

Seesaw Loss for Long-Tailed Instance Segmentation

Jiaqi Wang¹ Wenwei Zhang² Yuhang Zang² Yuhang Cao¹ Jiangmiao Pang⁵ Tao Gong⁶
Kai Chen^{3,4} Ziwei Liu² Chen Change Loy² Dahua Lin¹

¹SenseTime-CUHK Joint Lab, The Chinese University of Hong Kong

²S-Lab, Nanyang Technological University ³SenseTime Research

⁴Shanghai AI Laboratory ⁵Zhejiang University ⁶University of Science and Technology of China

{wj017, cy020, dhlin}@ie.cuhk.edu.hk {wenwei001, zang0012, ziwei.liu, ccloy}@ntu.edu.sg

{pangjiangmiao, gongtao950513}@gmail.com chen kai@sensetime.com

Abstract

Instance segmentation has witnessed a remarkable progress on class-balanced benchmarks. However, they fail to perform as accurately in real-world scenarios, where the category distribution of objects naturally comes with a long tail. Instances of head classes dominate a long-tailed dataset and they serve as negative samples of tail categories. The overwhelming gradients of negative samples lead to a biased learning process for classifiers. Consequently, objects of tail categories are more likely to be misclassified as backgrounds or head categories. To tackle this problem, we propose Seesaw Loss to dynamically re-balance gradients of positive and negative samples for each category, with two complementary factors, i.e., mitigation factor and compensation factor. The mitigation factor reduces punishments to tail categories w.r.t. the ratio of cumulative training instances between different categories. Meanwhile, the compensation factor increases the penalty of misclassified instances to avoid false positives of tail categories. We conduct extensive experiments on Seesaw Loss with mainstream frameworks and different data sampling strategies. With a simple end-to-end training pipeline, Seesaw Loss obtains significant gains over Cross-Entropy Loss, and achieves state-of-the-art performance on LVIS dataset without bells and whistles. Code is available at <https://github.com/open-mmlab/mmdetection>.

1. Introduction

Deep learning-based object detection and instance segmentation approaches have achieved immense success on datasets with relatively balanced category distribution, e.g., COCO dataset [26]. However, the distribution of categories in the real world is long-tailed [28]. There are a few head classes containing abundant instances, while most other classes comprise relatively few instances.

On long-tailed datasets, existing instance segmentation

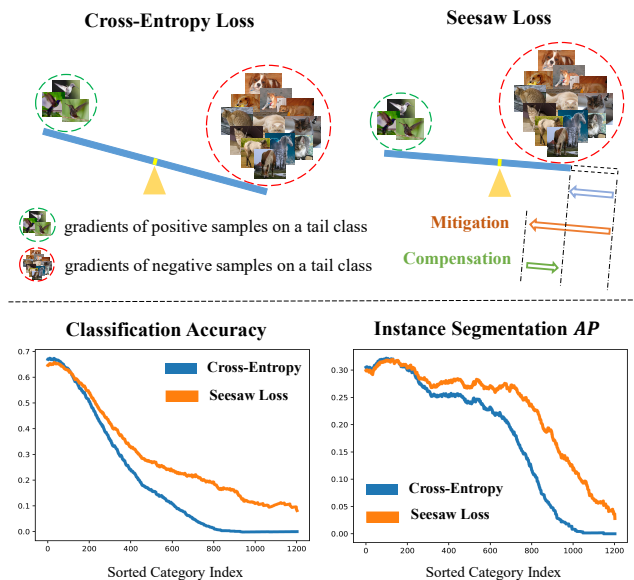


Figure 1: Seesaw Loss dynamically re-balances the gradients of positive and negative samples on a tail class with two complementary factors. It mitigates the overwhelming punishments on the tail class as well as compensates them to reduce the risk of inducing false positives. In Mask R-CNN [15], Seesaw Loss achieves remarkable higher classification accuracy of tail classes than Cross-Entropy Loss on LVIS [12] dataset. As a result, instance segmentation AP on tail classes is significantly improved, leading to better overall performance.

frameworks [1, 3, 15] fail to perform as accurately as on the datasets with balanced category distribution, exhibiting unsatisfactory performance on tail classes. Figure 1 shows the classification accuracy and instance segmentation performance of Mask R-CNN [15] on LVIS [12] dataset. The classifier in Mask R-CNN trained by Cross-Entropy Loss tends to misclassify tail categories as backgrounds or other confusing head classes, which leads to extremely low accuracy on tail classes.

The primary reason for this undesired phenomenon is that the instances from head classes are predominant in a long-tailed dataset. These instances contribute an overwhelmingly large quantity of negative samples for tail classes. Thus, the gradients of positive and negative samples on a tail class are heavily imbalanced, leading to a biased learning process for the classifier. One can imagine that gradients of positive and negative samples resemble two objects positioned on each end of a seesaw (see Fig. 1). To balance them, a viable solution is to shorten the arm of the heavier end in the seesaw, which is equivalent to scaling down the overwhelming gradients of negative samples on the tail class by a factor. Nevertheless, blindly reducing the gradients of negative samples increases the risk of inducing false positives of tail classes, since samples of other classes are less punished when they are misclassified as tail classes. Thus, a specialized mechanism is needed to compensate for the excessively reduced penalties on tail classes.

In this work, we propose **Seesaw Loss** that dynamically re-balances positive and negative gradients for each category with two complementary factors, *i.e.*, mitigation factor and compensation factor. According to the ratio between categories' cumulative sample numbers during training, the mitigation factor reduces the penalty to relatively rare classes. When a false positive sample of one category is observed, the compensation factor will increase the penalty to that category. The synergy of the two above factors enables Seesaw Loss to mitigate the overwhelming punishments to tail classes as well as compensate for the risk of misclassification caused by diminished penalties.

Seesaw Loss has three appealing properties. **1) Seesaw Loss is dynamic.** It explores the ratios of cumulative training sample numbers between different categories and instance-wise misclassification during training. This differs significantly to previous solutions that rely either on static group split [23] or loss reweighting with constant values [2, 6, 38]. **2) Seesaw Loss is self-calibrated.** The mitigation and the compensation factor synergize to relieve the overwhelming punishments on tail classes as well as avoid increasing false positives of tail categories. On the contrary, previous methods blindly reduce punishments on tail classes [38] or decrease the loss weights of head categories [6]. **3) Seesaw Loss is distribution-agnostic.** It does not rely on pre-computed datasets' distribution [2, 6, 28, 38], and it can operate well with any data sampler [12, 17]. By accumulating the number of samples in each class, Seesaw Loss gradually approximates the real data distribution during training to achieve more accurate balancing.

Through extensive experiments, we show consistent improvements of Seesaw Loss in different instance segmentation frameworks and data samplers. On the challenging LVIS [12] dataset, Seesaw Loss achieves significant

improvements of 6.0% AP and 2.1% AP upon Mask R-CNN [15] with random sampler and repeat factor sampler [12], respectively. Even if switching to the stronger Cascade Mask R-CNN [1], we still observe an impressive improvement of 6.4% AP and 2.3% AP with random sampler and repeat factor sampler. To show the versatility of Seesaw Loss, we integrate it into the long-tailed image classification task. Seesaw Loss significantly improves the classification accuracy by 6% on ImageNet-LT [28] dataset. Besides, we also explore the necessity of the decoupling training pipeline [20, 23] in Seesaw Loss. Experimental results demonstrate that Seesaw Loss provides a simpler and more effective solution to long-tailed instance segmentation without relying on complex training pipelines.

2. Related Work

Object Detection. Recent years have witnessed a remarkable improvement in object detection [7, 14, 44, 41, 42]. A leading paradigm in this area is the two-stage pipeline [10, 36], where the first stage generates a set of region proposals, and then the second stage classifies and refines the proposals. Unlike the two-stage approaches, the single-stage pipeline [25, 27, 34, 35] directly predicts bounding boxes. Classical single-stage approaches [25, 27] require densely populated anchors as a prior, while anchor-free methods [21, 22, 39, 45] manage to achieve similar or better performance without such prior. There are also attempts to apply cascade architecture [1, 8, 19, 31, 43] to refine the bounding boxes' predictions progressively.

Instance Segmentation. Instance segmentation is becoming popular in tandem with a surge in the interest in object detection. Early methods perform segmentation before object recognition [32, 33]. Via adding a mask prediction branch in the Faster R-CNN [36] architecture, Mask R-CNN [15] bridges the gap between object detection and instance segmentation. The idea is also adopted by [1, 3] in their cascading frameworks. More recent works [5, 47] introduce an even shorter pipeline by skipping the detection process and directly predicting mask for each instance. Seesaw Loss can easily cooperate with object detection and instance segmentation frameworks for the long-tailed datasets.

Long-Tailed Recognition. Long-tailed recognition tasks [28, 12, 13, 48] receive growing attention recently as the problems are closer to real-world applications. One representative solution to the problem is loss re-weighting [2, 6, 18]. Loss re-weighting methods adopt different re-weighting strategies [2, 6, 18, 25, 38] to adjust the loss of different classes based on each class's statistics [2, 6]. Other common approaches [13, 17, 30] re-balance the distribution of the instance numbers in each class, *e.g.*, repeat factor sampling [12] and class-balanced sampling [30], both are based on the sample numbers of classes. Different sampling

strategies can be adopted at different training stages to formulate a multi-stage training procedure [17, 20]. A recent work [20] proposes a decoupling training pipeline, which first trains a good representation network with natural sampling and then finetunes the classifier with class-balanced sampling. There are also attempts to modify the classifier to improve the performance on tail classes, *e.g.*, using different classifiers for different groups of classes [23], or use two classifiers trained with different data samplers [46].

3. Methodology

The classifier trained by the widely applied Cross-Entropy (CE) Loss (Sec. 3.1) is highly biased on long-tailed datasets, resulting in much lower accuracy of tail classes than head classes. The major reason is that gradients brought by positive samples are overwhelmed by gradients from negative samples on tail classes. Therefore, we propose Seesaw Loss to mitigate the overwhelming gradients of negative samples on tail classes as well as compensate the gradients of misclassified samples to avoid false positives (Sec. 3.2). We also explore some practical component designs to adopt Seesaw Loss in instance segmentation (Sec. 3.3).

3.1. Cross-Entropy Loss

We first revisit the most widely adopted Cross-Entropy (CE) Loss in existing frameworks [3, 15]. The formulation of CE Loss can be written as

$$L_{ce}(\mathbf{z}) = - \sum_{i=1}^C y_i \log(\sigma_i), \quad \text{with } \sigma_i = \frac{e^{z_i}}{\sum_{j=1}^C e^{z_j}}, \quad (1)$$

where $\mathbf{z} = [z_1, z_2, \dots, z_C]$ and $\boldsymbol{\sigma} = [\sigma_1, \sigma_2, \dots, \sigma_C]$ are the predicted logits and probabilities of the classifier, respectively. And $y_i \in \{0, 1\}, 1 \leq i \leq C$ is the one-hot ground truth label. Given a training sample of class i , the gradients on z_i and z_j are given by

$$\frac{\partial L_{ce}(\mathbf{z})}{\partial z_i} = \sigma_i - 1, \quad (2)$$

$$\frac{\partial L_{ce}(\mathbf{z})}{\partial z_j} = \sigma_j, \quad (3)$$

It shows that samples of class i punish the classifier of class j w.r.t. σ_j . In the case that the instance number of class i is enormously greater than that of class j , the classifier of class j will receive penalties in most samples and attains few positive signals during training. Thus the predicted probabilities of class j will be heavily suppressed, which results in a low classification accuracy of tail classes, as shown in Figure 1.

3.2. Seesaw Loss

To alleviate the above mentioned problem, one feasible solution is to decrease the gradients of negative samples in Eq. 3 imposed by head classes on a tail class. Therefore, we propose Seesaw Loss as

$$L_{seesaw}(\mathbf{z}) = - \sum_{i=1}^C y_i \log(\hat{\sigma}_i), \quad (4)$$

$$\text{with } \hat{\sigma}_i = \frac{e^{z_i}}{\sum_{j \neq i}^C \mathcal{S}_{ij} e^{z_j} + e^{z_i}}.$$

Then the gradient on z_j of negative class j in Eqn 3 becomes

$$\frac{\partial L_{seesaw}(\mathbf{z})}{\partial z_j} = \mathcal{S}_{ij} \frac{e^{z_j}}{e^{z_i}} \hat{\sigma}_i. \quad (5)$$

Here \mathcal{S}_{ij} works as a tunable balancing factor between different classes. By a careful design of \mathcal{S}_{ij} , Seesaw loss adjusts the punishments on class j from positive samples of class i . Seesaw loss determines \mathcal{S}_{ij} by a mitigation factor and a compensation factor, as

$$\mathcal{S}_{ij} = \mathcal{M}_{ij} \cdot \mathcal{C}_{ij}. \quad (6)$$

The mitigation factor \mathcal{M}_{ij} decreases the penalty on tail class j according to a ratio of instance numbers between tail class j and head class i . The compensation factor \mathcal{C}_{ij} increases the penalty on class j whenever an instance of class i is misclassified to class j .

Mitigation Factor. Seesaw Loss accumulates instance number N_i for each category i at each iteration in the whole training process. As shown in Fig. 2, given an instance with positive label i , for another category j , the mitigation factor adjusts the penalty for negative label j w.r.t. the ratio $\frac{N_j}{N_i}$

$$\mathcal{M}_{ij} = \begin{cases} 1, & \text{if } N_i \leq N_j \\ \left(\frac{N_j}{N_i}\right)^p, & \text{if } N_i > N_j \end{cases} \quad (7)$$

When category i is more frequent than category j , Seesaw Loss will reduce the penalty on category j , which is imposed by samples of category i , by a factor of $\left(\frac{N_j}{N_i}\right)^p$. Otherwise, Seesaw Loss will keep the penalty on negative classes to reduce misclassification. The exponent p is a hyper-parameter that adapts the magnitude of mitigation.

Note that Seesaw Loss accumulates the instance numbers during training, rather than get the statistics from the whole dataset ahead of time. This strategy brings two benefits. First, it can be applied when the distribution of the whole training set is unavailable, *e.g.*, training examples are obtained from a stream. Second, the training samples of each category can be affected by the adopted data sampler [12], and the online accumulation is robust to sampling

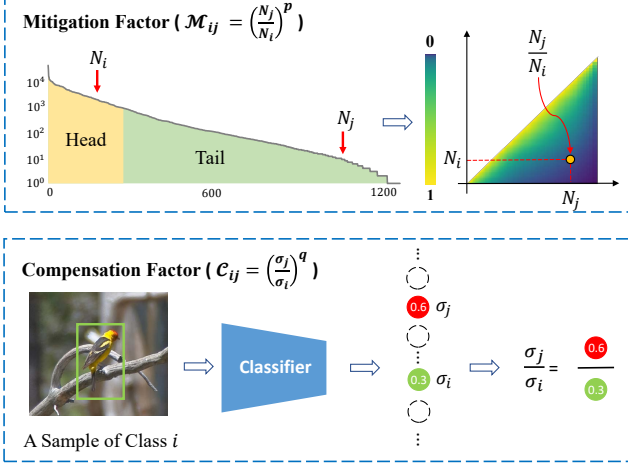


Figure 2: Seesaw Loss adjusts the punishments on tail classes with the mitigation factor \mathcal{M}_{ij} and the compensation factor \mathcal{C}_{ij} . The mitigation factor decreases the punishments w.r.t. the ratio of instance numbers between different categories. The compensation factor increases the penalty of misclassified instances w.r.t. the ratio of classification probabilities between the false positive and the ground-truth category.

methods. During training, the mitigation factor is uniformly initialized and smoothly updated to approximate the real data distribution.

Compensation Factor. The mitigation factor effectively balances the gradients of head and tail classes. Nevertheless, it may cause more false positives for tail classes due to less penalty. Moreover, the false positives cannot be eliminated by simply adjusting p in \mathcal{M}_{ij} , since it is applied to the whole category. We propose a compensation factor that focuses on misclassified samples instead of adjusting the whole category. As shown in Fig. 2, this factor compensates the diminished gradient when there is misclassification, *i.e.*, the predicted probability σ_j of negative label j is greater than σ_i . The compensation factor \mathcal{C}_{ij} is calculated as

$$\mathcal{C}_{ij} = \begin{cases} 1, & \text{if } \sigma_j \leq \sigma_i \\ \left(\frac{\sigma_j}{\sigma_i}\right)^q, & \text{if } \sigma_j > \sigma_i \end{cases} \quad (8)$$

For a training sample with positive label i , if the predicted probability of any negative class j is greater than class i , *i.e.*, $\sigma_j > \sigma_i$, the compensation factor increases the punishment on class j by a factor of $\left(\frac{\sigma_j}{\sigma_i}\right)^q$, where q is a hyper-parameter to control the scale. Otherwise, $\mathcal{C}_{ij} = 1$ and only the mitigation factor \mathcal{M}_{ij} is applied.

Normalized Linear Activation. The classifier in an object detector [36] usually predicts classification logits as $z = \mathcal{W}^T x + b$ on the dataset with balanced category distribution [26], where \mathcal{W} and b are the weights and bias of the linear layer and x is the input features. On long-tailed datasets, previous works [20, 23] find that the weight

norm of \mathcal{W}_i is highly related to the number of training instances in the corresponding category i . The more training samples of category i there are, the larger $\|\mathcal{W}_i\|$ will be. This phenomenon is also observed in the feature norm $\|x\|$. Therefore, we adopt a normalized linear activation (which [9, 28, 40] are related) to re-balance the scale of $\|\mathcal{W}_i\|$ and $\|x\|$ as $z = \tau \tilde{\mathcal{W}}^T \tilde{x} + b$, where $\tilde{\mathcal{W}}_i = \frac{\mathcal{W}_i}{\|\mathcal{W}_i\|_2}$, $i \in C$, $\tilde{x} = \frac{x}{\|x\|_2}$, and τ is a temperature factor. The normalized linear activation normalizes the weights \mathcal{W} and features x by their L_2 norm to reduce their scale variance for different categories. Thus, it effectively balances the distribution of predicted probabilities of different categories and improves the performance on a long-tailed dataset.

3.3. Model Design for Instance Segmentation

Objectness Branch. In contrast to image classification, the classifier in an object detector has two functionalities. It first determines if a bounding box is a foreground object then distinguishes which category the foreground instance belongs to. Previous practices [1, 15, 36] usually regard the background as an auxiliary category in the classifier. Given a dataset with C categories, the classifier in most detectors [15, 36] predicts logits of $C + 1$ classes. Although widely adopted, this design brings difficulty when adopting Seesaw Loss to balance long-tailed distribution. In general, most object candidates in a detector are backgrounds. Thus, all foreground categories are much rarer categories compared to the background category. Consequently, Seesaw Loss will significantly reduce punishments on all foreground categories. As a result, the classifier tends to misclassify more backgrounds as foregrounds and harms the performance. To tackle this problem, we decouple the two functionalities of the classifier in an object detector. Specifically, apart from the classifier with C classes, we adopt an extra objectness branch to distinguish the foregrounds and backgrounds. The objectness branch adopts the normalized linear activation to predict logits of two classes, *i.e.*, foreground and background, and is trained by cross-entropy loss. During inference, both the classification logit z_i^{class} of category $i \in C$ and logit of objectness z_i^{obj} are activated with a softmax function. The final detection probability σ_i^{det} for a bounding box of category i is $\sigma_i^{det} = \sigma_i^{class} \cdot \sigma_i^{obj}$.

Normalized Mask Prediction. Inspired by normalized linear activation, we further present a normalized mask prediction to alleviate the biased training process in mask head. In Mask R-CNN [15], a 1×1 convolution layer is applied in the end of the mask head, and the predicted logits are activated by a sigmoid function. We normalize the weights \mathcal{W} of the 1×1 convolution layer and the input features \mathcal{X} with L_2 normalization. Note that the spatial size of \mathcal{X} is $H \times W$, we denote the feature at (y, x) as $\mathcal{X}_{y,x}$. The formula of normalized mask prediction is $z = \tau \tilde{\mathcal{W}} * \tilde{\mathcal{X}} + b$, where

$\tilde{W}_i = \frac{W_i}{\|W_i\|_2}, i \in C, \tilde{X}_{y,x} = \frac{X_{y,x}}{\|X_{y,x}\|_2}, y \in H, x \in W$ and τ is a temperature factor.

4. Experiments

4.1. Experimental Settings

Datasets. We perform experiments on the challenging LVIS v1 dataset [12]. LVIS is a large vocabulary instance segmentation dataset containing 1203 categories with high-quality instance mask annotations. LVIS v1 provides a *train* split with 100k images, a *val* split with 19.8k images and a *test-dev* split with 19.8k images. According to the numbers of images that each category appears in the *train* split, the categories are divided into three groups: rare (1-10 images), common (11-100 images) and frequent (>100 images).

Evaluation metrics. The results of instance segmentation are evaluated with *AP* of mask prediction, which is averaged at different IoU thresholds (from 0.5 to 0.95) across categories. The *AP* for rare, common and frequent categories are denoted as AP_r , AP_c and AP_f . The *AP* for detection boxes is denoted as AP^{box} .

Implementation Details. We implement our method with mmdetection [4] and train Mask R-CNN [15], Cascade Mask R-CNN [1] using the 2x training schedule [4, 11]. The model is trained with batch size of 16 for 24 epochs. The learning rate is 0.02, and it will decrease by 0.1 after 16 and 22 epochs, respectively. ResNet-50 [16] with FPN [24] backbone is adopted if not further specified. Following the practice in mmdetection [4], we adopt multi-scale with horizontally flip augmentation during training. Specifically, we randomly resize the shorter edge of the image within {640, 672, 704, 736, 768, 800} pixels and keep the longer edge smaller than 1333 pixels without changing the aspect ratio. In inference, we adopt single-scale testing with image size of 1333×800 pixels and score thresholds of 10^{-3} without bells and whistles.

Apart from the standard random sampler that samples images in *train* split randomly, the repeat factor sampler (RFS) [12, 30] is also evaluated in experiments. RFS oversamples categories that appear in less than 0.1% of the total images and is effective to improve the overall *AP*. The ablation study is conducted with RFS if not further specified. We adopt Seesaw Loss in the box classification branch of Mask R-CNN [15] with hyper-parameter $p = 0.8$, $q = 2$, and $\tau = 20$. In Cascade Mask R-CNN [1], Seesaw Loss is adopted in box classification branches of all three stages with the same hyper-parameters as that in Mask R-CNN. We further evaluate the proposed Normalized Mask Prediction and integrate it into the mask head of Mask R-CNN and all the mask heads in Cascade Mask R-CNN. For simplicity, Normalized Mask Prediction adopts the same temperature, *i.e.* $\tau = 20$. We use the *train* split for training and report the performance on *val* split for ablation study. The perfor-

mance of our method is also reported on *test-dev* split.

4.2. Benchmark Results

To show the effectiveness of Seesaw Loss, we perform extensive experiments with different data samplers and instance segmentation frameworks. We adopt Mask R-CNN [15] with ResNet-101 [16] backbone with FPN [24] and train the models with the random sampler or the repeat factor sampler (RFS) by 2x schedule.

As shown in Table 1, Seesaw Loss significantly outperforms Cross-Entropy (CE) loss by 6.0% *AP* with random sampler and 2.1% *AP* on the stronger baseline with RFS. The improvements on AP_r , AP_c , and AP_f with both samplers reveals the effectiveness of Seesaw Loss on categories with different frequency. We further integrate the proposed Normalized Mask Prediction (Norm Mask) into Mask R-CNN [15] with Seesaw Loss. Without extra cost, the overall *AP* is improved from 26.6% to 27.1% and 27.6% to 28.1% with random sampler and RFS, respectively.

Apart from the CE loss baseline, we further compare Seesaw Loss with recent designs for long-tailed instance segmentation, *i.e.*, Equalization Loss (EQL) [38] and Balanced Group Softmax (BAGS) [23], in Table 1. Seesaw Loss outperforms EQL by 3.9% *AP* and 1.4% *AP*, and outperforms BAGS by 1.0% *AP* and 1.8% *AP* with random sampler and RFS, respectively. Seesaw Loss also achieves higher AP_r , AP_c and AP_f than the two methods consistently. Notably, EQL and BAGS achieve lower AP_f than the CE baseline while Seesaw Loss does not. This phenomenon indicates that these two methods improve the performance of rare and common categories while sacrificing frequent categories.

We further compare Seesaw Loss with previous methods [23, 38] with both random sampler and RFS on Cascade Mask R-CNN [1]. It’s a representative framework of cascade methods [3, 1] that outperforms Mask R-CNN [15]. As shown in Table 1, Seesaw Loss performs much superior to previous works [23, 38] on Cascade Mask R-CNN. Specifically, Seesaw Loss improves the baseline by 6.4% *AP* and 2.3% *AP* with random sampler and RFS, respectively. With Normalized Mask Prediction, Cascade Mask R-CNN with Seesaw Loss finally achieves 29.6% *AP* and 30.1% *AP* with the two samplers, respectively. Moreover, Seesaw Loss is also evaluated on *test-dev* split and consistently obtains significant gains over the CE baseline.

4.3. Ablation study

We conduct a comprehensive ablation study to verify the effectiveness of each design choice in the proposed method. **Components in Seesaw Loss.** There are three components in Seesaw Loss: mitigation factor, compensation factor, and normalized linear activation. We evaluated each component on Mask R-CNN with RFS (Table 2). The mitigation

Table 1: Performance comparison of Mask R-CNN [15] and Cascade Mask R-CNN [1] with Cross-Entropy (CE) Loss, Equalization Loss (EQL) [38], Balanced Group Softmax (BAGS) [23], and Seesaw Loss on LVIS v1 dataset [12]. The ResNet-101 [16] w/ FPN [24] is adopted as backbone. All models are trained with random sampler or repeat factor sampler (RFS) [12] by 2x schedule in an end-to-end pipeline. Norm Mask indicates the proposed Normalized Mask Prediction in Sec. 3.3.

Framework	Sampler	Loss	Split	AP	AP _r	AP _c	AP _f	AP ^{box}
Mask R-CNN [15]	Random	Cross-Entropy (CE)	val	20.6	0.8	19.3	30.7	21.7
		Equalization Loss (EQL) [38]		22.7	3.7	23.3	30.4	24.0
		Balanced Group Softmax (BAGS) [23]		25.6	17.3	25.0	30.1	26.4
		Seesaw Loss		26.6	18.1	25.8	31.2	27.4
		Seesaw Loss + Norm Mask		27.1	18.7	26.3	31.7	27.4
Mask R-CNN [15]	RFS [12]	Cross-Entropy (CE)	val	25.5	16.6	24.5	30.6	26.6
		Equalization Loss (EQL) [38]		26.2	17.0	26.2	30.2	27.6
		Balanced Group Softmax (BAGS) [23]		25.8	16.5	25.7	30.1	26.5
		Seesaw Loss (Ours)		27.6	20.6	27.3	31.1	28.9
		Seesaw Loss + Norm Mask (Ours)		28.1	20.0	28.0	31.8	28.9
Cascade Mask R-CNN [1]	Random	Cross-Entropy (CE)	val	22.6	2.4	22	32.2	25.5
		Equalization Loss (EQL) [38]		24.3	5.1	25.3	31.7	27.3
		Balanced Group Softmax (BAGS) [23]		27.9	19.6	27.7	31.6	31.5
		Seesaw Loss (Ours)		29.0	21.1	28.6	33.0	32.8
		Seesaw Loss + Norm Mask (Ours)		29.6	20.3	29.3	34.0	32.7
Cascade Mask R-CNN [1]	RFS [12]	Cross-Entropy (CE)	val	27.0	16.6	26.7	32.0	30.3
		Equalization Loss (EQL) [38]		27.1	17.0	27.2	31.4	30.4
		Balanced Group Softmax (BAGS) [23]		27.0	16.9	26.9	31.7	30.2
		Seesaw Loss (Ours)		29.3	21.7	29.2	32.8	32.8
		Seesaw Loss + Norm Mask (Ours)		30.1	21.4	30.0	33.9	32.8
Mask R-CNN [15]	RFS [12]	Cross-Entropy (CE)	test-dev	25.1	13.0	24.8	30.8	-
		Seesaw Loss + Norm Mask (Ours)		27.9	20.3	27.1	32.2	-
Cascade Mask R-CNN [1]	RFS [12]	Cross-Entropy (CE)	test-dev	26.4	15.5	25.5	32.3	-
		Seesaw Loss + Norm Mask (Ours)		30.0	23.0	29.3	34.1	-

Table 2: Ablation study of each design in Seesaw Loss with Mask R-CNN w/ R-50 FPN backbone and repeat factor sampler. MF, CF, NLA indicate mitigation factor, compensation factor, and normalized linear activation, respectively.

MF	CF	NLA	AP	AP _r	AP _c	AP _f	AP ^{box}
			23.7	13.5	22.8	29.3	24.7
✓			25.1	16.7	24.5	29.4	26.2
	✓		24.1	13.2	23.5	29.5	25.1
✓	✓		25.7	19.1	25.0	29.4	26.8
		✓	24.7	15.0	24.1	29.6	25.6
✓	✓	✓	26.4	19.6	26.1	29.8	27.4

factor that mitigates the overwhelming punishments on rare classes leads to a significant improvement from 23.7% AP to 25.1% AP. Notably, it improves the AP_r of rare classes by 2.8% AP. The compensation factor increases the punishments of a class when it observes false positives on that class to reduce misclassification. It improves the baseline by 0.4% AP. The combination of the mitigation and the compensation factors achieves 25.7% AP, outperforming the performance of mitigation factor by 0.6% AP. It reveals the effectiveness of instance-wise compensation to avoid misclassification. The normalized linear activation is another important component in Seesaw Loss, which reduces the scale invariance of weights and features across differ-

Table 3: The effectiveness of the normalized linear activation in different methods. EQL and BAGS indicates Equalization Loss and Balanced Group Softmax Loss, respectively.

Method	NLA	AP	AP _r	AP _c	AP _f	AP ^{box}
EQL		25.1	17.4	24.8	28.8	26.1
EQL	✓	25.4	17.8	25.2	29.1	26.5
BAGS		24.7	15.6	24.4	28.9	25.2
BAGS	✓	25.5	19.2	25.0	28.9	25.8
Seesaw		25.7	19.1	25.0	29.4	26.8
Seesaw	✓	26.4	19.6	26.1	29.8	27.4

ent categories. It improves the baseline performance from 23.7% to 24.7% AP. Seesaw Loss combining all these three components achieves 26.4% AP.

Normalized linear activation. We empirically find that normalized linear activation (NLA) helps to improve the performance of both CE Loss and Seesaw Loss. Therefore, we further integrate NLA with equalization loss (EQL) and balanced group softmax (BAGS) for fair comparisons. Results in Table 3 show that NLA improves the performance of EQL and BAGS by 0.3% and 0.8% AP, respectively. It is noteworthy that Seesaw Loss outperforms EQL and BAGS no matter whether NLA is adopted.

Cumulative Sample Numbers. Different from previous works [2, 38, 23] that rely on the pre-computed frequency

Table 4: Comparison of different approaches to obtain sample numbers of different categories in Seesaw Loss. From dataset indicates to obtain the distribution of instance numbers directly from the *train* split. Pre-Record indicates to load the cumulative sample numbers from a model trained with Seesaw Loss. Online indicates to accumulate the sample numbers during training.

Source	AP	AP _r	AP _c	AP _f	AP ^{box}
From dataset	26.1	19.7	25.6	29.5	27.2
Pre-Record	26.3	19.6	25.8	29.8	27.4
Online	26.4	19.6	26.1	29.8	27.4

Table 5: Ablation study of the hyper-parameter p in $\left(\frac{N_j}{N_i}\right)^p$ of mitigation factor. The normalized linear activation is not adopted in this table. $p = 0.8$ is the default setting in other experiments.

p	AP	AP _r	AP _c	AP _f	AP ^{box}
0.2	24.4	14.7	23.6	29.4	25.4
0.4	24.9	15.2	24.6	29.6	26.0
0.6	25.4	17.9	24.6	29.5	26.5
0.8	25.7	19.1	25.0	29.4	26.8
1.0	25.5	17.6	25.2	29.2	26.4
1.2	25.3	18.1	24.7	29.0	26.5

Table 6: Ablation study of the hyper-parameter q in $\left(\frac{\sigma_j}{\sigma_i}\right)^q$ of compensation factor. The normalized linear activation is not adopted in this table. $q = 2.0$ is the default setting.

q	AP	AP _r	AP _c	AP _f	AP ^{box}
0.5	25.4	17.6	25.0	29.4	26.3
1.0	25.5	17.5	25.1	29.4	26.6
1.5	25.5	17.8	25.1	29.3	26.8
2.0	25.7	19.1	25.0	29.4	26.8
2.5	25.6	17.7	25.2	29.4	26.5
3.0	25.4	17.3	24.9	29.5	26.4

distribution of categories in the dataset, Seesaw Loss accumulates the sample numbers of each category during training. We compare different approaches to obtain the sample numbers of categories for Seesaw Loss (Table 4). Directly using the statistics of *train* split decreases the performance of Seesaw Loss by 0.3% AP. The reason lies in that the data sampler, *e.g.*, repeat factor sampler, changes the frequency distribution of categories during training. We also explore loading the pre-recorded distribution of training samples from a model trained with Seesaw Loss. It achieves a similar performance with accumulating the training samples online (26.3% AP vs. 26.4% AP). These results verify the effectiveness and simplicity of online accumulating.

Hyper-parameters. We study the hyper-parameters, *i.e.*, p , q , τ , adopted in different components of Seesaw Loss. The normalized linear activation is not applied when studying the mitigation and compensation factors (25.7% AP with this setting). In Table 5, we explore p in $\left(\frac{N_j}{N_i}\right)^p$ of mitigation factor. p controls the magnitude to mitigate the punishments on rare classes. A higher value of p will reduce punishments more, as well as increase the risk of inducing

Table 7: Ablation study of the temperature term τ in Normalized Linear Activation. $\tau = 20$ is the default setting.

τ	AP	AP _r	AP _c	AP _f	AP ^{box}
10	24.7	16.7	24.1	28.8	25.7
15	26.0	19.0	25.5	29.6	27.0
20	26.4	19.6	26.1	29.8	27.4
25	26.2	19.2	25.7	29.9	27.4
30	25.9	16.4	26.2	29.8	27.0

Table 8: Ablation study of the objectness branch in CE loss baseline and Seesaw Loss. OBJ indicates whether the objectness branch is adopted.

Method	OBJ	AP	AP _r	AP _c	AP _f	AP ^{box}
CE		24.0	14.0	23.4	29.0	24.9
CE	✓	23.7	13.5	22.8	29.3	24.7
Seesaw		25.3	16.0	25.1	29.4	26.5
Seesaw	✓	26.4	19.6	26.1	29.8	27.4

false positives of tail classes. Therefore, it is critical to find a suitable p . Results show that $p = 0.8$ achieves the best performance. In Table 6, we explore q in $\left(\frac{\sigma_j}{\sigma_i}\right)^q$ of the compensation factor. q controls the magnitude to compensate the reduced punishments on tail classes when false positives are observed. We study the effectiveness of different q and find $q = 2.0$ achieves the best performance. Notably, q is robust across different values as the best value is only 0.3% AP better than the worst value. In Table 7, we study the temperature τ in normalized linear activation (NLA). τ determines the variance of the classifier’s predicted logit z . If τ is too small, the variance of z is insufficient to distinguish positive and negative samples. However, if τ is too big, the target of balancing the variance in weights and features between different categories will be sacrificed. We choose $\tau = 20$ in the NLA as it achieves the best performance.

Objectness Branch. In a common practice of object detection, the classifier predicts $C + 1$ scores for C foregrounds categories and one background category. Due to the extremely imbalanced distribution between foregrounds and backgrounds, Seesaw Loss will tend to misclassify more backgrounds as foregrounds with this design. Thus, we adopt an extra objectness branch as described in Sec. 3.3. The results in Table 8 shows that the objectness branch does not improve the Cross-Entropy loss baseline but is critical to Seesaw Loss. The objectness branch helps to avoid reducing the backgrounds’ punishments on C foreground categories. As a result, the objectness branch brings gains on Seesaw Loss across categories with different frequency, and improves the overall AP from 25.3% to 26.4%.

Training Pipeline. Apart from the end-to-end training pipeline, we further explore the popular decoupling training pipeline [20, 23] on Mask R-CNN. Specifically, we pre-train the Mask R-CNN with Cross-Entropy loss using either random sampler or repeat factor sampler for 2x sched-

Table 9: Explorations of decouple training pipelines [20, 23] for instance segmentation. Mask R-CNN with different classification loss is finetuned for 1x schedule with RFS sampler. We adopt pre-trained models with random (P-Rand) or repeat factor sampler (P-RFS) for 2x schedule. ‘DE-’ indicates the model is trained with decoupling training pipelines.

Loss	P-Rand	P-RFS	AP	AP _r	AP _c	AP _f	AP ^{box}
EQL			25.1	17.4	24.8	28.8	26.1
BAGS			24.7	15.6	24.4	28.9	25.2
Seesaw			26.4	19.6	26.1	29.8	27.4
DE-EQL	✓		23.9	12.9	23.7	28.9	25.4
DE-EQL		✓	25.2	15.9	25.5	28.9	26.5
DE-BAGS	✓		25.4	16.3	25.1	29.7	26.3
DE-BAGS		✓	25.6	17.0	25.3	29.8	26.6
DE-Seesaw	✓		25.1	16.6	24.6	29.5	26.2
DE-Seesaw		✓	25.8	18.7	25.3	29.6	26.9

ule. Then we finetune the final fully-connected layer of the classifier with all other components fixed. The 1x schedule and repeat factor sampler is adopted during finetuning. As shown in Table 9, Seesaw Loss with the pre-trained model on repeat factor sampler (P-RFS) achieves 25.8% AP, outperforming other methods with decoupling training pipeline. Notably, Seesaw Loss performs better with the end-to-end training pipeline than with the decoupling training pipeline. It indicates Seesaw Loss provides a simpler and more effective solution to long-tailed instance segmentation without relying on complex training pipelines.

4.4. Long-Tailed Image Classification

To show the versatility of Seesaw Loss, we apply it for long-tailed image classification task on ImageNet-LT [28] dataset. ImageNet-LT [28] is generated from the ImageNet-2012 [37] dataset with long-tailed distributed categories in training set. There are 115.8k images of 1000 categories with a maximum number of 1280 images and a minimum number of 5 images. The performance is evaluated with top-1 accuracy on all categories and the accuracies for Many Shot (> 100 images), Medium Shot (20~100 images) and Few Shot (< 20 images) categories are also reported.

We adopt two training pipelines: end-to-end training and decoupling training [20]. We use ResNeXt-50 [49] backbone and SGD optimizer with momentum of 0.9, initial learning rate of 0.2, batch size of 512, and cosine learning rate [29] following [20]. For the end-to-end training pipeline, the model is trained for 90 epochs. For decouple training pipeline [20], we load the pre-trained ResNeXt-50 [49] with Cross-Entropy Loss (CE), and finetune the classifier with class-balanced sampler while fixing all other layers for 10 epochs. Seesaw Loss in image classification mostly follows the hyper-parameters on the instance segmentation task except for q in the compensation factor. We adopt $q = 1$ for ImageNet-LT dataset. The study of q for ImageNet-LT dataset is shown in Table 11.

Table 10: Comparison of different methods on ImageNet-LT [28] test set. ResNeXt-50 [49] backbone is adopted in experiments. Decouple means using decouple training pipeline [20].

Method	Decouple	Overall	Many	Medium	Few
CE		44.4	65.9	37.5	7.7
Focal Loss [25]		43.3	64.5	36.3	7.8
CB-Focal [6]		45.3	60.4	40.6	19.2
EQL [6]		46.0	61.7	42.5	13.8
NCM [20]	✓	47.3	56.6	45.3	28.1
cRT [20]	✓	49.6	61.8	46.2	27.4
τ -norm [20]	✓	49.4	59.1	46.9	30.7
LWS [20]	✓	49.9	60.2	47.2	30.3
Seesaw	✓	49.7	60.7	46.8	28.9
Seesaw		50.4	67.1	45.2	21.4

Table 11: Comparison of hyper-parameter q in the compensation factor $\left(\frac{\sigma_j}{\sigma_i}\right)^q$ on ImageNet-LT [28].

q	Overall	Many	Medium	Few
0.5	49.6	66.2	44.4	20.9
1.0	50.4	67.1	45.2	21.4
1.5	49.6	66.4	44.3	20.7
2.0	49.4	66.5	44.0	20.3
2.5	48.4	65.8	42.9	18.7

We report the performance of Seesaw Loss in Table 10. Seesaw Loss improves top-1 accuracies of CE from 44.4% to 49.7% and 50.4% with the decoupling training and the end-to-end training pipeline, respectively. Similar to our observations on the instance segmentation task, Seesaw Loss performs better with the end-to-end training pipeline on image classification. The performance achieved by Seesaw Loss with the end-to-end pipeline is competitive among previous methods on ImageNet-LT [28].

5. Conclusion

In this paper, we propose Seesaw Loss for long-tailed instance segmentation. Seesaw Loss dynamically re-balances gradients of positive and negative samples for each category with two complementary factors. The mitigation factor reduces punishments to tail categories w.r.t. the ratio of cumulative training instances between categories. Meanwhile, the compensation factor increases the penalty of misclassified instances to avoid false positives. Experimental results demonstrate that Seesaw Loss provides a simpler and more effective solution to long-tailed instance segmentation without relying on complex training pipelines.

Acknowledgements. This research was conducted in collaboration with SenseTime. This work is supported by GRF 14203518, ITS/431/18FX, CUHK Agreement TS1712093, NTU NAP, A*STAR through the Industry Alignment Fund - Industry Collaboration Projects Grant, Shanghai Committee of Science and Technology, China (Grant No. 20DZ1100800).

References

- [1] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: High quality object detection and instance segmentation. *arXiv preprint arXiv:1906.09756*, 2019. 1, 2, 4, 5, 6
- [2] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arachiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *NeurIPS*, 2019. 2, 6
- [3] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. Hybrid task cascade for instance segmentation. In *CVPR*, 2019. 1, 2, 3, 5
- [4] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 5
- [5] Xinlei Chen, Ross B. Girshick, Kaiming He, and Piotr Dollár. TensorMask: A foundation for dense object segmentation. In *ICCV*, 2019. 2
- [6] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. *CVPR*, 2019. 2, 8
- [7] Golnaz Ghiasi, Tsung-Yi Lin, Ruoming Pang, and Quoc V. Le. NAS-FPN: Learning scalable feature pyramid architecture for object detection. *CoRR*, abs/1904.07392, 2019. 2
- [8] Spyros Gidaris and Nikos Komodakis. LocNet: Improving localization accuracy for object detection. In *CVPR*, 2016. 2
- [9] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *CVPR*, 2018. 4
- [10] Ross Girshick. Fast R-CNN. In *ICCV*, 2015. 2
- [11] Ross Girshick, Ilija Radosavovic, Georgia Gkioxari, Piotr Dollár, and Kaiming He. Detectron. <https://github.com/facebookresearch/detectron>, 2018. 5
- [12] Agrim Gupta, Piotr Dollar, and Ross Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 1, 2, 3, 5, 6
- [13] Haibo He and Edwardo A. Garcia. Learning from imbalanced data. *IEEE TKDE*, 2009. 2
- [14] Kaiming He, Ross Girshick, and Piotr Dollar. Rethinking ImageNet pre-training. In *ICCV*, 2019. 2
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask R-CNN. *ICCV*, 2017. 1, 2, 3, 4, 5, 6
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5, 6
- [17] Xinting Hu, Yi Jiang, Kaihua Tang, Jingyuan Chen, Chunyan Miao, and Hanwang Zhang. Learning to segment the tail. In *CVPR*, 2020. 2, 3
- [18] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Deep imbalanced learning for face recognition and attribute prediction. *IEEE TPAMI*, 2020. 2
- [19] Borui Jiang, Ruixuan Luo, Jiayuan Mao, Tete Xiao, and Yuning Jiang. Acquisition of localization confidence for accurate object detection. In *ECCV*, 2018. 2
- [20] 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. 2, 3, 4, 7, 8
- [21] Tao Kong, Fuchun Sun, Huaping Liu, Yuning Jiang, and Jianbo Shi. FoveaBox: Beyond anchor-based object detector. *CoRR*, abs/1904.03797, 2019. 2
- [22] Hei Law and Jia Deng. CornerNet: Detecting objects as paired keypoints. In *ECCV*, 2018. 2
- [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 *CVPR*, 2020. 2, 3, 4, 5, 6, 7, 8
- [24] 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. 5, 6
- [25] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 2, 8
- [26] 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, 4
- [27] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: single shot multibox detector. In *ECCV*, 2016. 2
- [28] 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, 4, 8
- [29] Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with warm restarts. In *ICLR*, 2017. 8
- [30] Dhruv Mahajan, Ross B. Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. Exploring the limits of weakly supervised pretraining. In *ECCV*, 2018. 2, 5
- [31] Mahyar Najibi, Mohammad Rastegari, and Larry S Davis. G-cnn: an iterative grid based object detector. In *CVPR*, 2016. 2
- [32] Pedro H. O. Pinheiro, Ronan Collobert, and Piotr Dollár. Learning to segment object candidates. In *NeurIPS*, 2015. 2
- [33] Pedro Oliveira Pinheiro, Tsung-Yi Lin, Ronan Collobert, and Piotr Dollár. Learning to refine object segments. In *ECCV*, 2016. 2
- [34] Joseph Redmon, Santosh Kumar Divvala, Ross B. Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016. 2
- [35] Joseph Redmon and Ali Farhadi. YOLO9000: Better, faster, stronger. In *CVPR*, 2017. 2
- [36] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 2, 4
- [37] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy,

- Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015. 8
- [38] 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, 5, 6
- [39] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: Fully convolutional one-stage object detection. *CoRR*, abs/1904.01355, 2019. 2
- [40] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. CosFace: Large margin cosine loss for deep face recognition. *CVPR*, 2018. 4
- [41] Jiaqi Wang, Kai Chen, Rui Xu, Ziwei Liu, Chen Change Loy, and Dahua Lin. CARAFE: Content-Aware ReAssembly of FEatures. In *ICCV*, 2019. 2
- [42] Jiaqi Wang, Kai Chen, Rui Xu, Ziwei Liu, Chen Change Loy, and Dahua Lin. Carafe++: Unified content-aware reassembly of features, 2020. 2
- [43] Jiaqi Wang, Kai Chen, Shuo Yang, Chen Change Loy, and Dahua Lin. Region proposal by guided anchoring. In *CVPR*, 2019. 2
- [44] Jiaqi Wang, Wenwei Zhang, Yuhang Cao, Kai Chen, Jiangmiao Pang, Tao Gong, Jianping Shi, Chen Change Loy, and Dahua Lin. Side-aware boundary localization for more precise object detection. In *ECCV*, 2020. 2
- [45] Ning Wang, Yang Gao, Hao Chen, Peng Wang, Zhi Tian, and Chunhua Shen. NAS-FCOS: Fast neural architecture search for object detection. *CoRR*, abs/1906.04423, 2019. 2
- [46] Tao Wang, Yu Li, Bingyi Kang, Junnan Li, Jun Hao Liew, Sheng Tang, Steven C. H. Hoi, and Jiashi Feng. The devil is in classification: A simple framework for long-tail instance segmentation. In *ECCV*, 2020. 3
- [47] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. SOLO: Segmenting objects by locations. *CoRR*, abs/1912.04488, 2019. 2
- [48] Tong Wu, Qingqiu Huang, Ziwei Liu, Yu Wang, and Dahua Lin. Distribution-balanced loss for multi-label classification in long-tailed datasets. In *ECCV*, 2020. 2
- [49] Saining Xie, Ross Girshick, Piotr Dollar, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017. 8