

DropLoss for Long-Tail Instance Segmentation

Ting-I Hsieh^{1*}, Esther Robb^{2*}, Hwann-Tzong Chen^{1,3}, Jia-Bin Huang²

¹ National Tsing Hua University ² Virginia Tech ³ Aeolus Robotics

Abstract

Long-tailed class distributions are prevalent among the practical applications of object detection and instance segmentation. Prior work in long-tail instance segmentation addresses the imbalance of losses between rare and frequent categories by reducing the penalty for a model incorrectly predicting a rare class label. We demonstrate that the rare categories are heavily suppressed by *correct background predictions*, which reduces the probability for all foreground categories with equal weight. Due to the relative infrequency of rare categories, this leads to an imbalance that biases towards predicting more frequent categories. Based on this insight, we develop DropLoss – a novel adaptive loss to compensate for this imbalance without a trade-off between rare and frequent categories. With this loss, we show state-of-the-art mAP across rare, common, and frequent categories on the LVIS dataset. Codes are available at <https://github.com/timy90022/DropLoss>.


Introduction

Object detection and instance segmentation have a wide array of practical applications. State-of-the-art object detection methods adopt a multistage framework (Girshick et al. 2014; Girshick 2015; Ren et al. 2015) trained on large-scale datasets with abundant examples for each object category (Lin et al. 2014). However, datasets used in real-world applications commonly fall into a long-tailed distribution over categories, i.e., the majority of classes have only a small number of training examples. Training a model on these datasets inevitably induces an undesired bias towards frequent categories. The limited diversity of rare-category samples further increases the risk of overfitting. Methods for addressing the issues involving long-tailed distributions commonly fall into several groups: *i*) resampling to balance the category frequencies, *ii*) reweighting the losses of rare and frequent categories, and *iii*) specialized architectures or feature transformations.

The instance segmentation problem presents unique challenges for learning long-tailed distributions, as it contains multiple training objectives to supervise region proposal, bounding box regression, mask regression, and object classification. Each of these losses contributes to the overall

*equal contribution

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



(a) Gradient percentage

(b) Prediction scores

Figure 1: Motivation. (a) Percentage of gradient updates from incorrect foreground classification (blue) and ground-truth background anchors (orange) on LVIS (Gupta, Dollár, and Girshick 2019). We divide the categories into “frequent” (white shading), “common” (orange shading), and “rare” (yellow shading). For rare categories, background gradients occupy a disproportionate percentage of total gradients. (b) The distribution of average foreground class prediction scores for ground-truth background bounding boxes at earlier (red) and later (blue) training stages. We find that, for background bounding boxes, the prediction scores of rare categories are more severely suppressed, and the training is biased towards predicting more frequent categories.

balance of model training. The prior state-of-the-art in long-tail instance segmentation (Tan et al. 2020) discovered a phenomenon where the predictions for rare categories are suppressed by *incorrect foreground class predictions*. To reduce these “discouraging gradients” and allow the network to explore the solution space for rare categories, the EQL method (Tan et al. 2020) removes losses to rare categories from incorrect foreground classification. However, we observe that most “discouraging gradients” in fact originate from *correct background classification* (where a bounding box does not contain any labeled objects). In the background case, the classification branch receives losses to suppress all foreground class prediction scores.

In Figure 1, we study the effect of such *discouraging gradients* on the different categories of a long-tail dataset, categorized by number of training images into *rare* (1-10 images), *common* (11-100), and *frequent* (> 100) categories. We find that these losses disproportionately affect rare and common categories, due to the infrequency of “encouraging gradients”

in which a bounding box contains the correct category label. Specifically, Figure 1(a) shows that 50-70% of discouraging gradients for rare categories originate from background predictions, compared with only 30-40% of discouraging gradients for frequent categories. Discouraging gradients from background classification (orange curve) contribute a much higher percentage of total discouraging gradients compared to that of incorrect foreground prediction (blue curve) as used in EQL (Tan et al. 2020). Figure 1(b) shows that using a ground-truth background anchor, a trained model predicts scores for rare categories with several orders-of-magnitude lower confidence than for frequent categories. This demonstrates a bias towards predicting more frequent categories.

Based on these observations, we develop a simple yet effective method to *adaptively* rebalance the ratio of background prediction losses between rare/common and frequent categories. Our proposed method *DropLoss* removes losses for rare and common categories from background predictions based on sampling a Bernoulli variable with parameters determined by batch statistics. *DropLoss* prevents suppression of rare and common categories, increasing opportunities for correct predictions of infrequent classes during training and reducing frequent class bias.

The contributions of this work are summarized as follows:

1. We provide an analysis of the unique characteristics of long-tailed distributions, particularly in the context of instance segmentation, to pinpoint the imbalance problem caused by disproportionate discouraging gradients from background predictions during training.
2. We develop a methodology for alleviating imbalances in the long-tailed setting by leveraging the ratio of rare and frequent classes in a sampled training batch.
3. We present state-of-the-art instance segmentation results on the challenging long-tail LVIS dataset (Gupta, Dollár, and Girshick 2019).

Related Work

Object Detection and Instance Segmentation. Two-stage detection architectures (Girshick et al. 2014; Girshick 2015; Ren et al. 2015; Lin et al. 2017a) have been successful in the object detection setting, where the first stage proposes a “region of interest” and the second stage refines the bounding box and performs classification. This decomposition was initially proposed in R-CNN (Girshick et al. 2014). Fast R-CNN (Girshick 2015) and Faster R-CNN (Ren et al. 2015) improve efficiency and quality for object detection. Mask R-CNN later adapts Faster R-CNN to the instance segmentation setting by adding a mask prediction branch in the second stage (He et al. 2017). Mask R-CNN has proven effective in a wide variety of instance segmentation tasks. Our work adopts this architecture. In contrast with two-stage methods, single-stage methods provide faster inference by eliminating the region proposal stage and instead predicting a bounding box directly from anchors (Liu et al. 2016; Redmon et al. 2016; Lin et al. 2017c). However, two-stage architectures generally provide better localization.

Learning Long-tailed Distributions. Techniques for learning long-tailed distributions generally fall into three groups: resampling, reweighting and cost-sensitive learning, and feature manipulation. We discuss each in the following sections.

Resampling Methods. Oversampling methods (Chawla et al. 2002; Han, Wang, and Mao 2005; Mahajan et al. 2018; Hensman and Masko 2015; Huang et al. 2016; He et al. 2008; Zou et al. 2018) duplicate rare class samples to balance out the class frequency distribution. However, oversampling methods tend to overfit to the rare categories, as this type of method does not address the fundamental lack of data. Several oversampling methods aim to address this by augmenting the available data (Chawla et al. 2002; Han, Wang, and Mao 2005), but undersampling methods are often preferred (Drummond, Holte et al. 2003). Undersampling methods (Drummond, Holte et al. 2003; Tsai et al. 2019; Kahn and Marshall 1953) remove frequent class samples from the dataset to balance the class frequency distribution. The loss of information from removing these samples can be mitigated through careful selection using statistical techniques (Tsai et al. 2019; Kahn and Marshall 1953). It can be beneficial to combine the advantages of undersampling and oversampling (Chawla et al. 2002). Dynamic methods adjust the sampling distribution throughout training based on loss or metrics (Pouyanfar et al. 2018). Class balance sampling (Kang et al. 2019; Shen, Lin, and Huang 2016) uses class-aware strategies to rebalance the data distribution for learning classifiers and representations. In the context of the dense instance segmentation problem, it is difficult to apply the above resampling methods because the number of class examples per image may vary.

Reweighting and Cost-sensitive Methods. Rather than rebalancing the sampling distribution, reweighting methods seek to balance the *loss weighting* between rare and frequent categories. Class frequency reweighting methods commonly use the inverse frequency of each class to weight the loss (Huang et al. 2016; Wang, Ramanan, and Hebert 2017; Cui et al. 2019). Cost-sensitive methods (Li, Liu, and Wang 2019; Lin et al. 2017d) aim to balance the model loss magnitudes between rare and frequent categories. An existing meta-learning method (Shu et al. 2019) explicitly learns loss weights based on the data. Our method provides a simple way to combine class frequency-aware sampling and cost-sensitive learning.

Feature Manipulation Methods. In contrast to resampling methods and reweighting methods that focus on modifying the loss based on class frequency, feature manipulation methods aim to design specific architectures or feature relationships to address the long-tail problem. Normalization can be used to control the distribution of deep features, preventing frequent categories from dominating training (Kang et al. 2019). Metric learning methods (Kang et al. 2019; Zhang et al. 2017) learn maximally-distant “prototypes” of deep features to improve performance on data-scarce categories, effectively transferring knowledge between head and tail categories. Similarly, knowledge transfer in feature space can be accomplished using memory-based prototypes (Liu et al. 2019) or transfer of intra-class variance (Yin et al. 2019).


Long-tail Learning Settings. Several methods have been proposed to handle the problem of learning from imbalanced datasets in other settings such as object classification (Cui et al. 2019; Cao et al. 2019; Jamal et al. 2020; Tan et al. 2020). In the long-tail object recognition setting, the prior state-of-the-art method (Tan et al. 2020) uses selective reweighting. Their work observed that rare categories receive significantly more “discouraging gradients” compared with frequent categories, and develop a method for rebalancing discouraging gradients from foreground misclassifications. Their method uses a binary 0 or 1 reweighting based on whether the class is rare or frequent. Unlike this work, our method focuses on the much more prevalent background classification losses in the instance segmentation setting, and we develop a new adaptive resampling and reweighting method which accounts for this imbalance and for the distribution of classes within a sample. Note that most of the methods for learning long-tailed distributions focus on the *image classification* setting where this is no background class (and therefore no losses associated with background). In object detection and instance segmentation, however, the background class plays a very dominant role in the loss. This inspires our design of a reweighting mechanism which specifically considers background class.

Method

Based on the observation that rare and common categories receive disproportionate discouraging gradients from background classifications (compared with frequent categories), the goal of our method is to prevent rare categories from being overly suppressed by reducing the imbalance of discouraging gradients. We are inspired by work in one-stage object detection (Li, Liu, and Wang 2019; Lin et al. 2017d), which encounters a similar problem of large gradients from negative anchors inhibiting learning. We first construct a baseline which modifies the equalization loss (Tan et al. 2020) to rebalance gradients from foreground and *background* region proposals for rare/common and frequent categories. We show that this baseline leads to improved results over (Tan et al. 2020) but requires careful hyperparameter selection and exhibits a clear tradeoff between rare/common and frequent categories. To alleviate this problem, we propose a stochastic *DropLoss* which improves the overall frequent-rare category performance as well as improving the tradeoff as measured by a Pareto frontier 3.

Revisiting the Equalization Loss

Equalization Loss. We start with a review of the equalization loss (Tan et al. 2020). Equalization loss modifies sigmoid cross-entropy to alleviate discouraging gradients from incorrect foreground predictions. Note that, in sigmoid cross-entropy, the ground-truth label y_j represents only a binary distribution for the foreground category j , and no extra class label for the background is included. That is, we have $y_j = 1$ if the ground-truth category of a region is j . On the other hand, if a region belongs to the background, we have $y_j = 0$ for all the categories. During training, a region proposal is labeled as “background” if its IoU with any ground-truth region of a foreground class is lower than 50%.



(a) Weight w_j as a function of the logarithm base b . (b) Frequent-rare category performance tradeoff

Figure 2: Background equalization loss. We present an extension of the equalization loss that specifically focuses on the background classification. (a) The curves of the $\mathcal{L}_{\text{BEQL}}$ weights in (4) with different choices for the logarithm base b . Smaller values of the logarithm base b reduce the effects of background more. (b) The experimental result of applying different logarithm bases shows a tradeoff between the mean average precision (mAP) of rare categories and the mAP of frequent categories with respect to different logarithm base settings. Note that the background equalization loss $\mathcal{L}_{\text{BEQL}}$ with a large base b reduces to the existing equalization loss \mathcal{L}_{EQL} .

Given a region proposal r , the equalization loss is formulated as follows:

$$\mathcal{L}_{\text{EQL}} = - \sum_{j=1}^C w_j \log(\hat{p}_j), \quad \hat{p}_j = \begin{cases} p_j, & \text{if } y_j = 1, \\ 1 - p_j, & \text{otherwise,} \end{cases} \quad (1)$$

$$w_j = 1 - E(r)T_\lambda(f_j)(1 - y_j), \quad (2)$$

where C is the number of categories, p_j is predicted logit, and f_j is the frequency of category j in the dataset. The indicator function $E(r)$ outputs 1 if r is a foreground region and 0 if it belongs to the background. More specifically, for a region proposal r that is considered a background region with all y_j being zero, we have $E(r) = 0$. $T_\lambda(f)$ is also a binary indicator function that, given a threshold λ , outputs 1 if $f < \lambda$ to indicate the category is of low frequency. It can be verified from (2) that, for a foreground region r (*i.e.*, $E(r) = 1$), the weight w_j is either 1 or $1 - T_\lambda(f_j)$, depending on the frequency of the ground-truth category j . Further, if a category j is of low frequency (rare category), the weight $w_j = 1 - T_\lambda(f_j)$ becomes zero and thus no penalty is given to incorrect foreground predictions. On the other hand, for frequent categories, the weight is 1 and the penalty of incorrect prediction remains $-\log(1 - p_j)$. By removing discouraging gradients to rare/common categories from incorrect foreground predictions, the equalization loss achieved state-of-the-art on the LVIS Challenge 2019. Foreground class label prediction is selected using the maximum logit, so entirely removing the loss allows the network to optimize rare categories without penalties, as long as the prediction logit is less than ground truth for the frequent categories. This approach removes large penalties for non-zero confidences in rare categories, which otherwise imbalance the training to suppress rare categories.

Background Equalization Loss. In contrast to the mechanism of equalization loss that prevents large penalties for non-zero confidences in rare categories, cost-sensitive learning methods only reduce (Lin et al. 2017d) or remove (Li, Liu, and Wang 2019) discouraging gradients if the magnitude of the loss falls below some threshold. Our core insight is that foreground and background categories require different approaches due to the differences in prediction criteria. For background categories, the network predicts background class if *all* logits p_j fall below a threshold. For foreground categories, the prediction is selected using the *maximum* logit p_j . Inspired by cost-sensitive loss and equalization loss, we present the *background equalization loss* as an extension to the original equalization loss:

$$\mathcal{L}_{\text{BEQL}} = - \sum_{j=1}^C w_j \log(\hat{p}_j), \quad \hat{p}_j = \begin{cases} p_j, & \text{if } y_j = 1, \\ 1 - p_j, & \text{otherwise,} \end{cases} \quad (3)$$

$$w_j = \begin{cases} 1 - T_\lambda(f_j)(1 - y_j), & \text{if } E(r) = 1, \\ 1 - T_\lambda(f_j) \cdot \min\{-\log_b(p_j), 1\}, & \text{otherwise.} \end{cases} \quad (4)$$

By comparing (2) and (4), we can see that the background equalization loss differs from the equalization loss in the weights for background regions. The equalization loss always penalizes a background region ($E(r) = 0$ and thus $w_j = 1$) even if the category is of low frequency. In contrast, our background equalization loss gives smaller weight to background predictions as long as their confidences are low. We use a logarithm base b to control the sensitivity of the weight concerning the confidence of background prediction. Figure 2(a) shows the curves of the $\mathcal{L}_{\text{BEQL}}$ weights in (4) by varying the value of the logarithm base b . For example, suppose we would want to focus on the performance of the rare category, we can set the value of $b = 2$. The main idea here is to alleviate the accumulation of small but non-negligible discouraging gradients from the background. When applying the proposed background equalization loss with different logarithm bases, however, we see a clear performance tradeoff between frequent and rare categories (see Figure 2(b)). The results show that the average precision of the rare categories behaves in the *opposite* way as the average precision of the frequent categories for different choices of logarithm bases.

DropLoss

While suppressing discouraging gradients from the background shows improvement for the rare categories, the background equalization loss has a drawback. The performance often sensitively depends on the choice of the logarithm base. It is difficult to choose an appropriate logarithm base that works for different long-tailed distributions without suffering from a tradeoff between frequent and rare categories. In light of this, we propose a new stochastic method, called *DropLoss*, which dynamically balances the influence of background discouraging gradients for rare/common/frequent categories.

Similar to the design of the background equalization loss, we seek to adjust weights on the logits of low-frequency categories for background region proposals. In DropLoss, we introduce a Bernoulli distribution and sample a binary value

from the distribution as the weight w_j if a region belongs to the background. Further, we determine the parameter of the Bernoulli distribution by a beta sampling distribution over the occurrence ratios of rare, common, and frequent categories for the regions generated by the Region Proposal Network (Ren et al. 2015) during training. The Bernoulli distribution with a Beta prior is suitable because we aim to model binary outcomes with varying biases in a stochastic manner.

Given a batch of region proposals, we compute the ratio between the occurrences of ‘rare + common’ categories to all foreground occurrences (*i.e.*, ‘rare + common + frequent’ categories). In other words, we treat a batch of region proposals as a sample of occurrence ratio that is drawn from a beta distribution to provide the parameter of the Bernoulli distribution. Our intuition behind such a scheme is simple: For region proposals of rare and common categories, their occurrences in a batch are of low frequency. Therefore, the discouraging gradients from the background predictions should be accordingly discounted for rare and common categories.

We formulate DropLoss as follows:

$$\mathcal{L}_{\text{Drop}} = - \sum_{j=1}^C w_j \log(\hat{p}_j), \quad \hat{p}_j = \begin{cases} p_j, & \text{if } y_j = 1, \\ 1 - p_j, & \text{otherwise,} \end{cases} \quad (5)$$

$$w_j = \begin{cases} 1 - T_\lambda(f_j)(1 - y_j), & \text{if } E(r) = 1, \\ w \sim \text{Ber}(\mu_{f_j}), & \text{otherwise,} \end{cases} \quad (6)$$

where a random sample $w \in \{0, 1\}$ is drawn from Bernoulli distribution $\text{Ber}(\mu_{f_j})$ if the region proposal r belongs to the background, *i.e.*, $E(r) = 0$. The parameter μ_{f_j} of the Bernoulli distribution is determined by the occurrence ratio of low-frequency (‘rare + common’) categories in the current batch of region proposals. We compute the parameter by

$$\mu_{f_j} = \begin{cases} (n_{\text{rare}} + n_{\text{common}})/n_{\text{all}}, & \text{if } T_\lambda(f_j) = 1, \\ n_{\text{frequent}}/n_{\text{all}}, & \text{otherwise,} \end{cases} \quad (7)$$

where n_{rare} , n_{common} , and n_{frequent} are the numbers of occurrences of rare, common, and frequent categories in the current training batch of foreground region proposals. The total number of foreground occurrences is $n_{\text{all}} = n_{\text{rare}} + n_{\text{common}} + n_{\text{frequent}}$. Implementation of the above DropLoss scheme is straightforward: For each batch, we derive the parameter μ_{f_j} depending on whether category j is rare/common or frequent. We can then simulate a flip of a biased coin with head probability μ_{f_j} and assign $w_j = 1$ if we get a head.

A region proposal is annotated as a background region if it does not overlap with any ground-truth foreground region, or if the IoU is lower than 50%. If the number of rare category occurrences in a given batch is large, discouraging gradients to that rare category are more likely to be kept (with a higher chance to get $w_j = 1$). On the other hand, if a rare category does not appear very often in a batch, it is highly probable that discouraging gradients to the rare category will be dropped. Therefore, our dropping strategy tends to neglect unrelated non-overlapping background proposals but would be inclined to keep more related ($0 < \text{IoU} < 0.5$) background proposals.

Architecture	Backbone	Loss	AP (%)	AP ₅₀	AP ₇₅	AP _r	AP _c	AP _f	AR	AP _{bbox}
Mask R-CNN	R-50-FPN	BCE	21.5	33.4	22.9	4.7	21.2	28.6	28.3	21
		EQL (Tan et al. 2020)	23.8	36.3	25.2	8.5	25.2	28.3	31.5	23.5
		DropLoss (Ours)	25.5	38.7	27.2	13.2	27.9	27.3	34.8	25.1
Mask R-CNN	R-101-FPN	BCE	23.6	36.5	25.1	5.6	24.2	30.1	30.9	23.3
		EQL (Tan et al. 2020)	26.2	39.5	27.9	11.9	27.8	29.8	33.8	26.2
		DropLoss (Ours)	26.9	40.6	28.9	14.8	29.7	28.3	36.4	26.8
Cascade R-CNN	R-50-FPN	BCE	21.4	32	23.1	3.4	20.4	29.8	27.6	22.8
		EQL (Tan et al. 2020)	24.2	35.9	25.8	7.8	25	29.7	31.4	26
		DropLoss (Ours)	25	37	26.9	9.1	27.2	28.7	34	26.9
Cascade R-CNN	R-101-FPN	BCE	23	34.4	24.7	3.5	22.8	31.2	29.9	24.9
		EQL (Tan et al. 2020)	25.4	37.3	27.3	7.2	26.6	31	33.1	27.2
		DropLoss (Ours)	26.4	39	28.1	11.5	28.5	29.7	35.5	28.6

Table 1: Comparison between architecture and backbone settings, evaluated on LVIS v0.5 validation set. We compare BCE (binary cross-entropy), EQL (equalization loss) and Drop (DropLoss). AP/AR refers to mask AP/AR, and subscripts ‘r’, ‘c’, and ‘f’ refer to rare, common, and frequent categories.

Experimental Results

In this section, we present the implementation details and experimental results. We compare DropLoss with the state-of-the-art long-tail instance segmentation baselines on the challenging LVIS dataset (Gupta, Dollár, and Girshick 2019). To validate the effectiveness of this approach, we compare across different architectures and backbones and integrate with additional long-tail resampling methods. We find that DropLoss demonstrates consistently improved results in AP and AR across all these experimental settings.

Dataset. Following the previous work *equalization loss* (Tan et al. 2020), we train and evaluate our model on LVIS benchmark dataset. LVIS is a large vocabulary instance segmentation dataset, containing 1,230 categories. In LVIS dataset, categories are sorted into three groups based on the number of images in which they appear: *rare* (1-10 images), *common* (11-100), and *frequent* (> 100). We report AP for each bin to quantify performance in the long-tailed distribution setting. We train our model on the 57K-image LVIS v0.5 training set and evaluate it on the 5K-image LVIS v0.5 validation set.

Implementation Details. For our experiments, we adopt the Mask R-CNN (He et al. 2020) architecture with Feature Pyramid Networks (Lin et al. 2017b) as a baseline model. We train the network using stochastic gradient descent with a momentum of 0.9 and a weight decay of 0.0001 for 90K iterations, with batch size 16 on eight parallel NVIDIA 2080 Ti GPUs. We initialize the learning rate to 0.2 and decay it by a ratio of 0.1 at iterations 60,000 and 80,000. We use the Detectron2 (Wu et al. 2019) framework with default data augmentation. The data augmentation includes scale jitter with a short edge of (640, 672, 704, 736, 768, 800) pixels and a long edge no more than 1,333 pixels horizontal flipping. In the Region Proposal Network (RPN), we sample 256 anchors with a 1:1 ratio between foreground and background to compute the RPN loss and choose 512 ROI-aligned proposals per image with a 1:3 foreground-background ratio for later pre-




Figure 3: Measuring the performance tradeoff. Comparison between rare, common, and frequent categories AP for baselines and our method. We visualize the tradeoff for ‘common vs. frequent’ and ‘rare vs. frequent’ as a Pareto frontier, where the top-right position indicates an ideal tradeoff between objectives. DropLoss achieves an improved tradeoff between object categories, resulting in higher overall AP.

dictions. Based on LVIS (Gupta, Dollár, and Girshick 2019), the prediction threshold is reduced from 0.05 to 0.0, and we set the top 300 bounding boxes as prediction results. This setting is widely used in LVIS training and evaluation. For all the experiments, we report the average results of three independent runs of model training. The variances in AP are generally small (approximately 0.1-0.2).

Comparisons with state-of-the-art methods. In our experiments, we use Mask R-CNN (He et al. 2017) as our architecture and compare it with two baseline training methods: standard Mask R-CNN and the equalization loss (Tan et al. 2020). To verify that DropLoss is effective across different settings, we validate on several different architectures and backbones. We test ResNet50 and ResNet101 (He et al. 2016) as backbones, and compare the Cascades R-CNN (Cai and Vasconcelos 2018) as an alternative architecture to Mask R-CNN (He et al. 2020). Table 1 reports the results, where all methods are tested using the same experiment settings and environment. We find that DropLoss achieves improved

Method	Use RFS	AP (%)	AP ₅₀	AP ₇₅	AP _r	AP _c	AP _f	AP _s	AP _m	AP _L	AP _{bbox}
Sigmoid	-	21.5	33.4	22.9	4.7	21.2	28.6	15.6	29.3	39	21
Softmax	-	21.3	33.1	22.6	3	21.2	28.6	15.8	28.5	39.2	21
EQL (Tan et al. 2020)	-	23.8	36.3	25.2	8.5	25.2	28.3	17.1	31.4	41.7	23.5
DropLoss (Ours)	-	25.5	38.7	27.2	13.2	27.9	27.3	17.7	32.7	43.2	25.1
Sigmoid	✓	23.8	36.3	25.2	8.5	25.2	28.3	17.1	31.4	41.7	23.5
Softmax	✓	24.3	37.8	25.9	14.1	24.3	28.3	16.5	31.6	41.2	23.8
EQL (Tan et al. 2020)	✓	25.5	39	27.2	16.7	26.3	28.1	17.5	33	43	25
DropLoss (Ours)	✓	26.4	40.3	28.4	17.3	28.7	27.2	17.9	33.1	44	25.8

Method	Use RFS	AP (%)	AP ₅₀	AP ₇₅	AP _r	AP _c	AP _f	AP _s	AP _m	AP _L	AP _{bbox}
Baseline	-	16.2	25.9	16.9	0.7	12.6	27	10.5	22.7	32.7	16.6
EQL (Tan et al. 2020)	-	18.4	28.6	19.4	2.5	16.5	27.4	11.9	25.4	35.6	18.9
DropLoss (Ours)	-	19.8	30.9	20.9	3.5	20	26.7	12.9	27.5	37.1	20.4
Baseline	✓	18.8	29.6	19.9	5.6	16.6	27.1	11.6	25.6	35.7	19.2
EQL (Tan et al. 2020)	✓	21	32.7	22.3	9.1	20.1	27.3	13.1	28.5	39.2	21.7
DropLoss (Ours)	✓	22.3	34.5	23.6	12.4	22.3	26.5	13.9	29.9	40	22.9

Table 2: Evaluation on LVIS v0.5 (top) and LVIS v1.0 (bottom) validation sets *with* and *without* Repeat Factor Sampling (RFS). Here we use Mask-RCNN and ResNet-50. DropLoss achieves the best overall AP across both settings.

performance (in terms of overall AP) compared with both baselines across all backbones and architectures. We are most interested in the AP_r, AP_c, AP_f and AR. Although the AP_f (frequent) decreases slightly in our method, our AP_r (rare) and AP_c (*common*) increase significantly. Our method improves the AP and AR by a large margin, indicating the overall performance across all categories is improved. In particular, using Mask R-CNN with ResNet-50 as the backbone, we achieve a 1.7 AP improvement over the state-of-the-art method (Tan et al. 2020) (winner of the LVIS 2019 challenge). Across all the settings, compared with the baselines, DropLoss can more successfully balance the tradeoff between rare and frequent categories, resulting in better performance in the long-tailed distribution dataset.

Incorporating with Resampling Methods. Here we show that our approach can be combined with state-of-the-art resampling methods to improve learning long-tailed distribution further. Specifically, we adopt the Repeat Factor Sampling (RFS) (Gupta, Dollár, and Girshick 2019) that uses the number of images per category to determine the sampling frequency. Table 2 shows the quantitative comparisons of different loss function choices on the LVIS v0.5 validation set.¹ We find that applying RFS generally improves the performance of all the methods. The proposed DropLoss compares favorably against other baseline methods either with or without using RFS. Note that the RFS method rebalances based on *overall* data distribution, while DropLoss reweights the

loss based on statistics in the *each batch*. The complementary nature of the two methods may explain why integrating RFS and the DropLoss leads to improved results.

Measuring the Frequent-rare Category Performance Tradeoff. Methods for learning long-tail distribution often involve a tradeoff between accuracy on rare, common, and frequent categories. Here we wish to quantify this tradeoff for various methods. We compare our proposed DropLoss against three baselines: equalization loss (Tan et al. 2020), background equalization loss, and fixed drop ratio.

Equalization loss and DropLoss have no tunable hyperparameters. Background equalization loss has the log base as a tunable hyperparameter. A fixed drop ratio has the drop ratio as a hyperparameter. These methods may be adjusted to measure the tradeoff between object categories. We can use the Pareto Frontier from multi-objective optimization to visualize this tradeoff, as seen in Figure 3. We observe that for reweighting methods with tunable hyperparameters, improvement in rare AP_r or common AP_c generally leads to a rapid decrease in frequent AP_f. Our proposed DropLoss does not have tunable hyperparameters, but Figure 3 demonstrates that DropLoss balances more effectively between AP_r, AP_c and AP_f, resulting in higher overall AP than other baselines.

DropLoss adapts to the sampling distribution so that if a rare category appears in a given batch, its loss is less likely to be dropped. However, if a rare category does not appear in a batch, the chance of its loss being dropped is very high. This allows the network to *dynamically* attend to the categories that it sees in a given batch, decreasing drop loss probability selectively for only those categories. We postulate that this allows the network to achieve a better overall balance between frequent and infrequent categories.

Quantitative Comparison Between the DropLoss and

¹Note that the EQL results (25.5 AP) are not consistent with the reported results (26.1 AP). We use the public implementation without changes and report the average over 3 runs. The difference may be due to number of GPUs used for training, resulting in different batch normalization. For fair comparisons, we use the same hardware and experimental setting to train all models.



Figure 4: Visual results comparison. Qualitative results of the Mask R-CNN trained with standard cross-entropy loss (*Top*) and the proposed DropLoss (*Bottom*). Instances with scores larger than 0.5 are shown.

Our Proposed Baseline BEQL. To validate that the DropLoss provides better pareto-efficiency over BEQL, in Table 3, we compare the DropLoss with two *best-performing* results (in term of overall AP) from BEQL. DropLoss still offers overall better performance (not only in AP but also in AR) despite not sweeping the parameters on the validation set to find the best performance as in BEQL.

Method	AP (%)	AP _r	AP _c	AP _f	AP _{bbox}	AR
BEQL ($b = 4$)	25.2	14.6	28.1	25.9	24.9	34.1
BEQL ($b = 5$)	25.1	13.4	27.4	26.9	24.8	34.3
DropLoss	25.5	13.2	27.9	27.3	25.1	34.8

Table 3: Comparison between DropLoss and two top-performing results from our another proposed method BEQL. Evaluation on LVIS v0.5 validation set. DropLoss offers overall better performance (in both AP and AR) compared with the BEQL baseline.

Visual Results. Figure 4 demonstrates the results on a dense instance segmentation example containing common/rare category. For example, the goose in the first image is a ‘common’ object category. We demonstrate the suppression of these less-frequent categories, as most of the geese in this image are classified as background or with low confidence. In contrast, training the model with the proposed loss correctly identifies all geese as foreground, and predicts category “goose” with high confidence and other waterbirds with lower confidence. Despite the stochastic removal of rare and common category losses for background proposals, we find that the network does not misclassify background regions as foreground. The

distinction between background and foreground is likely less difficult to learn than the distinction between foreground image categories, so reducing background gradients does not appear to significantly affect background/foreground classification. By reducing the suppression of rare and common categories via background predictions, our method allows for rare and common categories to improve prediction scores, decreasing bias towards frequent categories.

Conclusions

Through analysis of the loss gradient distributions over rare, common, and frequent categories, we discovered that disproportionate background gradients suppress less-frequent categories in the long-tailed distribution problem. To address this problem, we propose DropLoss, which balances the background loss gradients between different categories via random sampling and reweighting. Our method provides a sizable performance improvement across different backbones and architectures by improving the balance between object categories in the long-tail instance segmentation setting. While we focus on the challenging problem of instance segmentation in our experiments, we expect that DropLoss may be applicable to other visual recognition problems with long-tailed distributions. We leave the exploration to other tasks as future work.

Acknowledgments

This work was supported in part by Virginia Tech, the Ministry of Education BioPro A+, and MOST Taiwan under Grant 109-2634-F-001-012. We are particularly grateful to the National Center for High-performance Computing for computing time and facilities.

References

- Cai, Z.; and Vasconcelos, N. 2018. Cascade R-CNN: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6154–6162.
- Cao, K.; Wei, C.; Gaidon, A.; Arechiga, N.; and Ma, T. 2019. Learning imbalanced datasets with label-distribution-aware margin loss. In *Advances in Neural Information Processing Systems*.
- Chawla, N. V.; Bowyer, K. W.; Hall, L. O.; and Kegelmeyer, W. P. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16: 321–357.
- Cui, Y.; Jia, M.; Lin, T.-Y.; Song, Y.; and Belongie, S. 2019. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9268–9277.
- Drummond, C.; Holte, R. C.; et al. 2003. C4. 5, class imbalance, and cost sensitivity: why under-sampling beats over-sampling. In *Workshop on learning from imbalanced datasets II*, volume 11, 1–8. Citeseer.
- Girshick, R. 2015. Fast R-CNN. In *Proceedings of the IEEE international conference on computer vision*, 1440–1448.
- Girshick, R.; Donahue, J.; Darrell, T.; and Malik, J. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 580–587.
- Gupta, A.; Dollár, P.; and Girshick, R. B. 2019. LVIS: A Dataset for Large Vocabulary Instance Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*.
- Han, H.; Wang, W.-Y.; and Mao, B.-H. 2005. Borderline-SMOTE: a new over-sampling method in imbalanced data sets learning. In *International conference on intelligent computing*, 878–887. Springer.
- He, H.; Bai, Y.; Garcia, E. A.; and Li, S. 2008. ADASYN: Adaptive synthetic sampling approach for imbalanced learning. In *2008 IEEE international joint conference on neural networks (IEEE world congress on computational intelligence)*, 1322–1328. IEEE.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask R-CNN. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. B. 2020. Mask R-CNN. *IEEE Trans. Pattern Anal. Mach. Intell.* .
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Identity mappings in deep residual networks. In *European conference on computer vision*, 630–645. Springer.
- Hensman, P.; and Masko, D. 2015. The impact of imbalanced training data for convolutional neural networks. *Degree Project in Computer Science, KTH Royal Institute of Technology* .
- Huang, C.; Li, Y.; Loy, C. C.; and Tang, X. 2016. Learning Deep Representation for Imbalanced Classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5375–5384.
- Jamal, M. A.; Brown, M.; Yang, M.-H.; Wang, L.; and Gong, B. 2020. Rethinking Class-Balanced Methods for Long-Tailed Visual Recognition from a Domain Adaptation Perspective. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Kahn, H.; and Marshall, A. W. 1953. Methods of reducing sample size in Monte Carlo computations. *Journal of the Operations Research Society of America* 1(5): 263–278.
- Kang, B.; Xie, S.; Rohrbach, M.; Yan, Z.; Gordo, A.; Feng, J.; and Kalantidis, Y. 2019. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217* .
- Li, B.; Liu, Y.; and Wang, X. 2019. Gradient harmonized single-stage detector. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 8577–8584.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017a. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017b. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollar, P. 2017c. Focal Loss for Dense Object Detection. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017d. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. SSD: Single shot multibox detector. In *European conference on computer vision*, 21–37. Springer.
- Liu, Z.; Miao, Z.; Zhan, X.; Wang, J.; Gong, B.; and Yu, S. X. 2019. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2537–2546.
- Mahajan, D.; Girshick, R.; Ramanathan, V.; He, K.; Paluri, M.; Li, Y.; Bharambe, A.; and van der Maaten, L. 2018. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 181–196.
- Pouyanfar, S.; Tao, Y.; Mohan, A.; Tian, H.; Kaseb, A. S.; Gauen, K.; Dailey, R.; Aghajanzadeh, S.; Lu, Y.-H.; Chen, S.-C.; et al. 2018. Dynamic sampling in convolutional neural networks for imbalanced data classification. In *2018 IEEE conference on multimedia information processing and retrieval (MIPR)*, 112–117. IEEE.

- Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You Only Look Once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 779–788.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, 91–99.
- Shen, L.; Lin, Z.; and Huang, Q. 2016. Relay backpropagation for effective learning of deep convolutional neural networks. In *European conference on computer vision*, 467–482. Springer.
- Shu, J.; Xie, Q.; Yi, L.; Zhao, Q.; Zhou, S.; Xu, Z.; and Meng, D. 2019. Meta-weight-net: Learning an explicit mapping for sample weighting. In *Advances in Neural Information Processing Systems*, 1917–1928.
- Tan, J.; Wang, C.; Li, B.; Li, Q.; Ouyang, W.; Yin, C.; and Yan, J. 2020. Equalization Loss for Long-Tailed Object Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11662–11671.
- Tsai, C.-F.; Lin, W.-C.; Hu, Y.-H.; and Yao, G.-T. 2019. Under-sampling class imbalanced datasets by combining clustering analysis and instance selection. *Information Sciences* 477: 47–54.
- Wang, Y.-X.; Ramanan, D.; and Hebert, M. 2017. Learning to model the tail. In *Advances in Neural Information Processing Systems*, 7029–7039.
- Wu, Y.; Kirillov, A.; Massa, F.; Lo, W.-Y.; and Girshick, R. 2019. Detectron2. <https://github.com/facebookresearch/detectron2>.
- Yin, X.; Yu, X.; Sohn, K.; Liu, X.; and Chandraker, M. 2019. Feature transfer learning for face recognition with under-represented data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5704–5713.
- Zhang, X.; Fang, Z.; Wen, Y.; Li, Z.; and Qiao, Y. 2017. Range loss for deep face recognition with long-tailed training data. In *Proceedings of the IEEE International Conference on Computer Vision*, 5409–5418.
- Zou, Y.; Yu, Z.; Vijaya Kumar, B.; and Wang, J. 2018. Un-supervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, 289–305.