to certain balanced metrics (*e.g.*, mean accuracy).

In the two-stage framework, we typically consider a deep network model $\mathcal{M}$ with two main components: a feature extractor network $f(\cdot)$ and a classifier head $h(\cdot)$. The feature extractor f first extracts an input representation $\mathbf{x}$, which is then fed into the classifier head h to compute class prediction scores $\mathbf{z}$ as follows:

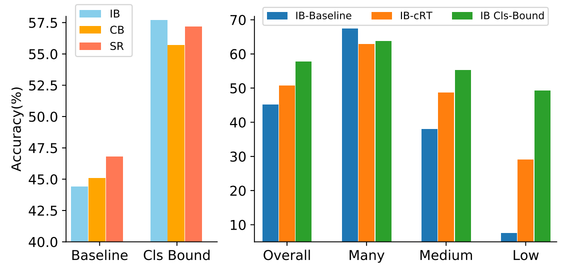

Figure 2: **Empirical analysis of the performance bottleneck.** *Left*: Baseline vs. ideal performance for representations learned with different sampling strategy. *Right*: Comparison of prior arts and ideal performance for the classifier head calibration. **Cls-Bound**: ideal performance bound. **IB**: instance-balanced sampling. **CB**: class-balanced sampling. **SR**: square-root sampling.

$$\mathbf{x} = f(\mathbf{I}, \theta_f) \in \mathbb{R}^d, \quad \mathbf{z} = h(\mathbf{x}, \theta_h) \in \mathbb{R}^K \tag{1}$$

where θ_f and θ_h are the parameter of $f(\cdot)$ and $h(\cdot)$, respectively. Here $\mathbf{z} = \{z_1, \cdots, z_K\}$ indicate the class prediction scores for K classes and the model predicts the class label by taking $y = \arg\max(\mathbf{z})$.

In this work, we instantiate the classifier head h as a linear classifier or a cosine similarity classifier as follows:

$$\text{Linear}: \quad z_j = \mathbf{w}_j^\intercal \mathbf{x} \tag{2}$$

$$\text{Cosine Similarity}: \quad z_j = s \cdot \frac{\mathbf{w}_j^\intercal \mathbf{x}}{||\mathbf{w}_j|| ||\mathbf{x}||} \tag{3}$$

where $\mathbf{w}_j \in \mathbb{R}^d$ is the parameter of j-th class and the s is a scale factor as in [29]. We note that the above formulation can be instantiated for multiple visual recognition tasks by changing the input $\mathbf{I}$: e.g., an image for image classification, an image with a pixel location for semantic segmentation, or an image with a bounding box for object detection.

Empirical Analysis on Performance Bound The two-stage learning method tackles the long-tail classification by decoupling the representation and the classifier head learning [19]. Specifically, it first learns the feature extractor f and classifier head h jointly, and then with the representation fixed, re-learns the classifier head with a class balancing strategy. While such design achieves certain success, an interesting question to ask is *which model component(s) impose a bottleneck on its balanced performance*. In the following, we attempt to address the question by exploiting the full set of the ImageNet dataset. Particularly, we follow the decoupling idea to conduct a series of ablative studies on two model components under an 'ideal' balanced setting.

We first investigate whether the feature representation learned on the imbalanced dataset is restrictive for the balanced performance. To this end, we start from learning the feature extractor on the imbalanced ImageNet-LT training set with several re-balancing strategies (*e.g.* instance-balanced, class-balanced, or square-root sampling). We then keep the representation fixed and re-train the classifier head with the ideal balanced ImageNet train set (excluding ImageNet-LT val set). Our results are shown in the left panel of Fig. 2, which indicate that *the first stage produces a strong feature representation that can potentially lead to large performance gain and the instance-based sampling achieves better overall results* (cf. [19]).

Moreover, we conduct an empirical study on the effectiveness of the recent decoupling method (*e.g.* cRT [19]) compared with the above 'ideal' classifier head learning. The right panel of Fig. 2 shows that there remains *a large performance gap* between the existing methods and the upper-bound. Those empirical results indicate that *the biased decision boundary in the feature space seems to be the performance bottleneck of the existing long-tail methods*. Consequently, a better strategy to address this problem would further improve the two-stage learning for the long-tail classification.

3.2. Distribution Alignment

To tackle the aforementioned issue, we now introduce a *unified* distribution alignment strategy to calibrate the classifier output via matching it to a reference distribution of classes that favors the balanced prediction. In this work, we adopt a two-stage learning scheme for all visual recognition tasks, which consists of a joint learning stage and a distribution calibration stage as follows.

1) Joint Learning Stage. The feature extractor $f(\cdot)$ and original classifier head (denoted as $h_o(\cdot)$ for clarity) are jointly learned on imbalanced $\mathcal{D}_{tr}$ with instance-balanced strategy in the first stage, where the original $h_o(\cdot)$ is severely biased due to the imbalanced data distribution.

2) Distribution Calibration Stage. For the second stage, the parameters of $f(\cdot)$ are frozen and we only focus on the classifier head to adjust the decision boundary. To this end, we introduce an *adaptive calibration function* (in Sec. 3.2.1) and a *distribution alignment strategy with generalized re-weighting* (in Sec. 3.2.2) to calibrate the class scores.

3.2.1 Adaptive Calibration Function

To learn the classifier head $h(\cdot)$ in the second stage, we propose an adaptive calibration strategy that fuses the original classifier head $h_o(\cdot)$ (parameters of $h_o(\cdot)$ are frozen) and a learned class prior in an input-dependent manner. As shown below, unlike previous work (*e.g.* cRT[19]), our design does not require a re-training of the classifier head from scratch and has much fewer free parameters. This enables us to reduce the adverse impact from the limited training data of the tail categories. Moreover, we introduce a flexible fusion mechanism capable of controlling the magnitude of calibration based on input features.

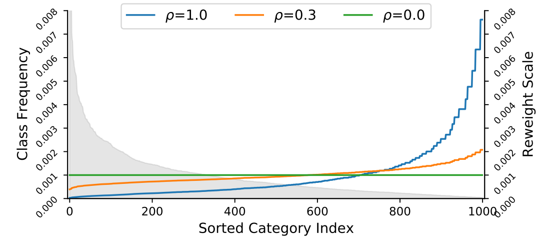

Figure 3: The category frequency r_i is plotted in gray color and the right axis denotes the re-weight coefficient at different scale ρ on ImageNet-LT dataset.

Specifically, denote the class scores from $h_o(\cdot)$ as $\mathbf{z}^o = [z_1^o, \cdots, z_K^o]$, we first introduce a class-specific linear transform to adjust the score as follows:

$$s_j = \alpha_j \cdot z_j^o + \beta_j, \quad \forall j \in \mathcal{C} \tag{4}$$

where α_j and β_j are the calibration parameters for each class, which will be learned from data. As mentioned above, we then define a confidence score function $\sigma(\mathbf{x})$ to adaptively combine the original and the transformed class scores:

$$\hat{z}_j = \sigma(\mathbf{x}) \cdot s_j + (1 - \sigma(\mathbf{x})) \cdot z_j^o \tag{5}$$
$$= (1 + \sigma(\mathbf{x})\alpha_j) \cdot z_j^o + \sigma(\mathbf{x}) \cdot \beta_j \tag{6}$$

where the confidence score has a form of $g(\mathbf{v}^\intercal \mathbf{x})$, which is implemented as a linear layer followed by a non-linear activation function (*e.g.*, sigmoid function) for all input $\mathbf{x}$. The confidence $\sigma(\mathbf{x})$ controls how much calibration is needed for a specific input $\mathbf{x}$. Given the calibrated class scores, we finally define a prediction distribution for our model with the Softmax function:

$$p_m(y = j|\mathbf{x}) = \frac{\exp(\hat{z}_j)}{\sum_{k=1}^{C} \exp(\hat{z}_k)}. \tag{7}$$

3.2.2 Alignment with Generalized Re-weighting

Given a train dataset $\mathcal{D}_{tr} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$, we introduce a calibration strategy based on distribution alignment between our model prediction $p_m(\cdot)$ and a reference distribution of classes that favors the balanced prediction.

Formally, denote the reference distribution as $p_r(y|\mathbf{x})$, we aim to minimize the expected KL-divergence between

Method	Align Type	Top-1 Accuracy@R-50				Top-1 Accuracy@X-50			
		Average	Many	Medium	Few	Average	Many	Medium	Few
Baseline[19]	-	41.6	64.0	33.8	5.8	44.4	65.9	37.5	7.7
Baseline*	-	48.4	68.4	41.7	15.2	49.2	68.9	42.8	15.6
NCM[19]	Hand-Craft	44.3	53.1	42.3	26.5	47.3	56.6	45.3	28.1
τ-Norm[19]		46.7	56.6	44.2	27.4	49.4	59.1	46.9	30.7
Logit Adjust(post)[28]		50.4	-	-	-	-	-	-	-
Deconfound*[35]		-	-	-	-	51.8	62.7	48.8	31.6
cRT[19]	Learnable	47.3	58.8	44.0	26.1	49.6	61.8	46.2	27.4
cRT*[19]		-	-	-	-	49.7	60.4	46.8	29.3
LWS[19]		47.7	57.1	45.2	29.3	49.9	60.2	47.2	30.3
DisAlign		51.3	59.9	49.9	31.8	52.6	61.5	50.7	33.1
DisAlign*		**52.9**	61.3	52.2	31.4	**53.4**	62.7	52.1	31.4

Table 1: **Quantitative results on ImageNet-LT.** $*$ denotes the model uses cosine classifier. **R-50** and **X-50** means the ResNet-50 and ResNeXt-50, respectively.

$p_r(y|\mathbf{x})$ and the model prediction $p_m(y|\mathbf{x})$ as follows:

$$\mathcal{L} = \mathbb{E}_{\mathcal{D}_{tr}}\left[\mathcal{KL}(p_r(y|\mathbf{x})||p_m(y|\mathbf{x}))\right] \tag{8}$$

$$\approx -\frac{1}{N}\sum_{i=1}^{N}\left[\sum_{y\in\mathcal{C}} p_r(y|\mathbf{x}_i)\log(p_m(y|\mathbf{x}_i))\right] + C \tag{9}$$

where the expectation is approximated by an empirical average on $\mathcal{D}_{tr}$ and C is a constant.

In this work, we adopt a re-weighting approach [9] and introduce a generalized re-weight strategy for the alignment in order to exploit the class prior. Formally, we represent the reference distribution as a weighted empirical distribution on the training set,

$$p_r(y=c|\mathbf{x}_i) = w_c \cdot \delta_c(y_i), \quad \forall c \in \mathcal{C} \tag{10}$$

where w_c is the class weight, and $\delta_c(y_i)$ is the Kronecker delta function(equals 1 if $y_i = c$, otherwise equals 0). We then define the reference weight based on the empirical class frequencies $\mathbf{r} = [r_1, \cdots, r_K]$ on the training set:

$$w_c = \frac{(1/r_c)^\rho}{\sum_{k=1}^{K}(1/r_k)^\rho}, \quad \forall c \in \mathcal{C} \tag{11}$$

where ρ is a scale hyper-parameter to provide more flexibility in encoding class prior. Note that our scheme reduces to the instance-balance re-weight method with $\rho = 0$, and to the class-balanced re-weight method with $\rho = 1$. We illustrate the curve of re-weight coefficients based on ImageNet-LT dataset in Fig. 3.

3.3. Connection with Recent Work

Below we discuss the connections between our proposed distribution alignment strategy and recent two-stage methods. Detailed comparison is reported in Tab. 2. Notably, Logit Adjustment[28] and Deconfound[35] introduce a hand-craft

Method	Align Method			
	Type	Balance	Magnitude	Margin
Joint	-	-	-	-
LWS[19]	L	CB-RS	α_j	0
τ-Normalized[19]	H	CB-RS	$1/\|\|\mathbf{w}_j\|\|^\tau$	0
Logit Adjust[28]	H	-	1.0	$-\lambda\log(r_j)$
Deconfound*[35]	H	-	1.0	$-\lambda d(\mathbf{x},\mathbf{e})\mathbf{w}_j^\intercal\mathbf{e}$
DisAlign	L	G-RW	$1+\sigma(\mathbf{x})\alpha_j$	$\sigma(\mathbf{x})\beta_j$
DisAlign*	L	G-RW	$1+\sigma(\mathbf{x})\alpha_j$	$\sigma(\mathbf{x})\beta_j$

Table 2: **Comparison with related methods.** $*$ denotes cosine classifier, **L**: learnable, **H**:hand-craft, **CB-RS**: class-balanced resampling, **G-RW**: generalized re-weight, r_j: class frequency for the j-th class, λ: hypper-parameter, $\mathbf{e}$: mean feature of training data, $d(\cdot)$: cosine distance.

margin to adjust the distribution while keep the magnitude as 1.0, and incorporate the class prior directly in r_i or $\mathbf{w}_i$ without re-training. LWS[19] and τ-normalized[19] try to achieve a similar goal by learning a magnitude scale and discarding the margin adjustment.

All these methods can be considered as the special cases of our DisAlign approach, which provides a unified and simple form to model the distribution mismatch in a learnable way. Moreover, the resample based strategy is not easy to be applied for the instance-level (object detection/instance segmentation) or pixel-level (semantic segmentation) tasks, our generalized re-weight provides an alternative solution to incorporate the class prior in a simple and effective manner. Experimental results in Sec. 4 also demonstrate the strength of our method compared with the aforementioned works.

4. Experiments

In this section, we conduct a series of experiments to validate the effectiveness of our method. Below we present our experimental analysis and ablation study on the image

Method	ResNet-50		ResNet-152	
	90 E	**200 E**	**90 E**	**200 E**
LDAM[3]	68.0	-	-	-
Baseline	61.7	65.8	65.0	69.0
Baseline*	64.8	66.2	67.3	69.0
cRT[19]	65.2	68.2	68.5	71.2
τ-norm[19]	65.6	69.3	68.8	72.5
LWS[19]	65.9	69.5	69.1	72.1
BBN[47]	66.3	69.6	-	-
DisAlign	67.8	**70.6**	71.3	**74.1**
DisAlign*	**69.5**	70.2	**71.7**	72.8

Table 3: **Average accuracy on iNaturalist-2018.** $*$ denotes the cosine classifier.

Method	ResNet-152			
	Average	**Many**	**Medium**	**Few**
Focal Loss[26]	34.6	41.1	34.8	22.4
Range Loss[26]	35.1	41.1	35.4	23.2
OLTR[26]	35.9	44.7	37.0	25.3
Feature Aug[8]	36.4	42.8	37.5	22.7
Baseline	30.2	**45.7**	27.3	8.2
NCM	36.4	40.4	37.1	27.3
cRT	36.7	42.0	37.6	24.9
LWS	37.6	40.6	39.1	28.6
τ-norm	37.9	37.8	40.7	**31.8**
DisAlign	**39.3**	40.4	**42.4**	30.1

Table 4: **Results on Place365-LT with ResNet-152.**

classification task in Sec. 4.1, followed by our results on semantic segmentation task in Sec. 4.2. In addition, we further evaluate our methods on object detection and instance segmentation tasks in Sec. 4.3.

4.1. Image Classification

Experimental Details To demonstrate our methods, we conduct experiments on three large-scale long-tail datasets, including ImageNet-LT [26], iNaturalist 2018 [36], and Places-LT [26]. We follow the experimental setting and implementation of [19] [1]. For the ImageNet-LT dataset, we report performance with ResNet/ResNeXt-{50,101,152} as backbone, and mainly use ResNet-50 for ablation study. For iNaturalist 2018 and Places-LT, our comparisons are performed under the settings of ResNet-{50,101,152}.

Comparison with previous methods 1) ImageNet-LT. We present the quantitative results for ImageNet-LT in Tab. 1. Our approach achieves **52.9%** in per-class average accuracy based on ResNet-50 backbone and **53.4%** based on ResNeXt-50, which outperform the state-of-the-art methods by a significant margin of **2.5%** and **1.6%**, respectively. **2) iNaturalist.** In Tab. 3, our method DisAlign with cosine classifier achieves **69.5%** per-class average accuracy using ResNet-50 backbone and 90 epochs of training, surpassing the prior art LDAM by a large margin at **1.5%**. It also shows that our performance can be further improved with larger backbone and/or more training epochs. **3) Places-LT.** In Tab. 4, we show the experimental results under the same setting as [19] on Places-LT. Our method achieves **39.3%** per-class average accuracy based on ResNet-152, with a notable performance gain at **1.4%** over the prior methods. We also report the detailed performance of these three datasets with ResNet-{50,101,152} in the supplementary materials.

[1]Detailed configuration and results are provided in the supplementary materials.

GR	MT	MG	Average	Many	Medium	Few
✗	✗	✗	41.6	64.0	33.8	5.8
✓	✓	✗	50.1	60.4	48.0	28.8
✓	✗	✓	49.9	63.9	46.9	21.2
✓	✓	✓	51.3	59.9	49.9	31.8

Table 5: **Ablation study of DisAlign. GR** means the generalized reweight strategy. **MT** means the learnable magnitude parameter $(1+\sigma(\mathbf{x})\alpha)$ and **MG** is the learnable margin parameter $\sigma(\mathbf{x})\beta$.

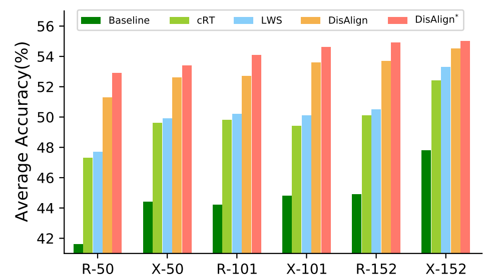

Figure 4: **Performance of DisAlign with different backbone on ImageNet-LT.** Detailed results will be reported in supplementary material.

Ablation Study 1) Different Backbone: We validate our method on different types of backbone networks, ranging from ResNet-{50,101,152} to ResNeXt-{50, 101, 152}, reported in Fig. 4. Our method achieves **54.9%** with ResNet-152, and **55.0%** with ResNeXt-152. It's worth noting that even when adopting stronger backbones, the gain of DisAlign compared to the state-of-the-art methods is still significant. This demonstrates that our DisAlign is complementary to the capacity of backbone networks. **2) Model Components:** We conduct a series of ablation studies to evaluate the importance of each component used in our DisAlign method. Tab. 5 summarizes the results of our ablation experiments, in which we compare our full model with several partial model settings. From the table, we find the learnable magnitude has a significant improvement compared with

Framework	Method	B	Mean IoU(%)				Mean Accuracy(%)			
			Average	**Head**	**Body**	**Tail**	**Average**	**Head**	**Body**	**Tail**
FCN[32]	Baseline	R-50	38.1	64.6	40.0	29.6	46.3	78.6	49.3	35.4
	DisAlign		40.1(+2.0)	65.0(+0.4)	42.8(+2.8)	31.3(+1.7)	51.4(+5.1)	78.6(+0.0)	56.1(+6.8)	40.6(+5.2)
	Baseline	R-101	41.4	67.0	43.3	33.2	50.2	80.6	52.9	40.1
	DisAlign		43.7(+2.3)	67.4(+0.4)	46.1(+2.8)	35.7(+2.5)	55.9(+5.7)	80.6(+0.0)	59.7(+6.8)	46.4(+6.3)
	Baseline	S-101	46.2	67.6	48.0	39.1	57.3	79.4	61.7	48.2
	DisAlign		46.9(+0.7)	67.7(+0.1)	48.2(+0.2)	40.3(+1.2)	60.1(+2.8)	79.7(+0.3)	64.2(+2.5)	51.9(+3.7)
DeepLabV3$^+$[7]	Baseline	R-50	44.9	67.7	48.3	36.4	55.0	80.1	60.8	44.1
	DisAlign		45.7(+0.8)	67.7(+0.0)	48.6(+0.3)	37.8(+1.4)	57.3(+2.3)	80.8(+0.7)	63.0(+2.2)	46.9(+2.8)
	Baseline	R-101	46.4	68.7	49.0	38.4	56.7	80.9	61.5	46.7
	DisAlign		47.1(+0.7)	68.7(+0.0)	49.4(+0.4)	39.6(+1.2)	59.5(+2.8)	81.4(+0.5)	64.2(+2.7)	50.3(+3.6)
	Baseline	S-101	47.3	69.0	49.7	39.7	58.1	80.8	63.4	48.2
	DisAlign		47.8(+0.5)	68.9(-0.1)	49.8(+0.1)	40.7(+1.0)	60.1(+2.0)	81.0(+0.2)	65.5(+2.1)	52.0(+3.8)

Table 6: **Performance of semantic segmentation on ADE-20K:** All baseline models are trained with an image size of 512×512 and 160K iteration in total. **B** is backbone network(R-50, R-101, S-101 denote ResNet-50, ResNet-101 and ResNeSt-101, respectively).

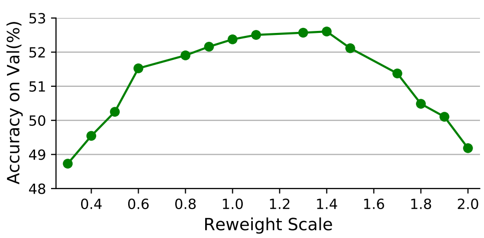

Figure 5: **Effects of the different generalized re-weight scale.** Performance is reported on ImageNet-LT val split.

baseline and the learnable margin also achieves competitive results at 49.9%, which demonstrate the effectiveness of individual modules in our design. **3) Generalized Re-weight Scale** We also investigate the influence of the generalized re-weight scale on the validation set of ImageNet-LT and plot the accuracy-scale curve in Fig. 5. It is evident that adjusting generalized reweight is able to achieve significant performance improvement. Moreover, we find the setting of $\rho > 1$ is able to outperform the class-balanced re-weight ($\rho = 1$), which indicates that the generalized re-weight is more effective in coping with long-tail distributions.

4.2. Semantic Semgnetaion on ADE20k Dataset

To further validate our method, we apply DisAlign strategy to segmentation networks and report our performance on the semantic segmentation benchmark, ADE20k [49].

Dataset and Implementation Details Follow a similar protocol as in image classification, we divide the 150 categories into 3 subsets according to the percentage of pixels in every category over the entire dataset. Specifically, we define three disjoint subsets as follows: *head classes* (each with more than 1.0% of total pixels), *body classes* (each with a percentage ranging from 0.1% to 1% of total pixels) and *tail classes* (each with less than 0.1% of total pixels). [2]

B	Method	Mask R-CNN		Cascade R-CNN	
		AP$_{bbox}$	**AP**$_{mask}$	**AP**$_{bbox}$	**AP**$_{mask}$
R-50	Baseline	20.8	21.2	25.2	23.0
	DisAlign	23.9(+3.1)	24.2(+3.0)	28.7(+3.5)	26.1(+3.1)
	Baseline*	22.8	23.8	28.8	26.2
	DisAlign*	25.6(+2.8)	26.3(+2.5)	32.2(+3.4)	29.4(+3.2)
R-101	Baseline	22.2	22.6	26.1	24.0
	DisAlign	25.6(+3.4)	25.8(+3.2)	29.7(+3.6)	27.3(+3.3)
	Baseline*	24.5	25.1	30.4	28.1
	DisAlign*	27.5(+3.0)	28.2(+3.1)	33.7(+3.3)	30.9(+2.8)
X-101	Baseline	24.5	25.0	28.4	26.1
	DisAlign	26.8(+2.3)	27.4(+2.4)	31.3(+2.9)	28.7(+2.6)
	Baseline*	26.9	27.7	32.6	29.8
	DisAlign*	29.5(+2.6)	30.0(+2.3)	34.7(+2.1)	31.8(+2.0)

Table 7: **Results on LVIS v0.5 dataset with different backbones and different architectures.** The results are reported based on the Detectron2[44, 50] framework. We refer the reader to the supplementary material for the detailed comparison with the state of art.

Quantitative Results We evaluate our method using two widely-adopted segmentation models (FCN [32] and DeepLabV3$^+$ [7]) based on different backbone networks, ranging from ResNet-50, ResNet-101 to the latest ResNeSt-101, and report the performance in Tab. 6. Our method achieves **2.0** and **2.3** improvement in mIoU using FCN-8s with ResNet-50 and ResNet-101, respectively. The performance on the body and tail are improved significantly. Moreover, our method outperforms the baseline with large margin at **5.7** in mean accuracy with ResNet-101 backbone. Even with a stronger backbone: **ResNeSt-101** [45], our method also achieves **0.7** mIoU and **2.8** improvement in mean accu-

[2]The complete list of the split is reported in supplementary material.

Pre-Train	Method	BBox AP				Mask AP			
		AP_{bbox}	AP^r_{bbox}	AP^c_{bbox}	AP^f_{bbox}	AP_{mask}	AP^r_{mask}	AP^c_{mask}	AP^f_{mask}
ImageNet	Baseline	20.8	3.3	19.5	29.4	21.2	3.7	21.6	28.4
	Baseline*	22.8	10.2	21.1	30.1	23.8	11.5	23.7	28.9
	Focal Loss[24]	21.9	-	-	-	21.0	9.3	21.0	25.8
	SimCal[38]	22.6	13.7	20.6	28.7	23.4	16.4	22.5	27.2
	LST[17]	22.6	-	-	-	23.0	-	-	-
	RFS[14]	23.6	12.8	22.3	29.4	24.3	14.6	24.0	28.5
	EQL[34]	23.3	-	-	-	22.8	11.3	24.7	25.1
	DisAlign	23.9	7.5	25.0	29.1	24.3	8.5	26.3	28.1
	DisAlign*	**25.6**	13.7	25.6	30.5	**26.3**	14.9	27.6	29.2
COCO	Baseline	22.8	2.6	21.8	32.0	23.9	2.8	23.4	30.5
	Baseline*	25.0	10.2	23.9	32.3	25.3	11.0	25.5	30.7
	GroupSoftmax[23]	25.8	15.0	25.5	30.4	26.3	18.0	26.9	28.7
	DisAlign	25.5	8.2	26.3	32.4	25.7	9.4	27.6	29.7
	DisAlign*	**27.6**	14.8	27.9	32.4	**27.9**	16.2	29.3	30.8

Table 8: **Comparison with the-state-of-art on LVIS with Mask-R-CNN-FPN(ResNet-50 backbone).** All results are evaluated on the LVIS v0.5 validation set with the score threshold at 0.0001. ($*$ denotes cosine classifier for bbox classification.)

Backbone	Method	BBox AP				Mask AP			
		AP_{bbox}	AP^r_{bbox}	AP^c_{bbox}	AP^f_{bbox}	AP_{mask}	AP^r_{mask}	AP^c_{mask}	AP^f_{mask}
ResNet-50	Baseline*	26.5	8.7	25.0	36.0	23.5	8.1	22.4	31.5
	DisAlign*	30.5	17.9	30.1	36.5	27.0	15.7	27.0	31.9
ResNet-101	De-confound[35]	25.8	-	-	-	23.5	5.2	22.7	32.3
	De-confound TDE[35]	30.0	-	-	-	27.1	16.0	26.9	32.1
	Baseline*	28.9	11.8	27.7	37.8	25.6	10.5	24.9	33.0
	DisAlign*	32.7	20.5	32.8	38.1	28.9	18.0	29.3	33.3
ResNeXt-101	Baseline*	30.7	14.2	29.3	39.6	27.3	13.0	26.4	34.6
	DisAlign*	33.7	21.4	33.1	39.7	29.7	18.4	29.7	34.7

Table 9: **Results on LVIS v1.0 dataset with Cascade R-CNN.** * denotes cosine classifier head.

racy, where the tail categories have a performance gain of **1.2** in mIoU and **3.7** in mean accuracy. We further validate our method using DeepLabV3$^+$, which is a more powerful semantic segmentation model. Our DisAlign improves the performance of DeepLabV3$^+$ by a margin of **0.5** based on ResNeSt-101, which achieves the new state-of-the-art (**47.8** in mIoU) on the ADE20k dataset.

4.3. Object Detection and Instance Segmentation

Experimental Configuration We conduct experiments on LVIS [14] dataset. For evaluation, we use a COCO-style average precision (AP) metric that averages over categories and different box/mask IoU threshold [25].

Quantitative Results and Ablation Study We first compare our method with recent work and report quantitative results in Tab. 8. We find our DisAlign with cosine classifier head achieves **25.6** in AP_{bbox}, and **26.3** in AP_{mask} when applied to the Mask R-CNN+FPN with the ImageNet pre-trained ResNet-50 backbone. Moreover, our strategy can be further improved to achieve **27.6** in AP_{bbox} and **27.9** in AP_{mask} based on the COCO pre-trained model. In both cases, our method is able to maintain the performance of the frequent (also called *head*) categories, and gain significant improvement on common (also called *body*) and rare (also called *tail*) categories. We also report performance with more power detection framework (*e.g.*Cascade R-CNN) and stronger backbones (*e.g.* ResNet-50/101, and ResNeXt-101) in Tab. 7 and Tab. 9. It is worth noting that even with the stronger backbones or frameworks, the performance gain of our DisAlign over the baseline is still significant.

5. Conclusion

In this paper, we have presented a unified two-stage learning strategy for the large-scale long-tail visual recognition tasks. To tackle the biased label prediction, we develop a confidence-aware distribution alignment method to calibrate initial classification predictions. In particular, we design a generalized re-weight scheme to leverage the category prior for the alignment process. Extensive experiments show that our method outperforms previous works with a large margin on a variety of visual recognition tasks(image classification, semantic segmentation, and object detection/segmentation).

References

[1] Samy Bengio. Sharing representations for long tail computer vision problems. In *Proceedings of the ACM on International Conference on Multimodal Interaction*, 2015. 1

[2] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 2018. 1, 2

[3] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *Advances in Neural Information Processing Systems(NeurIPS)*, 2019. 1, 2, 6

[4] 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. 2

[5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence(TPAMI)*, 2018. 3

[6] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint*, 2017. 3

[7] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision(ECCV)*, 2018. 3, 7

[8] Peng Chu, Xiao Bian, Shaopeng Liu, and Haibin Ling. Feature space augmentation for long-tailed data. In *Proceedings of the European Conference on Computer Vision(ECCV)*, 2020. 2, 6

[9] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2019. 1, 2, 5

[10] Chris Drummond, Robert C Holte, et al. C4. 5, class imbalance, and cost sensitivity: why under-sampling beats over-sampling. In *Workshop on Learning from Imbalanced Datasets, II*, 2003. 2

[11] Yuan Gao, Xingyuan Bu, Yang Hu, Hui Shen, Ti Bai, Xubin Li, and Shilei Wen. Solution for large-scale hierarchical object detection datasets with incomplete annotation and data imbalance. *arXiv preprint*, 2018. 2

[12] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision(ICCV)*, 2015. 2

[13] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2014. 2

[14] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2019. 1, 2, 8

[15] Hui Han, Wen-Yuan Wang, and Bing-Huan Mao. Borderline-smote: a new over-sampling method in imbalanced data sets learning. In *International Conference on Intelligent Computing*, 2005. 2

[16] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision(ICCV)*, 2017. 2

[17] Xinting Hu, Yi Jiang, Kaihua Tang, Jingyuan Chen, Chunyan Miao, and Hanwang Zhang. Learning to segment the tail. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2020. 8

[18] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2016. 2

[19] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *International Conference on Learning Representations(ICLR)*, 2020. 1, 2, 3, 4, 5, 6

[20] Salman Khan, Munawar Hayat, Syed Waqas Zamir, Jianbing Shen, and Ling Shao. Striking the right balance with uncertainty. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2019. 2

[21] Salman H Khan, Munawar Hayat, Mohammed Bennamoun, Ferdous A Sohel, and Roberto Togneri. Cost-sensitive learning of deep feature representations from imbalanced data. *IEEE Transactions on Neural Networks and Learning Systems(TNNLS)*, 2017. 2

[22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems(NeurIPS)*, 2012. 1

[23] Yu Li, Tao Wang, Bingyi Kang, Sheng Tang, Chunfeng Wang, Jintao Li, and Jiashi Feng. Overcoming classifier imbalance for long-tail object detection with balanced group softmax. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2020. 1, 2, 8

[24] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision(ICCV)*, 2017. 2, 8

[25] 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 *Proceedings of the European Conference on Computer Vision(ECCV)*, 2014. 8

[26] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2019. 1, 2, 6

[27] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European Conference on Computer Vision(ECCV)*, 2018. 2

[28] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. *arXiv preprint*, 2020. 1, 2, 5

[29] Hang Qi, Matthew Brown, and David G Lowe. Low-shot learning with imprinted weights. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2018. 3

[30] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning(ICML)*, 2018. 2

[31] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems(NeurIPS)*, 2015. 2, 3

[32] Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional networks for semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence(TPAMI)*, 2017. 3, 7

[33] Li Shen, Zhouchen Lin, and Qingming Huang. Relay backpropagation for effective learning of deep convolutional neural networks. In *Proceedings of the European Conference on Computer Vision(ECCV)*, 2016. 1, 2

[34] Jingru Tan, Changbao Wang, Buyu Li, Quanquan Li, Wanli Ouyang, Changqing Yin, and Junjie Yan. Equalization loss for long-tailed object recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2020. 2, 8

[35] Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. Long-tailed classification by keeping the good and removing the bad momentum causal effect. In *Advances in Neural Information Processing Systems(NeurIPS)*, 2020. 1, 2, 5, 8

[36] Grant Van Horn, Oisin 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2018. 1, 2, 6

[37] Hao Wang, Qilong Wang, Fan Yang, Weiqi Zhang, and Wangmeng Zuo. Data augmentation for object detection via progressive and selective instance-switching, 2019. 2

[38] Tao Wang, Yu Li, Bingyi Kang, Junnan Li, Junhao Liew, Sheng Tang, Steven Hoi, and Jiashi Feng. The devil is in classification: A simple framework for long-tail instance segmentation. In *Proceedings of the European Conference on Computer Vision(ECCV)*, 2020. 1, 2, 8

[39] Xin Wang, Thomas E. Huang, Trevor Darrell, Joseph E Gonzalez, and Fisher Yu. Frustratingly simple few-shot object detection. In *International Conference on Machine Learning(ICML)*, 2020. 2

[40] Yu-Xiong Wang, Ross Girshick, Martial Hebert, and Bharath Hariharan. Low-shot learning from imaginary data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2018. 2

[41] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. In *Advances in Neural Information Processing Systems(NeurIPS)*, 2017. 1, 2

[42] Jialian Wu, Chunluan Zhou, Qian Zhang, Ming Yang, and Junsong Yuan. Self-mimic learning for small-scale pedestrian detection. In *Proceedings of the 28th ACM International Conference on Multimedia(ACM MM)*, 2020. 2

[43] Tz-Ying Wu and Pedro Morgado. Solving long-tailed recognition with deep realistic taxonomic classiï¬ er. In *European Conference on Computer Vision(ECCV)*, 2020. 2

[44] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. `https://github.com/facebookresearch/detectron2`, 2019. 7

[45] Hang Zhang, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Zhi Zhang, Haibin Lin, Yue Sun, Tong He, Jonas Muller, R. Manmatha, Mu Li, and Alexander Smola. Resnest: Split-attention networks. *arXiv preprint*, 2020. 7

[46] Yaoyao Zhong, Weihong Deng, Mei Wang, Jiani Hu, Jianteng Peng, Xunqiang Tao, and Yaohai Huang. Unequal-training for deep face recognition with long-tailed noisy data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition(CVPR)*, 2019. 2

[47] Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2020. 6

[48] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence(TPAMI)*, 2017. 1

[49] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2017. 1, 2, 7

[50] Benjin Zhu*, Feng Wang*, Jianfeng Wang, Siwei Yang, Jianhu Chen, and Zeming Li. cvpods: All-in-one toolbox for computer vision research, 2020. 7

[51] Xiangxin Zhu, Dragomir Anguelov, and Deva Ramanan. Capturing long-tail distributions of object subcategories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR)*, 2014. 1