

FASA: Feature Augmentation and Sampling Adaptation for Long-Tailed Instance Segmentation

Yuhang Zang¹ Chen Huang² Chen Change Loy¹✉

¹S-Lab, Nanyang Technological University ²Carnegie Mellon University

{zang0012, ccloy}@ntu.edu.sg chen-huang@apple.com

Abstract

Recent methods for long-tailed instance segmentation still struggle on rare object classes with few training data. We propose a simple yet effective method, Feature Augmentation and Sampling Adaptation (FASA), that addresses the data scarcity issue by augmenting the feature space especially for rare classes. Both the Feature Augmentation (FA) and feature sampling components are adaptive to the actual training status — FA is informed by the feature mean and variance of observed real samples from past iterations, and we sample the generated virtual features in a loss-adapted manner to avoid over-fitting. FASA does not require any elaborate loss design, and removes the need for inter-class transfer learning that often involves large cost and manually-defined head/tail class groups. We show FASA is a fast, generic method that can be easily plugged into standard or long-tailed segmentation frameworks, with consistent performance gains and little added cost. FASA is also applicable to other tasks like long-tailed classification with state-of-the-art performance.^{1 2}

1. Introduction

A growing number of methods are proposed to learn from long-tailed data in vision tasks like face recognition [17], image classification [29] and instance segmentation [13]. We focus on the problem of long-tailed instance segmentation, which is particularly challenging due to the class imbalance issue. State-of-the-art methods [1, 4, 14, 25] often fail to handle this issue and witness a large performance drop on rare object classes. Fig. 1 exemplifies the struggle of the competitive Mask R-CNN [14] baseline on highly imbalanced LVIS v1.0 dataset [13]. We can see that there are about 300 rare classes in the tail that have no more than 10^2 positive object instances. Such scarce training data result in poor performance for most of the tail classes, with

near-zero class probabilities predicted for them.

To address the data scarcity issue, one intuitive option is to over-sample images that contain the tail-class objects [13]. But the downside is that the over-sampled images will include more head-class objects at the same time due to the class co-occurrence within images. Hence for the instance segmentation task, re-sampling at the instance level is more desirable than at image level. Another option is data augmentation for considered objects, either in image space (e.g., random flipping) or in feature space (i.e., feature augmentation, on object regional features). Along this line, many methods have proven effective by augmenting the image/feature space of rare classes. The feature augmentation methods already show benefits for face recognition [44, 26], person Re-ID [26] and classification [20, 6]. However, these methods require manual design of class groups, e.g., using heuristics like class size. And they often involve two stages of feature training and feature transfer, which induce additional cost.

Feature augmentation remains little explored for long-tailed instance segmentation tasks. In this paper, we present an efficient and effective method, called Feature Augmentation and Sampling Adaptation (FASA). FASA does not require any elaborate transfer learning or loss design [23, 34, 41]. As a result, FASA keeps its simplicity, has extremely low complexity, while remains highly adaptive during training. In the proposed method, we perform online Feature Augmentation (FA) for each class. The FA module generates class-wise virtual features using a distributional prior with its statistics calculated from previously observed real samples. This allows FA to capture class distributions and adapts to evolving feature space.

The augmented features can still be limiting for rare classes. This is because the observed feature variance is often small with few training examples. To overcome this problem, we adapt the sampling density of virtual features for each class. This way, our synthesized virtual features still live in the true feature manifold, just with different sampling probabilities to avoid under-fitting or over-fitting. We propose an adaptive sampling method here: when aug-

¹GitHub: <https://github.com/yuhangzang/FASA>.

²Project page: <https://www.mmlab-ntu.com/project/fasa/index.html>.




Figure 1: **Class imbalance and the comparison of Mask R-CNN [14] baseline with and without FASA on LVIS v1.0 dataset.** (a) By adaptive feature augmentation and sampling, our method FASA largely alleviates the imbalance issue, especially for rare classes. (b) Compare the prediction results of FASA vs Mask R-CNN baseline regarding average category probability scores. The baseline model predicts near-zero scores for rare classes. While with FASA, rare-class scores are significantly boosted which merits final performance. (c) FASA brings consistent improvements over different backbone models in mask AP_r defined on rare classes. Such gains come at a very low cost (training time only increases by $\sim 3\%$ on average).

mented features improve the corresponding class performance in validation loss, the feature sampling probabilities get increased, otherwise decreased. Such a loss-optimized sampling method is effective to re-balance the model predicting performance, see Fig 1(b).

It is noteworthy that our full FASA approach handles virtual features only, thus it can serve as a plug-and-play module applicable to existing methods that learn real data with any re-sampling schemes or loss functions (standard or tailored for class imbalance). Comprehensive experiments of instance segmentation on LVIS v1.0 [13] and COCO-LT [38] datasets demonstrate FASA as a generic component that can provide consistent improvement to other methods. Take LVIS dataset as an example. FASA improves Mask R-CNN [14] by 9.0% and 3.3% in mask AP metrics for rare and overall classes respectively, and improves a contemporary loss design [34] by 10.3% and 2.3%. Moreover, these gains come at the cost of merely a $\sim 3\%$ increase in training time, see Fig. 1(c). Furthermore, FASA can generalize beyond the instance segmentation task, achieving state-of-the-art performance on long-tailed image classification as well.

To summarize, the main **contributions** of this work is a fast and effective feature augmentation and sampling approach for long-tailed instance segmentation. The proposed FASA can be combined with existing methods in a plug-and-play manner. It achieves consistent gains and surpasses more sophisticatedly designed state-of-the-art methods. FASA also generalizes well to other long-tailed tasks.

2. Related Work

Long-tailed classification. For the long-tailed classification task, there is a rich body of widely used methods including data re-sampling [3] and re-weighting [2, 7]. Recent works [19, 48] reveal the effectiveness of using different sampling schemes in decoupled training stages.

Instance-balanced sampling is found useful for the first feature learning stage, which is followed by a classifier fine-tuning stage with class-balanced sampling.

Long-tailed instance segmentation. To handle class imbalance in instance segmentation task, recent methods still heavily rely on the ideas of data re-sampling [13, 31, 16, 41], re-weighting [34, 30, 33, 37, 49, 39, 47] and decoupled training [23, 38]. For re-sampling, class-balanced sampling [31] and Repeat Factor Sampling (RFS) [13] are conducted at the image level. However, image-level re-sampling sometimes worsens the imbalance at instance level due to the instance co-occurrence within images. The methods of data-balanced replay [16] and NMS re-sampling [41] belong to the category of instance-level re-sampling. For the re-weighting scheme, Equalization loss v1 [34] and v2 [33] are such representative methods that re-weigh the sigmoid loss. More recent works [23, 41] attempt to divide imbalanced classes into relatively balanced class groups for robust learning. However, the class grouping process relies on static heuristics like class size or semantics which is not optimal. Tang *et al.* [35] studied the sample co-existence effect in the long-tailed setting and proposed de-confounded training. Seesaw Loss [37] dynamically rebalances the gradients of positive and negative samples, especially for rare categories. Surprisingly, data augmentation as a simple technique, has been barely studied for long-tailed instance segmentation. In this paper, we show that competitive results can be obtained from smart data augmentation and re-sampling in the feature space, while being intuitively simple and computationally efficient. Also our approach is orthogonal to prior works, and can be easily combined with them to achieve consistent improvements.

Data augmentation. To avoid over-fitting and improve generalization, data augmentation is often used during network training. In the context of long-tailed recognition, data augmentation can also be used to supplement effective training data of under-represented rare classes, which helps

to re-balance the performance across imbalanced classes.

There are two main categories of data augmentation methods: augmentation in the image space and feature space (*i.e.*, feature augmentation). Commonly used image-level augmentation methods include random image flipping, scaling, rotating and cropping. A number of advanced methods have also been proposed, such as Mixup [46] and Cutmix [45]. For our considered instance segmentation task, the copy-and-paste technique such as InstaBoost [11] and Ghiasi *et al.* [12] and prove effective. More recent works try to synthesize novel images using GAN [28] or semi-supervised learning method [42]. Compared to the well established image-level data augmentation technique in deep learning, feature augmentation has not yet received enough attention. On the other hand, Feature Augmentation (FA) directly manipulates the feature space, hence it can reshape the decision boundary for rare classes. Classic FA is built on SMOTE [3]-typed methods that interpolate neighboring feature points. More recently, researchers propose to Manifold Mixup [36] or MoEx [21] for better performance. ISDA [40] augments data samples by translating CNN features along the semantically meaningful directions. Such methods are often not directly applied for imbalanced class discrimination and suffer from high complexity.

Some recent works have shown that feature augmentation benefits long-tail tasks such as face recognition [44], person Re-ID [26], or long-tail classification [6, 20, 27, 22]. However, we observe that these methods have limitations when apply them to the long-tailed instance segmentation datasets such as LVIS [13]. Due to the high computational cost of instance segmentation task, some methods [44, 20, 6] rely on two-phase pipeline or large memory of historical features[27] incur high cost in time and memory thus less efficient. The special background class exists for the instance segmentation task (no class anchor), makes method [26] that rely on margin-based classification loss less effective. In addition, the small batch size of the instance segmentation frameworks limits the performance of approach [6] that rely on mining confusing categories. Our approach is specially designed for long-tailed instance segmentation, waives the need for complex two-phase training and the associated costs. We compare with them in our experiments, see Sec 4.1.

3. Methodology

In this section, we introduce the proposed Feature Augmentation and Sampling Adaptation (FASA) approach consisting of two components: 1) adaptive Feature Augmentation (FA) that generates virtual features to augment the feature space of all classes (especially for rare classes), 2) adaptive Feature Sampling (FS) that dynamically adjusts the sampling probability of virtual features for each class.

To better illustrate how FASA works for long-tailed in-

stance segmentation, we take the Mask R-CNN [14] framework as our baseline and show an example of combining FASA and the segmentation baseline. The overall pipeline is shown in Fig. 2. Note FASA is a standalone module for feature augmentation, leaving the baseline module unaltered. Therefore, FASA serves as a plug-and-play module and it can combine with and facilitate stronger baselines other than Mask R-CNN. We will show this flexibility in our experiments.

Under the Mask R-CNN framework, the standard multi-task loss is defined for each Region of Interest (RoI):

$$\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{box} + \mathcal{L}_{mask}. \quad (1)$$

For simplicity, we apply FASA to the classification branch only, which is the most vulnerable branch on long-tailed data as supported in Fig. 1(b). More discussion is provided in the supplementary materials. It is possible to augment other branches using FASA as well. We leave that to future work and expect further improvements.

3.1. Adaptive Feature Augmentation

Mask R-CNN is originally trained on “real” feature embeddings of positive region proposals generated by the Region Proposal Network (RPN). For each class, we aim to augment its real features that are likely scarce (*e.g.*, for rare classes). An ideal FA component should have the following properties: 1) generate diverse virtual features to enrich the feature space of corresponding class, 2) the generated virtual features capture intra-class variations accurately and do not deviate much from the true manifold which hinders learning, 3) adapt to the actual feature distribution that is evolving during training, 4) high efficiency.

To this end, we maintain an online Gaussian prior from previously observed real features, which fulfills the above-mentioned requirements. We found such prior suffices for FA purpose by generating up-to-date and diverse features even when the actual features are not Gaussian. Concretely, for each foreground class c within the current batch, we can calculate the corresponding feature mean $\mu_c^t \in \mathbb{R}^d$ and standard deviation $\sigma_c^t \in \mathbb{R}^d$ which together define a Gaussian feature distribution. Given the noisy nature of μ_c^t and σ_c^t , we use them to continuously update the more robust estimates μ_c and σ_c instead, via a momentum mechanism:

$$\begin{aligned} \mu_c &\leftarrow (1 - m) \cdot \mu_c + m \cdot \mu_c^t, \\ \sigma_c &\leftarrow (1 - m) \cdot \sigma_c + m \cdot \sigma_c^t, \end{aligned} \quad (2)$$

where m is set to 0.1 in all experiments. Then according to a Gaussian prior with up-to-date μ_c and σ_c , we generate the class-wise virtual features \hat{x}_c by random perturbation under the feature independence assumption:

$$\hat{x}_c = \mu_c + \sigma_c \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, I_d). \quad (3)$$

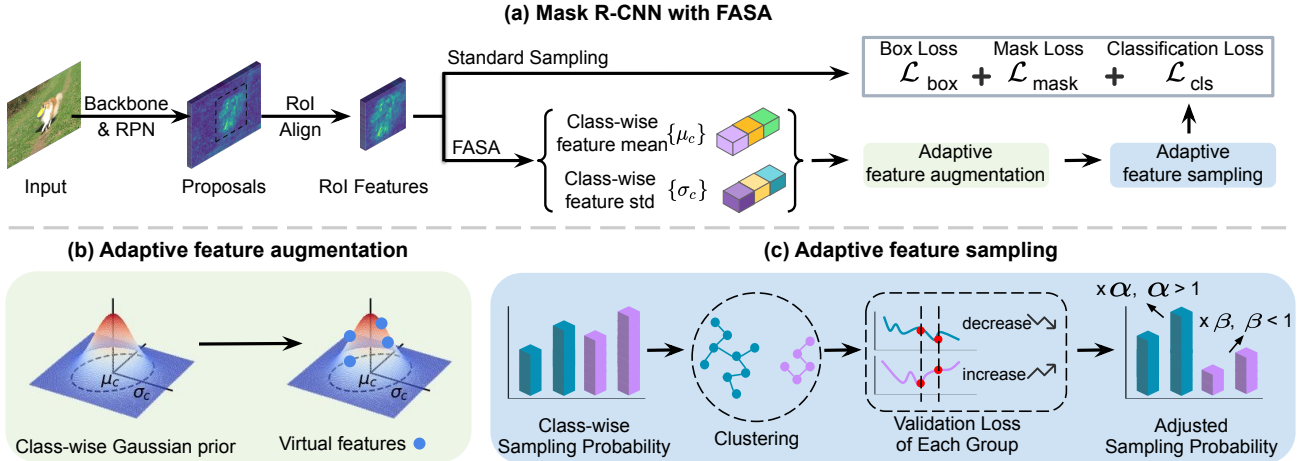


Figure 2: (a) The pipeline of Mask R-CNN combined with the proposed FASA, a standalone module that generates virtual features to augment the classification branch for better performance on long-tailed data. FASA maintains class-wise feature mean and variance online, followed by (b) adaptive feature augmentation and (c) adaptive feature sampling.

The feature independence assumption allows for efficient FA, where augmented features are treated as IID random variables. We consider the covariance matrix of the Gaussian to be diagonal, thus largely reduce the complexity from d^2 to d . Such generated virtual features $\{\hat{x}_c\}$ will need to be re-sampled (detailed later). Finally, both re-sampled $\{\hat{x}_c\}$ and real features $\{x_c\}$ (with its own sampling strategy in any baseline method) are sent to compute \mathcal{L}_{cls} in Eq. (1).

3.2. Adaptive Feature Sampling

For a wise use of the generated virtual features for each class, we propose an adaptive Feature Sampling (FS) scheme to effectively avoid under-fitting or over-fitting from FA. The sampling process operates in a *relative* fashion: if virtual features improve the corresponding class performance, their sampling probabilities get increased, otherwise decreased. Such relative adjustments of virtual feature sampling probabilities cater to the changing needs of FA during training. We can imagine that FA may be useful for rare classes during some training periods, but other training stages may call for decreased amount of FA to avoid over-fitting. By contrast, a static predefined sampling distribution would be independent of training dynamics and thus suboptimal.

Parametric sampling formulation. Note we still need an initial feature sampling distribution so as to adjust it later on. Obviously, we can benefit from a good initialization that avoids initial biased FA and costly adjustments. We choose to initialize the sampling distribution simply based on the inverse class frequency. It favors rare classes with higher sampling probabilities p_c and is useful in our class-imbalanced setting with minimal assumption about data distribution. Now we have a predefined skewed sampling distribution at the beginning. Then we dynamically scale the

per-class sampling probability p_c as follows:

$$p_c = s_c \cdot \left(\frac{1}{N_c} \right) / \sum_{c=1}^C \left(\frac{1}{N_c} \right), \quad (4)$$

where s_c is the scaling factor to be estimated at each time, and N_c denotes the class size of class c .

Adaptive sampling scheme. Recall the insight behind adapting the sampling probability p_c is that: if FA does improve the performance of class c , then we should generate more virtual features and increase p_c ; if we observe worse performance from FA, then p_c is probably big enough and we should decrease it to avoid over-fitting with the augmented features. In practice, we use multiplicative adjustments to update p_c every epoch. Specifically, we increase p_c to $\min(1, p_c \cdot \alpha)$ with $\alpha = 1.1$, and decrease p_c to $\max(0, p_c \cdot \beta)$ with $\beta = 0.9$.

One can adopt different performance metrics to guide the adjustment of p_c . Existing adaptive learning systems are either based on the surrogate loss [32] or the more ideal actual evaluation metric [18]. However for large-scale instance segmentation tasks, frequent evaluation of metrics like mAP on large datasets is very expensive. Therefore, we use the classification loss \mathcal{L}_{cls} on validation set to adjust p_c for our large-scale segmentation task. For the validation set, we apply Repeat Factor Sampling (RFS) [13] method to balance the class distribution and provide more meaningful losses. The aforementioned setting works reasonably well in our experiments.

Group-wise adaptation. To make our method more robust to different scenarios, we need to address two common challenges in the long-tailed instance segmentation task. First, the per-class loss can be extremely noisy on limited evaluation data, *e.g.*, for rare classes. In this case, both loss and

Table 1: Ablation study of the proposed FASA method on LVIS v1.0 *validation* set. ‘FA’ denotes adaptive Feature Augmentation, and ‘FS’ denotes adaptive Feature Sampling. The metrics AP, AP_r, AP_c and AP_f denote the mask mAP for overall, rare, common and frequent class groups.

FA	FS	AP	AP _r	AP _c	AP _f
✗	✗	20.8	8.0	20.2	27.0
✓	✗	22.3 (+1.5)	12.7 (+5.7)	21.7	26.9
✓	✓	23.7 (+2.9)	17.8 (+9.8)	22.9	27.2

metric cannot serve as reliable performance indicators. Second, the evaluation data may simply not be available for all classes. For example, the validation set of LVIS dataset [13] only has 871 classes out of 1203 from training set. This makes it impossible to evaluate loss for the other 332 classes on validation set.

To solve the aforementioned issues, we propose to cluster all training classes into super-groups. Then we compute the average validation loss of within-group classes, and adjust their feature sampling probabilities together. In other words, we adjust the sampling probabilities by a single scaling factor (α or β) depending on the average of group-wise loss. By doing so, those classes with missing evaluation data can be safely ignored when computing the loss average, but their sampling probabilities can still be updated along with other classes within the same group. Moreover, the group-wise update is less noisy since it is based on the loss average computed on bigger data (from multiple classes).

For clustering, rather than using predefined heuristics like class size or semantics as in [23, 41], we use the on-line class-wise feature mean μ_c and standard deviation σ_c . We adopt the density-based [10] clustering algorithm using the following distance based on Fisher’s ratio:

$$d_{ij} = \frac{(\mu_i - \mu_j)^2}{\sigma_i^2 + \sigma_j^2}. \quad (5)$$

The resulting super-groups are much more adaptive and meaningful than the predefined ones (*e.g.*, rare, common, and frequent class groups), and facilitate better group-wise feature re-sampling. Please refer to the supplementary materials for the visualization of some super-groups of semantically similar classes.

4. Experiments

Datasets. Our experiments are conducted on two datasets: LVIS v1.0 [13] with 1203 categories and COCO-LT [38] with 80 categories. Both of them are designed for long-tailed instance segmentation with highly class-imbalanced distribution. We choose LVIS v1.0 dataset over LVIS v0.5 since the former has more labeled data for meaningful evaluations and comparisons. LVIS v1.0 dataset defines three class groups: rare [1, 10), common [10, 100), and frequent [100, -) based on the number of images that contain at least

Table 2: Comparing our FASA with 1) **interpolation** based methods SMOTE [3] and MoEx [21], 2) **copy-and-paste** base methods InstaBoost [11], 3) **feature augmentation** based methods Yin *et al.* [44], Liu *et al.* [26], Chu *et al.* [6] on LVIS v1.0 *validation* set. The baseline (top row) denotes Mask R-CNN [14] without any augmentation.

Augmentation	AP	AP _r	AP _c	AP _f
✗	20.8	8.0	20.2	27.0
SMOTE [3]	21.5 (+0.7)	10.2 (+2.2)	20.9	27.1
MoEx, CVPR’21 [21]	21.2 (+0.4)	9.2 (+1.2)	20.6	27.1
InstaBoost, CVPR’19 [11]	21.4 (+0.6)	10.3 (+2.3)	20.7	27.2
Yin <i>et al.</i> , CVPR’19 [44]	21.6 (+0.8)	11.1 (+3.1)	20.9	27.1
Liu <i>et al.</i> , CVPR’20 [26]	21.0 (+0.2)	9.6 (+1.6)	20.1	26.8
Chu <i>et al.</i> , ECCV’20 [6]	21.4 (+0.6)	9.7 (+1.7)	21.0	27.0
FASA (ours)	23.7 (+2.9)	17.8 (+9.8)	22.9	27.2

one instance of the corresponding class. Similarly, COCO-LT dataset defines four class groups [1, 20), [20, 400), [400, 8000), [8000, -). We use the standard mean average precision (mAP) as the evaluation metric. Using this metric on different class groups can well characterize the long-tailed class performance. Following [13], we denote the corresponding performance metrics of rare, common and frequent class groups as AP_r, AP_c and AP_f. Recently, Dave *et al.* [8] proposed that the mAP metric is sensitive to changes in cross-category ranking, and introduced two complementary metrics AP^{Fixed} and AP^{Pool}. We also report the performance of FASA under the AP^{Fixed} and AP^{Pool} metrics in the supplementary materials.

Implementation details. Our implementation is based on the MMDetection [5] toolbox. We follow the same experimental setting as in [13, 34] for fair comparison. Please refer to supplementary materials for other details.

4.1. Ablation Study on LVIS

We first perform ablation studies on the large-scale LVIS dataset. We report the validation performance to ablate the core modules of our FASA approach.

Effectiveness of FA and FS. Table 1 verifies the critical roles of our adaptive modules for Feature Augmentation (FA) and Feature Sampling (FS). The baseline (top row) only performs repeat factor sampling [13] for real features, without using any FASA components. Our FA module (second row) significantly improves the performance, both for overall and rare classes. Our adaptive FS further boosts the performance, especially in the rare class groups (AP_r^m 12.7% \rightarrow 17.8%) with other groups being competitive. The results suggest the effectiveness of both the FA and FS components in improving the training performance.

Comparison with other augmentation methods. To further show the simplicity and efficacy of our method, we compare with the classic SMOTE method [3], MoEx [21], InstaBoost [11] and state-of-the-art methods [26, 6, 44] that

Table 3: Comparison between our adaptive feature sampling strategy and a static one (with fixed scaling factor s_c in Eq. (4)).

Sampling Method	AP	AP _r	AP _c	AP _f
Static ($s = 1$)	21.7	12.2	20.8	27.0
Static ($s = 5$)	22.3	12.7	21.7	26.9
Static ($s = 15$)	21.3	12.0	20.2	26.6
Adaptive	23.7	17.8	22.9	27.2

are specially designed for the long-tailed setting. Since [26, 6, 44] only report results on face recognition and person Re-ID, while no public codes are available, we re-implemented them and optimized their parameters and performance for LVIS (see the supplementary materials for implementation details).

Table 2 shows favorable results of our FASA when compared to others. Our gains are particularly apparent for AP_r and AP_c. Since SMOTE [3], MoEx [21] and InstaBoost [11] are not directly designed for long-tailed setting, we verify that our FASA obtains a favorable performance over them. For more related feature augmentation approaches [6, 26, 44], we observe some limitations when transferring them into the instance segmentation task. Liu *et al.* [26] are based on margin-based face recognition loss such as ArcFace [9] that constrains the margin between each instance and its class anchor. Unfortunately, since the instance segmentation task has to deal with the special *background* class that no distinct anchor, the margin-based loss does not perform well on LVIS. In contrast, our approach is not restricted by the form of loss function. As for Chu *et al.* [6], the performance is limited by the small batch size of the instance segmentation task, which does not guarantee the selection of top confusing samples. Both [6] and [44] apply feature transfer from head to tail classes. They use a two-stage training pipeline that requires a pre-trained model to extract features. In contrast, our proposed FA method can be trained end-to-end, much faster than [6, 44] and only incurs a small extra memory which is constant.

Analysis of adaptive FS. Recall in Eq. (4), our class-wise feature sampling probabilities are adaptively adjusted by the scaling factor s_c in an adaptive manner. Table 3 validates our design, showing that sampling is poor with a static sampling strategy. Specifically, we observe that $s_c = 5$ works best for AP_r, but is not optimal for AP_f that needs $s_c = 1$ instead. On the other hand, our adaptive FS tunes s_c online to effectively re-balance the performance across categories.

Fig. 3(a) depicts how the class-wise sampling probabilities change during training. Overall, the rare classes exhibit high sampling probabilities compared with the common and frequent classes. The rare-class sampling probabilities typically get increased to use more virtual features at the beginning. Then they gradually decrease to avoid over-fitting. Upon convergence, when the learning rate decreases, the sampling probabilities of rare categories see a bump to help

with the corresponding classifier ‘fine-tuning’. By contrast, the changes of sampling probabilities for the common and frequent classes are smaller. Fig. 3(b) further shows the dynamic changes of sampling probabilities within each class group.

4.2. Comparison with State of the Arts on LVIS

In this section, we evaluate the full FASA method against state-of-the-art methods on the LVIS v1.0 dataset. We consider the following representative methods: 1) *Repeat Factor Sampling (RFS)* [13] is an image-level data re-sampling technique. We use the same repeat factor $1e-3$ as proposed in the original paper. 2) *Equalization Loss (EQL)* [34] is a loss re-weighting method that aims to ignore the harmful gradients from rare classes. 3) *Classifier Re-training (cRT)* [19] first uses random sampling for feature representation learning, and then re-trains the classifier with repeat factor sampling. 4) *Balanced Group Softmax (BAGS)* [23] first performs class grouping and then makes the classification loss relatively balanced across those class groups. Grouping simply relies on the class sizes that are independent of training dynamics and suboptimal. 5) *Seesaw Loss* [37] balances ratio of cumulative gradients for the positive samples and negative samples of different categories.

Comparing using Mask R-CNN baseline. The aforementioned methods have shown solid improvements when combined with the Mask R-CNN baseline. Here we conduct a more comprehensive comparison with even stronger versions of these methods under different training schedules. Specifically, we experimented with the default 12 and 24 epochs schedules [5], as well as the decoupled two-stage training schedule [19, 23]. In the first stage, we train the model for 12 epochs with standard random data sampling and cross-entropy loss. Then in the second stage, we fine-tune for 12 epochs using those advanced re-sampling or re-weight approaches, such as RFS and BAGS. We refer to this schedule as ‘12+12’. We compare with all of these methods by plugging in our FASA module and comparing the performance difference.

Table 4 summarizes the comparison results. Specifically, we repeat each experiment three times with different random seeds and report the mean value of results. When we combine our FASA with vanilla Mask R-CNN, we observe huge 3.3% gains in the overall metric AP. Our benefits are particularly evident in the rare class performance AP_r, with gains of 9.0%. This validates FASA’s superior ability to deal with long-tailed tasks. We also verify that the benefits of FASA remain consistent over multiple runs. When combined with stronger methods (RFS/EQL/cRT/BAGS/Seesaw) or training schedules, FASA can still obtain consistent improvements in overall AP, where the rare-class improvements in AP_r dominate. Such gains come without compromising the common- and




Figure 3: **Visualization of the class-wise sampling probabilities during training.** (a) We show the average sampling probabilities for the rare/common/frequent class groups. (b) The changing behavior of class-wise sampling probabilities within each group.

Table 4: Comparing state-of-the-art methods with and without our FASA on the LVIS v1.0 *validation* dataset. We compare with the Mask R-CNN baseline, and state-of-the-art re-sampling approach Repeat Factor Sampling (RFS) [13], Equalization Loss (EQL) [34], Classifier Re-training (cRT) [19], Balanced Group Softmax (BAGS) [23] and Seesaw Loss [37]. The ‘Uniform’ method indicates random and uniform sample images. These methods are trained under different training schedules (24 or 12+12 epochs) using the public codes. All methods use ResNet-50 [15] as the backbone for fair comparison.

Loss	Sampler	#Epochs	FASA	AP	AP _r	AP _c	AP _f
Softmax CE	Uniform	24	✗	19.3	1.2	17.4	29.3
			✓	22.6 (+3.3)	10.2 (+9.0)	21.6	29.2
	RFS, CVPR’19 [13]	24	✗	22.8	12.9	21.6	28.3
			✓	24.1 (+1.3)	17.3 (+4.4)	22.9	28.5
EQL, CVPR’20 [34]	Uniform	24	✗	22.1	5.1	22.4	29.3
			✓	24.4 (+2.3)	15.4 (+10.3)	23.5	29.4
cRT, ICLR’20 [19]	Uniform / RFS, CVPR’19 [13]	12+12	✗	22.4	12.2	20.4	29.1
			✓	23.6 (+1.2)	15.1 (+2.9)	22.0	29.1
BAGS’ CVPR’20 [23]	Uniform / RFS, CVPR’19 [13]	12+12	✗	22.8	12.4	22.2	28.3
			✓	24.0 (+1.2)	15.2 (+2.8)	23.4	28.3
Seesaw, CVPR’21 [37]	RFS, CVPR’19 [13]	24	✗	26.4	19.6	26.1	29.8
			✓	27.5 (+1.1)	21.0 (+1.4)	27.5	30.1

frequent-class metrics AP_c and AP_f, where FASA performs better or remains on par.

Long training schedules and large backbones. FASA works well even with large backbones (ResNet 101 [15], ResNeXt 101-32-8d [43]) and advanced instance segmentation framework Cascade Mask R-CNN [1]. As shown in Table 5, the benefits of FASA still hold under different settings. On the validation set, FASA improves the performance of Mask R-CNN baseline (ResNet-101 backbone) by a large margin. We observe improvements in overall AP and rare-class AP_r by 1.9% and 5.9% respectively. For the stronger ResNeXt101 backbone, we observe similar trend. FASA significantly improves the rare class performance AP_r by 4.6%. For Cascade Mask R-CNN [1] framework and Seesaw loss [37], FASA offers consistent performance boost in rare classes while remaining strong on both common and frequent classes.

Analyzing classifier weight norm. As discussed in [19, 23], the weight norm of a classifier is correlated with the imbalanced learning performance, with the weight norms of tail classes often being much smaller on long-tailed data. Fig. 4 visualizes the weight norm of classifiers trained with and without FASA. It can be seen that FASA leads to more balanced class distributions of weight norms than Cross-




Figure 4: Comparison of the classifier weight norm of Cross-Entropy and Seesaw Loss [37] (with and without FASA). Our FASA leads to a more class-balanced weight norm distribution.

Entropy baseline and Seesaw Loss [37]. This partially explains why FASA greatly improves the rare-class performance of these methods.

4.3. Evaluation on COCO-LT

We evaluated FASA on COCO-LT [38] dataset to examine the generalizability of our approach. For a fair comparison, we follow the same experiment setting as SimCal [38]. There are two main differences compared to the implemen-

Table 5: Comparison state-of-the-art methods with and without FASA, using large backbones (ResNet-101 [15], ResNeXt-101-32x8d [43]) and advanced Cascade Mask R-CNN [1] framework.

Method	Loss	Sampler	Backbone	FASA	AP	AP _r	AP _c	AP _f
Mask R-CNN, ICCV'17 [14]	Softmax CE	RFS, CVPR'19 [13]	R101	\times ✓	24.4 26.3 (+1.9)	13.2 19.1 (+5.9)	24.7 25.4	30.3 30.6
Mask R-CNN, ICCV'17 [14]	Softmax CE	RFS, CVPR'19 [13]	X101	\times ✓	26.1 27.7 (+1.6)	16.1 20.7 (+4.6)	24.9 26.6	32.0 32.0
Cascade Mask R-CNN, TPAMI'19 [1]	Softmax CE	RFS, CVPR'19 [13]	R101	\times ✓	25.4 27.7 (+2.3)	13.7 19.8 (+5.9)	24.8 27.3	31.4 31.6
Cascade Mask R-CNN, TPAMI'19 [1]	Seesaw, CVPR'21 [37]	RFS, CVPR'19 [13]	R101	\times ✓	30.1 31.5 (+1.4)	21.4 24.1 (+2.7)	30.0 31.9	33.9 34.0

Table 6: Results on COCO-LT [38] *minimal* set. AP^m and AP^b denote the Mask mAP and Bbox mAP, respectively. AP₁^m, AP₂^m, AP₃^m, AP₄^m refer to bin of [1, 20), [20, 400), [400, 8000), [8000, -) training instances.

Method	AP ^m	AP ₁ ^m	AP ₂ ^m	AP ₃ ^m	AP ₄ ^m	AP ^b
Mask R-CNN	18.7	0.0	8.2	24.4	26.0	21.4
SimCal [38]	21.8	15.0	16.2	24.3	26.0	24.6
FASA (Ours)	23.4	13.5	19.0	25.2	27.5	26.0

Table 7: Comparison of classification accuracy on CIFAR-LT-100 [2] dataset with imbalance ratio 100. † denotes that the results are copied from [20].

Backbone	Method	Augmentation	Accuracy (%)
ResNet 32	LDAM [2] †	\times	42.0 ± 0.09
	LDAM [2]	M2M [20]	43.5 ± 0.22
	LDAM [2]	FASA(Ours)	43.7 ± 0.25
	De-confound-TDE [35]	\times	44.1
	De-confound-TDE [35]	FASA (Ours)	45.2 ± 0.16
ResNet 34	Cross entropy	Chu <i>et al.</i> [6]	48.5
	Cross entropy	FASA(Ours)	49.1 ± 0.23

tation details we described in the LVIS dataset: (1) Single-scale training is used. We set the input image’s short side to 800 pixels and do not use multi-scale jittering. (2) The bounding box head and mask head are class agnostic.

Table 6 summarizes the results. The top row shows results of the Mask R-CNN [14] with ResNet50 [15] backbone and FPN [24] neck. The second row shows results obtained by the previous state-of-the-art method SimCal [38], which involves classification calibration training and then dual head inference. The third row shows results of the Mask R-CNN baseline augmented with our FASA. Note how FASA improves the performance of Mask R-CNN considerably, especially for those rare classes (AP₁^m and AP₂^m). Overall, FASA performs better than SimCal [38] without any form of decoupled training.

4.4. Evaluation on CIFAR-LT-100

To demonstrate the generalizability of FASA in other domains, we further apply FASA on the long-tailed image classification task. We follow [2] to conduct experiments on the long-tailed CIFAR-100 dataset with a sub-

stantial imbalance ratio of $\rho = 100$ (the ratio between sample sizes of the most frequent and least frequent class, $\rho = \max_i \{N_i\} / \min_i \{N_i\}$). We train for 200 epochs and use the top-1 accuracy as the evaluation metric. We compare FASA with two state-of-the-art feature-level augmentation approaches M2M [20] and Chu *et al.* [6]. When comparing with M2M [20], we use LDAM [2] loss as the baseline for a fair comparison. We use the same backbones of [20, 6] to experiment under the same settings.

From Table 7, we observe that our FASA brings 1.7% accuracy improvement over the LDAM [2] baseline and is comparable with M2M [20]. When applied to the stronger baseline De-confound-TDE [35], FASA’s benefits still hold. FASA is also better performing than Chu *et al.* [6] with the cross entropy loss. Other than that, our FASA is more time-efficient than M2M [20] and Chu *et al.* [6] because they need extra time for classifier pre-training while FASA can be applied online.

5. Conclusion

We have presented a simple yet effective method, Feature Augmentation and Sampling Adaptive (FASA), to address the class imbalance issue in the long-tailed instance segmentation task. FASA generates virtual features on the fly to provide more positive samples for rare categories, and leverages a loss-guided adaptive sampling scheme to avoid over-fitting. FASA achieves consistent gains on the Mask R-CNN baseline under different backbones, learning schedules, data samplers and loss functions, with minimal impact to the training efficiency. It shows large gains for rare classes without compromising the performance for common and frequent classes. Notably, as an orthogonal component, it improves other more recent methods such as RFS, EQL, cRT, BAGS and Seesaw. Compelling results are shown on two challenging instance segmentation datasets, LVIS v1.0 and COCO-LT, as well as an imbalanced image classification benchmark, CIFAR-LT-100.

Acknowledgements. This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF- ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- [1] Zhaowei Cai and Nuno Vasconcelos. Cascade R-CNN: High quality object detection and instance segmentation. *TPAMI*, 2019. 1, 7, 8
- [2] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arachiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *NIPS*, 2019. 2, 8
- [3] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. SMOTE: Synthetic minority over-sampling technique. *JAIR*, 16:321–357, 2002. 2, 3, 5, 6
- [4] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiao-xiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, et al. Hybrid task cascade for instance segmentation. In *CVPR*, 2019. 1
- [5] 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:1906.07155*, 2019. 5, 6
- [6] Peng Chu, Xiao Bian, Shaopeng Liu, and Haibin Ling. Feature space augmentation for long-tailed data. In *ECCV*, 2020. 1, 3, 5, 6, 8
- [7] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *CVPR*, 2019. 2
- [8] Achal Dave, Piotr Dollár, Deva Ramanan, Alexander Kirillov, and Ross Girshick. Evaluating large-vocabulary object detectors: The devil is in the details. *arXiv preprint arXiv:2102.01066*, 2021. 5
- [9] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *CVPR*, pages 4690–4699, 2019. 6
- [10] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, volume 96, pages 226–231, 1996. 5
- [11] Hao-Shu Fang, Jianhua Sun, Runzhong Wang, Minghao Gou, Yong-Lu Li, and Cewu Lu. Instaboost: Boosting instance segmentation via probability map guided copy-pasting. In *CVPR*, pages 682–691, 2019. 3, 5, 6
- [12] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *CVPR*, 2021. 3
- [13] Agrim Gupta, Piotr Dollar, and Ross Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 1, 2, 3, 4, 5, 6, 7, 8
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017. 1, 2, 3, 5, 8
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 7, 8
- [16] Xinting Hu, Yi Jiang, Kaihua Tang, Jingyuan Chen, Chunyan Miao, and Hanwang Zhang. Learning to segment the tail. In *CVPR*, pages 14045–14054, 2020. 2
- [17] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Deep imbalanced learning for face recognition and attribute prediction. *TPAMI*, 2019. 1
- [18] Chen Huang, Shuangfei Zhai, Walter Talbott, Miguel Ángel Bautista, Shih-Yu Sun, Carlos Guestrin, and Joshua M. Susskind. Addressing the loss-metric mismatch with adaptive loss alignment. In *ICML*, 2019. 4
- [19] 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, 6, 7
- [20] Jaehyung Kim, Jongheon Jeong, and Jinwoo Shin. M2m: Imbalanced classification via major-to-minor translation. In *CVPR*, pages 13896–13905, 2020. 1, 3, 8
- [21] Boyi Li, Felix Wu, Ser-Nam Lim, Serge Belongie, and Kilian Q Weinberger. On feature normalization and data augmentation. In *CVPR*, 2021. 3, 5, 6
- [22] Shuang Li, Kaixiong Gong, Chi Harold Liu, Yulin Wang, Feng Qiao, and Xinjing Cheng. Metasaug: Meta semantic augmentation for long-tailed visual recognition. In *CVPR*, pages 5212–5221, 2021. 3
- [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*, pages 10991–11000, 2020. 1, 2, 5, 6, 7
- [24] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 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 *ECCV*, 2014. 1
- [26] Jialun Liu, Yifan Sun, Chuchu Han, Zhaopeng Dou, and Wenhui Li. Deep representation learning on long-tailed data: A learnable embedding augmentation perspective. In *CVPR*, pages 2970–2979, 2020. 1, 3, 5, 6
- [27] Jialun Liu, Jingwei Zhang, Wenhui Li, Chi Zhang, and Yifan Sun. Memory-based jitter: Improving visual recognition on long-tailed data with diversity in memory. *arXiv preprint arXiv:2008.09809*, 2020. 3
- [28] Lanlan Liu, Michael Muelly, Jia Deng, Tomas Pfister, and Li-Jia Li. Generative modeling for small-data object detection. In *ICCV*, pages 6073–6081, 2019. 3
- [29] 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*, pages 2537–2546, 2019. 1
- [30] Jiawei Ren, Cunjun Yu, Shunan Sheng, Xiao Ma, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Balanced meta-softmax for long-tailed visual recognition. In *NIPS*, 2020. 2
- [31] Li Shen, Zhouchen Lin, and Qingming Huang. Relay back-propagation for effective learning of deep convolutional neural networks. In *ECCV*, pages 467–482. Springer, 2016. 2

- [32] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, pages 761–769, 2016. 4
- [33] Jingru Tan, Xin Lu, Gang Zhang, Changqing Yin, and Quanquan Li. Equalization loss v2: A new gradient balance approach for long-tailed object detection. In *CVPR*, 2021. 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 *CVPR*, pages 11662–11671, 2020. 1, 2, 5, 6, 7
- [35] Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. Long-tailed classification by keeping the good and removing the bad momentum causal effect. In *NIPS*, 2020. 2, 8
- [36] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In *ICML*, pages 6438–6447. PMLR, 2019. 3
- [37] Jiaqi Wang, Wenwei Zhang, Yuhang Zang, Yuhang Cao, Jiangmiao Pang, Tao Gong, Kai Chen, Ziwei Liu, Chen Change Loy, and Dahua Lin. Seesaw loss for long-tailed instance segmentation. In *CVPR*, 2021. 2, 6, 7, 8
- [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 *ECCV*, 2020. 2, 5, 7, 8
- [39] Tong Wang, Yousong Zhu, Chaoyang Zhao, Wei Zeng, Jinqiao Wang, and Ming Tang. Adaptive class suppression loss for long-tail object detection. In *CVPR*, pages 3103–3112, 2021. 2
- [40] Yulin Wang, Xuran Pan, Shiji Song, Hong Zhang, Gao Huang, and Cheng Wu. Implicit semantic data augmentation for deep networks. *NeurIPS*, 32:12635–12644, 2019. 3
- [41] Jialian Wu, Liangchen Song, Tiancai Wang, Qian Zhang, and Junsong Yuan. Forest R-CNN: Large-vocabulary long-tailed object detection and instance segmentation. In *ACM MM*, pages 1570–1578, 2020. 1, 2, 5
- [42] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *CVPR*, pages 10687–10698, 2020. 3
- [43] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017. 7, 8
- [44] Xi Yin, Xiang Yu, Kihyuk Sohn, Xiaoming Liu, and Manmohan Chandraker. Feature transfer learning for face recognition with under-represented data. In *CVPR*, 2019. 1, 3, 5, 6
- [45] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, pages 6023–6032, 2019. 3
- [46] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. 3
- [47] Songyang Zhang, Zeming Li, Shipeng Yan, Xuming He, and Jian Sun. Distribution alignment: A unified framework for long-tail visual recognition. In *CVPR*, pages 2361–2370, 2021. 2
- [48] Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. BBN: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In *CVPR*, pages 9719–9728, 2020. 2
- [49] Xingyi Zhou, Koltun Vladlen, and Philipp Krähenbühl. Joint COCO and LVIS workshop at ECCV 2020: LVIS challenge track technical report: CenterNet2, 2020. 2