

# PLM: PARTIAL LABEL MASKING FOR IMBALANCED MULTI-LABEL CLASSIFICATION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Neural networks trained on real-world datasets with long-tailed label distributions are biased towards frequent classes and perform poorly on infrequent classes. The imbalance in the ratio of positive and negative samples for each class skews network output probabilities further from ground-truth distributions. We propose a method, Partial Label Masking (PLM), which utilizes this ratio during training. By stochastically masking labels during loss computation, the method balances this ratio for each class, leading to improved recall on minority classes and improved precision on frequent classes. The ratio is estimated adaptively based on the network’s performance by minimizing the KL divergence between predicted and ground-truth distributions. Whereas most existing approaches addressing data imbalance are mainly focused on *single-label* classification and do not generalize well to the multi-label case, this work proposes a general approach to solve the data imbalance issue for *multi-label* classification. PLM is versatile: it can be applied to most objective functions and it can be used alongside other strategies for class imbalance. Our method achieves strong performance when compared to existing methods on both *multi-label* (MultiMNIST and MSCOCO) and *single-label* (imbalanced CIFAR-10 and CIFAR-100) image classification datasets.

## 1 INTRODUCTION

The impressive performance of deep learning methods has led to the creation of many large-scale datasets (Deng et al., 2009; Gu et al., 2018; Lin et al., 2014; Krasin et al., 2017). Due to the naturally imbalanced distribution of objects within the world, these datasets contain imbalanced numbers of samples for different classes. The class labels in these datasets form a long-tailed distribution: several classes appear frequently (the head classes), while many classes contain few samples (the tail classes). This imbalance causes classifiers to perform poorly, especially on classes which are infrequent in training. Finding a solution to this problem is necessary to successfully scale deep networks to larger real-world datasets which tend to have long-tail data distributions.

Previous works on data imbalance (Buda et al., 2018; Cui et al., 2019; Cao et al., 2019) tend to focus on single-label classification datasets, but these approaches do not generalize well to the multi-label case, where a single sample may have multiple class labels. Two common methods for learning long-tailed distributions, reweighting and resampling, were not designed for data with multiple labels. Our experiments show that reweighting based on the inverse number of samples performs poorly on multi-label datasets; also, it is difficult to resample multi-label data due to the co-occurrence of labels within individual samples. Therefore, in this work, we propose a general solution for long-tailed imbalance which works for both multi-label and single-label classification.

Classifiers trained on imbalanced multi-label datasets tend to over-predict frequent classes and under-predict minority classes. This behaviour is displayed in Figure 1. When the class is not present within the image, the probability output for the frequent classes is skewed towards 1; conversely, when an infrequent class is present within the sample, the classifier outputs a low probability score. These output probability distributions differ greatly from the ideal distribution (i.e. the ground-truth distribution where all positive samples are labeled 1 and all negative samples are labeled 0). *We argue, that this behaviour is caused not only by an imbalance in the number of positive samples between different classes, but also by the ratio of positive and negative samples for each class.*

As dataset imbalance increases, the ratio of positive samples to negative samples increases for head classes and decreases for tail classes. We find that the change in this ratio greatly impacts the classifier’s ability to generalize. If a class has large ratio of positive samples to negative samples, the classifier over-predicts the given class, leading to an increase of false positive predictions; conversely, a small ratio leads to under-prediction and an increase of false negatives. Assuming there is an optimal ratio which can minimize the over/under-predictions, an algorithm can estimate and leverage this ratio to improve network performance.

We present Partial Label Masking (PLM): a novel approach for training classifiers on imbalanced multi-label datasets which improves network generalization by leveraging this ratio. By partially masking positive and negative labels for frequent and infrequent classes respectively; our method reduces the discrepancy between the classifiers’ output probability distribution and the ground-truth distribution (as seen in Figure 1). Our method performs this masking stochastically for each sample and it continually adapts the target ratio based on the classifier’s output probabilities. This leads to improved precision on classes with many samples and improved recall on classes with few samples. Moreover, our method consistently improves performance on difficult classes, regardless of the number of samples.

Our contributions include: (i) we present a general solution for data imbalance which balances the ratio between positive and negative samples, (ii) we propose an adaptive strategy to determine the ideal ratio which minimizes the difference between predicted probability and ground-truth distributions, (iii) we empirically evaluate our method on both multi-label datasets (imbalanced MultiMNIST and MSCOCO) and single-label datasets (CIFAR10 and CIFAR100), and (iv) we thoroughly analyse our method’s ability to improve classifiers’ performance on both difficult and infrequent classes.

## 2 METHOD

**Notation** We denote a dataset with  $N$  samples and  $C$  class categories as  $\mathcal{D} = \{(x^{(i)}, \mathbf{y}^{(i)})\}$  where  $x^{(i)}$  is the  $i^{th}$  input and  $\mathbf{y}^{(i)} = [y_1^{(i)}, \dots, y_C^{(i)}]$  is its corresponding binary label vector; note that multiple elements in this label may be non-zero in multi-label classification. The number of positive samples for a class  $c$  is  $n_c^+ = |\{i : i \in \{1, \dots, N\}, \text{ where } y_c^{(i)} = 1\}|$ , and the number of negative samples is  $n_c^- = N - n_c^+$ . The ratio of positive samples to negative samples for a class  $c$  is  $r_c = \frac{n_c^+}{n_c^-}$ . We denote the “optimal” ratio, which minimizes over/under-predictions as  $\bar{r}_c$ .

### 2.1 PARTIAL LABEL MASKING

Given a network prediction,  $\hat{\mathbf{y}}^{(i)}$ , the total loss is sum of the losses over all  $C$  classes:

$$L(\mathbf{y}^{(i)}, \hat{\mathbf{y}}^{(i)}) = \sum_{j=1}^C \ell(y_j^{(i)}, \hat{y}_j^{(i)}). \tag{1}$$

This is a general form (e.g.  $\ell(y, \hat{y}) = -[y \log(\hat{y}) + (1 - y) \log(1 - \hat{y})]$  for binary cross-entropy).

Partial Label Masking (PLM) masks the loss computed for certain classes. For each image  $i$  we generate a binary mask  $\mathbf{g}^{(i)} = [g_1^{(i)}, \dots, g_C^{(i)}]$  that masks loss calculations for certain classes as follows:

$$L(\mathbf{y}^{(i)}, \hat{\mathbf{y}}^{(i)}, \mathbf{g}^{(i)}) = \sum_{j=1}^C g_j^{(i)} \ell(y_j^{(i)}, \hat{y}_j^{(i)}). \tag{2}$$

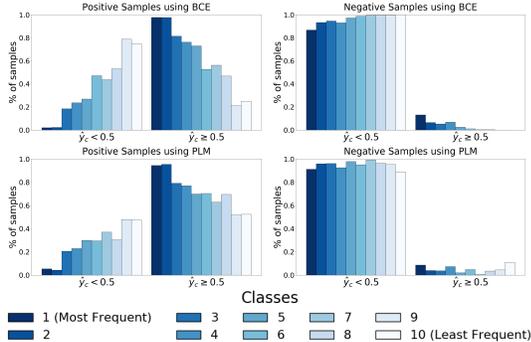


Figure 1: The output probability ( $\hat{y}_c$ ) distributions of a ResNet-32 classifier trained on artificially imbalanced CIFAR-10 for positive (left) and negative samples (right). For frequent classes, the predicted distribution skews towards 1; for infrequent classes, it skews towards 0. Classifiers trained using PLM (bottom) reduces this bias, when compared to classifiers trained with binary cross-entropy (top).

Note that equation 2 is a generalized form of 1 when all masks equal one, i.e.  $g_c^{(i)} = 1$ . Therefore, PLM can be used with any loss which sums over different classes.

**Masks Generation** With this formulation, one can generate masks so that frequent class samples are ignored by the loss calculation and an equal number of samples for each class is used to train the network, which is typically employed in under-sampling. The issue with under-sampling samples, however, is that frequent and infrequent classes can co-occur within the same sample, so removing it from training would not meaningfully change the imbalance. The benefit of the PLM formulation is that individual positive, or negative, labels can be masked to ensure a specific number of labels for each class is used to train the network regardless of co-occurrence.

Instead of under-sampling based on the number of positive samples, we consider the ratio of positive to negative samples for each class,  $r_c$ . Assuming there is some ideal ratio,  $\bar{r}_c$ , which minimizes over/under-predictions, we can generate the masks,  $\mathbf{g}^{(i)}$ , stochastically as a function of the label:

$$g_c^{(i)} = f\left(y_c^{(i)}\right) = \begin{cases} \mathbb{1}\left[\frac{\bar{r}_c}{r_c}\right], & r_c > \bar{r}_c \text{ and } y_c^{(i)} = 1 \\ \mathbb{1}\left[\frac{r_c}{\bar{r}_c}\right], & r_c < \bar{r}_c \text{ and } y_c^{(i)} = 0 \\ 1, & \text{otherwise} \end{cases} \quad (3)$$

where  $\mathbb{1}[p] = 1$  with probability  $p$ , and 0 with probability  $1 - p$ . The probabilities used,  $\bar{r}_c/r_c$  and  $r_c/\bar{r}_c$ , ensure that the ratio between positive and negative samples on which loss is calculated is  $\bar{r}_c$ . With partial label masking, we are able to train a classifier with a set ratio between positives and negatives for each class. Now the question arises: How does one determine this ideal ratio?

**Ratio Selection** The ratio between positives and negatives for each class should be set to minimize over-predictions and under-predictions for the head and tail classes respectively. Since decreasing the ratio reduces the number of positive samples used for loss calculation and increasing the ratio reduces the number of negative samples, the ratio selected for the head classes should be decreased (i.e.  $\bar{r}_c < r_c$ ). On the other hand, the ratio for the tail classes should be increased ( $\bar{r}_c > r_c$ ).

One simple approach for setting the ratio would be to use the average ratio of all the classes. This ensures a decrease in the ratio for head classes and increase for tail classes. Although, this ratio leads to improved results on some datasets, it is sensitive to different levels of imbalance within the dataset as well as the relative difficulty of the various classes. Ideally this hyper-parameter would change based on the performance of the classifier on the dataset on-hand. To this end, we propose a method to estimate the ratio adaptively based on the output probability distribution of the classifier.

**Ratio Adaptation** For a class  $c$ , the set of output probabilities for the positive and negative samples are represented by  $\hat{S}_c^+$  and  $\hat{S}_c^-$ , respectively. The ground-truth sets are  $S_c^+$  and  $S_c^-$ , where members of this set will be 1 for a positive sample and 0 for a negative sample. Formally, these sets can be defined as:

$$\begin{aligned} S_c^+ &= \{y_c^{(i)} | y_c^{(i)} = 1, \forall i\}, \\ \hat{S}_c^+ &= \{\hat{y}_c^{(i)} | y_c^{(i)} = 1, \forall i\}, \\ S_c^- &= \{y_c^{(i)} | y_c^{(i)} = 0, \forall i\}, \\ \hat{S}_c^- &= \{\hat{y}_c^{(i)} | y_c^{(i)} = 0, \forall i\}. \end{aligned} \quad (4)$$

A discrete distribution is formed using these sets by placing the probabilities in  $\tau$  bins of width  $1/\tau$ . For each set, we denote these distributions as  $P_c^+$ ,  $\hat{P}_c^+$ ,  $P_c^-$ , and  $\hat{P}_c^-$  respectively. When data imbalance is present, the classifiers output probabilities skew further from the ground-truth distributions. For head classes, the difference between  $\hat{P}_c^-$  and  $P_c^-$  increases as the probability outputs for negative samples become larger; conversely, for tail classes, probabilities are pushed closer to 0 for positive samples, so the difference between  $\hat{P}_c^+$  and  $P_c^+$  increases.

We utilize this difference to change the ratio until the output distributions better resemble the ground-truth distribution. To compute the difference between the distributions, we use the Kullback-Leibler divergence:

$$D_c^+ = D_{\text{KL}}\left(\hat{P}_c^+ || P_c^+\right); \text{ and } D_c^- = D_{\text{KL}}\left(\hat{P}_c^- || P_c^-\right). \quad (5)$$

---

**Algorithm 1** PLM training algorithm. Inputs are a network  $F_\theta$  and the dataset  $\mathcal{D} = \{(x^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^N$ .

---

- 1: Initialize  $\bar{r}_{c,1}$
  - 2: **for**  $t = 1$  **to**  $T$  **epochs do**
  - 3:   Generate masks  $g_c^{(i)} = f(y_c^{(i)})$ , using  $\bar{r}_{c,t}$  ▷ Equation 3
  - 4:   **repeat**
  - 5:      $\mathcal{B} \leftarrow \text{SampleMiniBatch}(\mathcal{D})$
  - 6:      $\hat{\mathbf{y}} \leftarrow F_\theta(\mathcal{B})$  ▷ Perform forward pass on mini-batch
  - 7:      $F_\theta \leftarrow F_\theta - \alpha \nabla_{\theta} L(\mathbf{y}, \hat{\mathbf{y}}, \mathbf{g})$  ▷ One SGD step with loss in equation 2
  - 8:   **until** all data has been sampled
  - 9:   Generate  $P_c^+, P_c^-, \hat{P}_c^+, \hat{P}_c^-$  from all  $\mathbf{y}^{(i)}$  and  $\hat{\mathbf{y}}^{(i)}$
  - 10:   Compute  $D_c = \tilde{D}_c^+ - \tilde{D}_c^-$  using equations 5 and 6
  - 11:    $\bar{r}_{c,t+1} \leftarrow e^{\lambda D_c} \bar{r}_{c,t}$  ▷ Update ratio
- 

We normalize the divergence

$$\tilde{D}_c^+ = \frac{D_c^+ - \mu^+}{\sigma^+}; \text{ and } \tilde{D}_c^- = \frac{D_c^- - \mu^-}{\sigma^-}, \quad (6)$$

with their means ( $\mu^+$  and  $\mu^-$ ) and standard deviations ( $\sigma^+$  and  $\sigma^-$ ). After normalization,  $\tilde{D}_c^+$  tends to be positive when the network under-predicts class  $c$  and negative when the network over-predicts. The inverse holds for  $\tilde{D}_c^-$ : it is positive when the network over-predicts, and negative otherwise.

The ratio should be adjusted during training until the divergence scores in equation 6 become balanced ( $D_c^+ = D_c^-$ ). At each epoch  $t$ , the ratio  $\bar{r}_{c,t}$  becomes

$$\bar{r}_{c,t} = e^{\lambda D_c} \bar{r}_{c,t-1}, \quad (7)$$

where  $\bar{r}_{c,1}$  is the initial ratio (e.g. the dataset’s ratio,  $r_c$ ),  $D_c = \tilde{D}_c^+ - \tilde{D}_c^-$ , and  $\lambda$  is a hyper-parameter which controls the rate of change of the ratio. The exponential term in equation 7 increases, or decreases, the current ratio for each class based on the divergence between ground-truth and output probability distributions: if  $D_c > 0$ , which tends to occur for infrequent classes, then the ratio increases; if  $D_c < 0$ , which occurs for frequent classes, then the ratio decreases. If  $D_c = 0$  (i.e.  $D_c^+ = D_c^-$ ), then the ratio remains unchanged as some optimal ratio has been reached. The training procedure used for partial label masking is described in Algorithm 1.

### 3 EXPERIMENTAL EVALUATION

We evaluate our proposed method on several image classification benchmarks with varying levels of imbalance. Imbalance,  $\rho$ , is denoted as the ratio between the number of samples for the most frequent class and the least frequent class:  $\rho = \max_i\{n_i^+\}/\min_i\{n_i^+\}$ .

#### 3.1 MULTI-LABEL CLASSIFICATION

**Datasets** We evaluate our method on two multi-label datasets: MultiMNIST and the large-scale real-world multi-label dataset MSCOCO (Lin et al., 2014). We construct the imbalanced MultiMNIST dataset by superimposing two MNIST (LeCun et al., 1998) digits into a single image. We sample the MNIST training set to obtain a long-tail distribution over the different digits. Due to cooccurrence of the same digit within a sample, the final imbalance of the dataset is  $\rho = 90.33$ . The test set consists of 90000 samples; it is roughly balanced, with all classes being present a similar number of samples<sup>1</sup>. MSCOCO is a multi-label image classification benchmark, which contains a large amount of imbalance ( $\rho = 352.92$ ). We use the standard train/test split which consists of 82,783 training and 40,504 evaluation images. It contains 80 classes, with an average of 2.9 labels per sample.

**Metrics** We evaluate classifier performance using multiple standard multi-label classification metrics: per-class precision, per-class recall, F1-score, and 0-1 exact match accuracy. To evaluate performance on tail classes, we measure some metrics averaged over the  $K$  most infrequent classes. A description of these metrics is in Section A.1 of the Appendix.

<sup>1</sup>Additional details of the dataset’s construction and the dataset statistics are in section A.2 of the Appendix.

Loss	Und.	Precision	Recall			F1 Score			0-1 Acc.
		Avg.	Avg.	K=5	K=3	Avg.	K=5	K=3	
BCE	✗	86.96	78.65	63.39	51.50	80.35	73.87	65.73	53.42
Focal	✗	86.75	78.28	62.53	50.44	79.93	73.33	64.88	52.64
BCE+CB	✗	85.86	76.55	61.02	49.12	78.05	72.32	63.81	48.51
BCE	✓	87.97	79.00	65.36	54.16	81.39	75.57	67.92	54.46
PLM (BCE)	✗	87.41	<b>81.55</b>	71.91	65.31	<b>83.73</b>	79.17	74.27	<b>56.95</b>
PLM (Focal)	✗	87.50	80.95	70.27	62.31	83.22	78.13	72.61	56.45
PLM (BCE+CB)	✗	<b>89.37</b>	75.24	65.31	56.25	80.52	75.64	69.28	49.95
PLM (BCE)	✓	88.54	80.38	<b>71.29</b>	<b>65.86</b>	83.56	<b>79.44</b>	<b>75.10</b>	56.53

Table 1: Results for multi-label image classification on MultiMNIST. “Und.” denotes undersampling and “+CB” denotes class-based reweighting (Cui et al., 2019).

Loss	Und.	Precision	Recall			F1 Score			0-1 Acc.
		Avg.	Avg.	K=20	K=10	Avg.	K=20	K=10	
BCE	✗	<b>69.35</b>	39.89	30.02	16.00	48.74	37.94	21.80	<b>21.76</b>
Focal	✗	68.01	38.62	28.92	15.10	47.38	36.52	20.46	20.88
BCE+CB	✗	67.82	37.23	35.25	23.68	46.10	43.64	30.53	19.43
BCE	✓	65.96	37.75	35.18	19.50	46.06	41.59	24.16	18.60
PLM (BCE)	✗	67.75	<b>41.29</b>	39.69	29.91	<b>50.16</b>	44.24	31.26	20.51
PLM (Focal)	✗	66.89	39.48	37.63	27.67	48.49	42.38	29.56	19.66
PLM (BCE+CB)	✗	65.71	37.06	39.06	29.25	46.39	<b>45.68</b>	<b>33.18</b>	17.41
PLM (BCE)	✓	63.72	36.23	<b>40.21</b>	<b>31.46</b>	44.85	44.49	32.72	16.71

Table 2: Results for multi-label image classification on MSCOCO. “Und.”, denotes undersampling and “CB” denotes class-based reweighting (Cui et al., 2019).

**Baselines** Since recent class imbalance works tend to focus on single-label classification, we compare our method with those approaches which can be applied to multi-label classification. We use the following baselines: i) binary cross-entropy, ii) focal loss (Lin et al., 2017), iii) reweighting by inverse number of classes (CB), and iv) undersampling. For reweighting, we use the formulation:

$$L(\mathbf{y}^{(i)}, \hat{\mathbf{y}}^{(i)}) = \sum_{j=1}^C w_j \ell(y_j^{(i)}, \hat{y}_j^{(i)}), \quad (8)$$

where  $w_j$  is the inverse number of samples for class  $j$ , as calculated in (Cui et al., 2019). Undersampling that removes all imbalance is a non-trivial operation when multiple labels can be present in each sample. Therefore, for each epoch, we select  $S$  samples from each class (without duplication) to reduce the imbalance seen by the network ( $S=500$  for MultiMNIST and  $S=1000$  for MSCOCO).

**Implementation** For experiments on MultiMNIST, we train a ResNet-12 (He et al., 2016) model using stochastic gradient descent (SGD) with momentum of 0.9 for 90 epochs. The initial learning rate is set to 0.1 which is decayed by a factor of 0.1 at epochs 60 and 80. We employ a linear learning rate warm-up (Goyal et al., 2017) for the first 5 epochs. The network is model with a batch size of 128. On MSCOCO, a ResNet-50 model is trained with a similar training procedure; however, the initial learning rate is set to 0.4 which is decayed epochs 30, 60, and 80, and the batch size is increased to 200. For both datasets, we initialize the ratio with the dataset’s ratio (i.e.  $\bar{r}_{c,1} = r_c$ ) and set  $\lambda$  to 0.01. Since the imbalance in MSCOCO is large, we clip the ratios within the range  $[0, 1]$ .

**Results on MultiMNIST** We present results for MultiMNIST in Table 1. Undersampling improves results on MultiMNIST across all metrics, but the use of reweighting (CB) leads to a decrease in performance since the technique is formulated for the single-label case. The use of PLM during training improves results across all metrics. The F1-score improves by 2.34% and the exact match accuracy improves by 2.49% over the next best baseline. Notably, PLM greatly improves performance on the minority classes - for the three most infrequent digits, the recall increases by 11.7% and the F1-score increases by 7.18%.

**Results on MSCOCO** Results for MSCOCO are presented in Table 2. We find conventional methods designed for single-label data imbalance (focal loss, class-based reweighting, and undersampling) tend to underperform standard binary cross-entropy (BCE) training. Although reweighting

Dataset	CIFAR-10				CIFAR-100			
	100		10		100		10	
	Overall	K=5	Overall	K=5	Overall	K=5	Overall	K=5
BCE	29.43	46.46	12.71	15.98	59.88	79.18	43.35	52.12
Focal	29.69	46.88	12.63	16.74	60.04	81.62	43.04	53.74
CB	25.43	-	12.51	-	60.40	-	42.01	-
LDAM-DRW	22.97	-	11.84	-	57.96	-	41.29	-
PLM (BCE)	24.86	36.05	12.35	14.98	55.98	73.14	41.71	49.78
PLM (Focal)	25.19	36.53	12.20	14.60	56.19	72.39	41.54	48.83

Table 3: Error rates for single-label image classification on imbalanced CIFAR-10 and CIFAR-100. We measure the average error across all classes (Overall) as well as the error on the K classes with the fewest number of samples (K=5 and K=50 for CIFAR-10 and CIFAR-100 respectively).

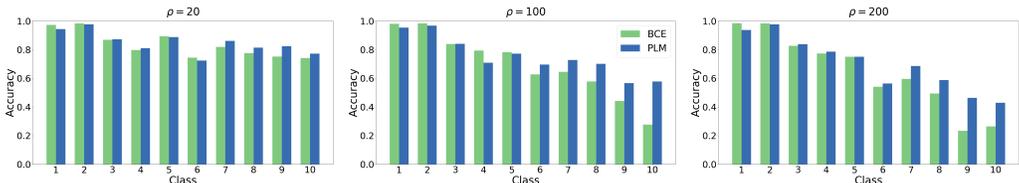


Figure 2: The test accuracy for each class on CIFAR-10 with varying levels of imbalance. PLM leads to improved generalization on minority classes, especially when there is a large amount of imbalance.

and under-sampling lead to improvements on tail classes, they dramatically reduce performance on head classes leading to an overall performance decrease. PLM, on the other hand, achieves an overall improvement on recall (1.4%) and f1-score (1.42%). There is a decrease in precision due to the drastic improvement of recall on minority classes - this is further discussed in section 3.3. Whereas PLM leads to a 0-1 accuracy improvement on MultiMNIST, there is a decrease on MSCOCO due to the imbalance of the MSCOCO test set and the bias of this metric to frequently occurring classes. Overall, PLM greatly improves performance on tail classes: when compared to the BCE baseline the 10 most infrequent classes have a 13.91% increase in recall and a 9.46% increase in f1-score. Moreover, these experiments highlight the versatility of our method: PLM can be used on top of existing data imbalance methods as well as different losses.

### 3.2 SINGLE-LABEL CLASSIFICATION

**Datasets** CIFAR-10 and CIFAR-100 (Krizhevsky et al., 2009) contain  $32 \times 32$  images with 10 and 100 classes respectively. To artificially create imbalance, we reduce the training examples for certain classes, while keeping the 10,000 validation images unchanged. Following the experimental setup in (Cui et al., 2019), we evaluate on varying levels of long-tailed imbalance, ranging from 10 to 200.

**Baselines** We compare our method with the following baselines: i) binary-cross entropy (BCE), ii) focal loss, iii) reweighting (CB) (Cui et al., 2019), and iv) LDAM-DRW (Cao et al., 2019). The later two are recent methods for dealing with imbalanced single-label datasets.

**Implementation** For the CIFAR experiments, we train a ResNet-32 following the training procedure described in (Cui et al., 2019). Again, we initialize the ratio with the dataset’s ratio. The hyper-parameter  $\lambda$  is set to 1.0 and 0.01 for the CIFAR-10 and CIFAR-100 experiments respectively. These values were selected empirically - we analyse the effect of this hyper-parameter in section 3.3.

**Results** Table 3 contains PLM’s results on CIFAR-10 and CIFAR-100, respectively, for imbalance factors<sup>2</sup>  $\rho = 10$  and  $\rho = 100$ . We achieve improvement when compared to the common methods for data imbalance - focal loss and reweighting based on the inverse effective number of samples (CB). Furthermore, we achieve comparable accuracy when compared to the recent LDAM-DRW approach.

<sup>2</sup>We present results on other imbalance factors  $\rho = 20$ ,  $\rho = 50$ , and  $\rho = 200$  in section A.3 of the Appendix.

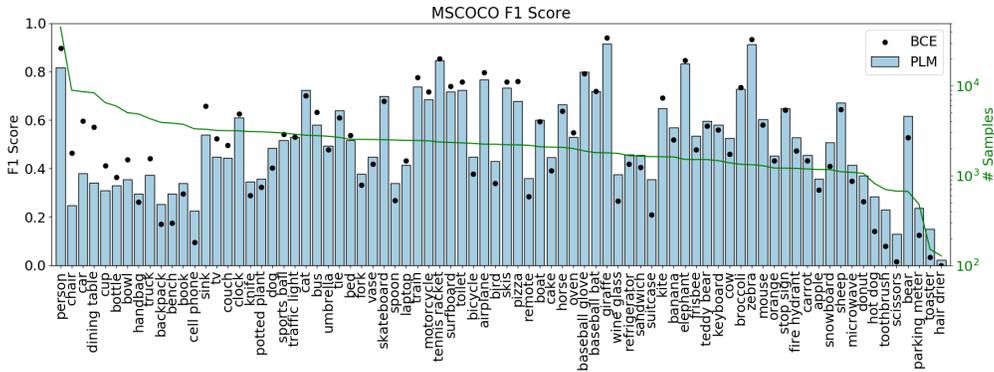


Figure 3: The class-wise F1 score on the MSCOCO dataset. The classes are ordered based on the number of samples in the training set. The bars represent the results achieved by using PLM, and the points are the results with BCE. Not only does PLM improves results on the 15 most infrequent classes, but also it improves scores on difficult classes, like *backpack*, *bench*, and *cell phone*.

### 3.3 ADDITIONAL ANALYSIS

**Generalization on minority classes** To perform well on long-tailed datasets, classifiers must generalize well on classes which have few training samples. All of our experiments show that PLM greatly improves performance on minority classes. Figure 2 displays the class-level accuracy on the CIFAR-10 dataset for various levels of imbalance; as the imbalance factor increases, there is a corresponding relative improvement on tail classes. Similarly, the MSCOCO experiments show an improvement on all 15 of the least frequent classes in terms of F1-score, as seen in Figure 3.

**Effects on Precision and Recall** In imbalanced datasets, classifiers tend to have poor precision on frequent classes and poor recall on infrequent classes. This is due, in part, to the large difference in the ratio of positive to negative samples. By minimizing the KL divergence between the ground-truth and predicted output distributions, PLM effectively balances this ratio, leading to improved performance. We present per-class precision and recall graphs in section A.3 of the Appendix. In summary, on MSCOCO, we find that our classifier’s precision over the 5 most frequent classes improves by 12.73% whereas its recall over the 5 least frequent classes improves by 13.04%. In general, PLM leads to improved recall over most classes, which suggests that our method would be best suited for applications which prioritize true positives over false positives.

**Ratio Adaptation** Since our approach attempts to estimate an optimal ratio for training a network, we investigate how the ratio changes throughout training. We present the effect of ratio’s rate of change ( $\lambda$ ) on the CIFAR-10 and CIFAR-100 datasets in Table 3.3. We find that all tested values lead to an improvement over no change (i.e. using standard BCE). On CIFAR-10, a higher rate of change ( $\lambda = 1.0$ ) leads to best performance. Meanwhile, a lower rate of change ( $\lambda = 0.01$ ) leads to best results on CIFAR-100 (this is also the case on MultiMNIST and MSCOCO). Figure 4 depicts how the ratio changes throughout training on the imbalanced MultiMNIST dataset. The ratio of head classes (digits 0, 1, and 2) decrease, while the ratio of the tail classes (digits 7, 8, and 9) increase. Over the 90 epochs, the ratio of the most frequent class (digit 0) decreases from 1.82 to 0.24 (decreasing by a factor of 7.58), and the ratio of the least frequent class (digit 9) increases from 0.007 to 0.047 (increasing by a factor of 6.54). This demonstrates that the original ratios (i.e. the dataset’s ratios) are not optimal for performance; by adapting the ratio, PLM greatly improves the training process.

**Improvements on difficult classes** Another benefit of using PLM during training is the improvements observed on difficult classes. This behaviour can be seen in Figure 3, with categories *backpack*, *bench*, and *cell phone*. These categories contain many samples, but achieve less than 20% F1-score when trained with BCE; when the network is trained with PLM, each of these categories improve by an average of 11.16%. This improvement is caused by the increasing of the ratio  $\bar{r}_c$  during training, which leads to masking more negative labels for these classes and an increase in recall.

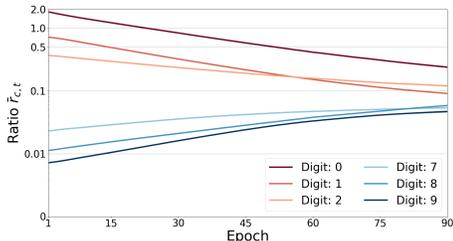


Figure 4: The change of the ratio over time on imbalanced MultiMNIST. The ratio for tail classes tend to increase while the ratio for head classes tend to decrease.

## 4 RELATED WORK

Many works have studied the problem of long-tailed data imbalance. Extensive overviews of the problem can be found for both classical methods (Japkowicz & Stephen, 2002; He & Garcia, 2009) and convolutional neural networks (Buda et al., 2018).

**Resampling** Resampling techniques either over-sample the infrequent classes (Chawla et al., 2002; Han et al., 2005; He et al., 2008; Mullick et al., 2019), or under-sample the frequent classes (Drummond et al., 2003; Liu et al., 2008). By resampling, the amount of imbalance seen by the network during training is greatly reduced. One benefit of PLM over conventional resampling is that it can be successfully applied to the multilabel case, where the co-occurrence of labels within individual samples causes issues with traditional resampling approaches. If tail and head classes co-occur frequently within the same inputs, then resampling would not result in a meaningful change of the imbalance. However, PLM masks individual labels, which allows for the head class’ label to be masked (i.e. removed from training) while keeping the tail class, resulting in more balanced training.

**Reweighting** Reweighting methods apply weights to certain classes, or samples, during the objective calculation. The standard approach is to apply class weights based on the inverse number of samples (Huang et al., 2016; Wang et al., 2017; Khan et al., 2019; Huang et al., 2019), or the square-root of the inverse number of samples (Mikolov et al., 2013). Cui et al. (2019) proposed class weights based on the effective number of samples, which improved results on more frequent classes. The focal loss (Lin et al., 2017) is another popular technique, which reweights the loss based on the classifiers’ output probabilities. Our method can be viewed as a type of loss reweighting where the weights are binary and generated for each sample stochastically. Furthermore, PLM is versatile and can be applied on top of existing reweighting and resampling approaches.

**Partial Label Learning** Our approach takes inspiration from partial label learning, in which learning is performed on data with incomplete, or missing, labels (Bucak et al., 2011; Sun et al., 2010). Recently, (Durand et al., 2019) proposed an effective method for training a convolutional neural network on partial labels. Whereas the goal of partial label learning is to learn from sparsely annotated samples, our work deals with learning completely annotated imbalanced datasets; our method assumes all labels are complete, and we partially mask some labels to reduce the effect of data imbalance on network training.

## 5 CONCLUSION

In this work, we propose a general approach to the long-tail data imbalance problem. Our partial label masking algorithm leverages the ratio of positive and negative samples for each class to greatly improve performance on both infrequent and difficult classes. The method can be used alongside most existing methods for data imbalance and is designed for both single-label and multi-label classification. We evaluate PLM on multiple datasets and perform extensive analysis to verify its effectiveness.

$\lambda$	CIFAR-10	CIFAR-100
1.0	<b>24.86</b>	56.92
0.1	25.08	57.67
0.01	25.64	<b>55.98</b>
0.0 (BCE)	29.43	59.88

Figure 5: Error rates on CIFAR-10 and CIFAR-100 (both with  $\rho = 100$ ). This ablation measures the effect of  $\lambda$ . For these experiments, the initial ratio is the dataset’s ratio  $\bar{r}_{c,1} = r_c$ .

## REFERENCES

- Serhat Selcuk Bucak, Rong Jin, and Anil K Jain. Multi-label learning with incomplete class assignments. In *CVPR 2011*, pp. 2801–2808. IEEE, 2011.
- Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106:249–259, 2018.
- Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *Advances in Neural Information Processing Systems*, pp. 1565–1576, 2019.
- Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9268–9277, 2019.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Chris Drummond, Robert C Holte, et al. C4. 5, class imbalance, and cost sensitivity: why under-sampling beats over-sampling. In *Workshop on learning from imbalanced datasets II*, volume 11, pp. 1–8. Citeseer, 2003.
- Thibaut Durand, Nazanin Mehrasa, and Greg Mori. Learning a deep convnet for multi-label classification with partial labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 647–657, 2019.
- Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.
- Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6047–6056, 2018.
- Hui Han, Wen-Yuan Wang, and Bing-Huan Mao. Borderline-smote: a new over-sampling method in imbalanced data sets learning. In *International conference on intelligent computing*, pp. 878–887. Springer, 2005.
- Haibo He and Eduardo A Garcia. Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering*, 21(9):1263–1284, 2009.
- Haibo He, Yang Bai, Eduardo A Garcia, and Shutao Li. Adasyn: Adaptive synthetic sampling approach for imbalanced learning. In *2008 IEEE international joint conference on neural networks (IEEE world congress on computational intelligence)*, pp. 1322–1328. IEEE, 2008.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5375–5384, 2016.
- Chen Huang, Yining Li, Change Loy Chen, and Xiaoou Tang. Deep imbalanced learning for face recognition and attribute prediction. *IEEE transactions on pattern analysis and machine intelligence*, 2019.

- Nathalie Japkowicz and Shaju Stephen. The class imbalance problem: A systematic study. *Intelligent data analysis*, 6(5):429–449, 2002.
- Salman Khan, Munawar Hayat, Syed Waqas Zamir, Jianbing Shen, and Ling Shao. Striking the right balance with uncertainty. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 103–112, 2019.
- Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Andreas Veit, Serge Belongie, Victor Gomes, Abhinav Gupta, Chen Sun, Gal Chechik, David Cai, Zheyun Feng, Dhyanesh Narayanan, and Kevin Murphy. Openimages: A public dataset for large-scale multi-label and multi-class image classification. Dataset available from <https://github.com/openimages>, 2017.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Yann LeCun, Corinna Cortes, and Chris Burges. Mnist handwritten digit database, 1998. 1998.
- 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 *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pp. 2980–2988, 2017.
- Xu-Ying Liu, Jianxin Wu, and Zhi-Hua Zhou. Exploratory undersampling for class-imbalance learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(2): 539–550, 2008.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119, 2013.
- Sankha Subhra Mullick, Shounak Datta, and Swagatam Das. Generative adversarial minority oversampling. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1695–1704, 2019.
- Yu-Yin Sun, Yin Zhang, and Zhi-Hua Zhou. Multi-label learning with weak label. In *Twenty-fourth AAAI conference on artificial intelligence*, 2010.
- Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. In *Advances in Neural Information Processing Systems*, pp. 7029–7039, 2017.

## A APPENDIX

In the appendix we include a description of different multi-label evaluation metrics used (section A.1) and additional details on the imbalanced MultiMNIST dataset (section A.2). Furthermore, we include additional results on the MSCOCO and CIFAR datasets as well as an ablation study on the initialization of the ratio (section A.3).

### A.1 METRICS

In this section, we explain the metrics used in our multi-label classification experiments. In an evaluation set, there are  $N$  images with corresponding labels  $\{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(N)}\}$  where  $\mathbf{y}^{(i)} = [y_1^{(i)}, \dots, y_C^{(i)}]$  is the  $i^{\text{th}}$  binary label vector; note that multiple elements in this label vector may be non-zero in multi-label classification. For each sample, the classifier predicts class probabilities to which a threshold (0.5 in our experiments) is applied to obtain binary prediction vectors  $\{\hat{\mathbf{y}}^{(1)}, \dots, \hat{\mathbf{y}}^{(N)}\}$ .

**Per-class Precision and Recall** The precision and recall of a classifier on a single class  $c$  is given by

$$\begin{aligned} P(c) &= \frac{\sum_{i=1}^N \mathbb{1} \left[ y_c^{(i)} = \hat{y}_c^{(i)} \right]}{\sum_{i=1}^N \mathbb{1} \left[ \hat{y}_c^{(i)} = 1 \right]}, \\ R(c) &= \frac{\sum_{i=1}^N \mathbb{1} \left[ y_c^{(i)} = \hat{y}_c^{(i)} \right]}{\sum_{i=1}^N \mathbb{1} \left[ y_c^{(i)} = 1 \right]}, \end{aligned} \tag{9}$$

respectively, where  $\mathbb{1} [\dots]$  is the indicator function. To obtain the reported metric, we average the precision (or recall) over all classes:

$$\begin{aligned} \text{Precision} &= \frac{1}{C} \sum_{c=1}^C P(c), \\ \text{Recall} &= \frac{1}{C} \sum_{c=1}^C R(c). \end{aligned} \tag{10}$$

In these metrics, each class is treated equally regardless of the number of samples in the test set.

**F1-score** The F1-score is the harmonic mean between precision and recall:

$$F_1(c) = \frac{2P(c)R(c)}{P(c) + R(c)}. \tag{11}$$

The final F1-score is averaged over all classes:

$$\text{F1-score} = \frac{1}{C} \sum_{c=1}^C F_1(c). \tag{12}$$

**0-1 Accuracy** This metric measures how often the network is able to output predictions which exactly match the ground-truth labels. It is computed as follows:

$$\text{0-1 Accuracy} = \frac{1}{N} \sum_{i=1}^N \mathbb{1} \left[ \mathbf{y}^{(i)} = \hat{\mathbf{y}}^{(i)} \right]. \tag{13}$$

Since this metric is averaged over all samples, as opposed to all classes, it tends to be biased towards more frequent classes. Therefore, this metric tends to be a poor measure of model performance if there is imbalance present in the test set.

## A.2 IMBALANCED MULTIMNIST DATASET

**Dataset** We construct the Imbalanced MultiMNIST dataset by superimposing two MNIST LeCun et al. (1998) digits into a single image. We begin by sampling the MNIST training set similarly to how Cui et al. (2019) create the Imbalanced CIFAR dataset with an imbalance of  $\rho = 100$ . These images are padded to become  $32 \times 32$  and shifted randomly by  $-6$  to  $6$  pixels, both vertically and horizontally. Then, each image is randomly paired with another digit in the training set, which is also padded and shifted randomly. These two digits are superimposed upon each other to create a  $32 \times 32$  image with two digits. Due to cooccurrence of the same digit, within a sample, the final imbalance of the dataset becomes 90.33.

For evaluation, we select all samples in the MNIST test set and perform the same padding and shifting operations. For each digit, we select other random test digits that have different labels, which are superimposed to make 9 samples with two digits each. This results in 90000 test samples. The test set is roughly balanced, with all classes having a similar number of samples. The number of samples for each class can be found in Table 4.

Digit	Train	Test
0	9485	17840
1	6167	19080
2	3967	18256
3	2476	18080
4	1508	17856
5	875	17136
6	538	17664
7	331	18224
8	164	17792
9	105	18072
Total	14694	90000

Table 4: The number of samples for each digit in the Imbalanced MultiMNIST dataset.

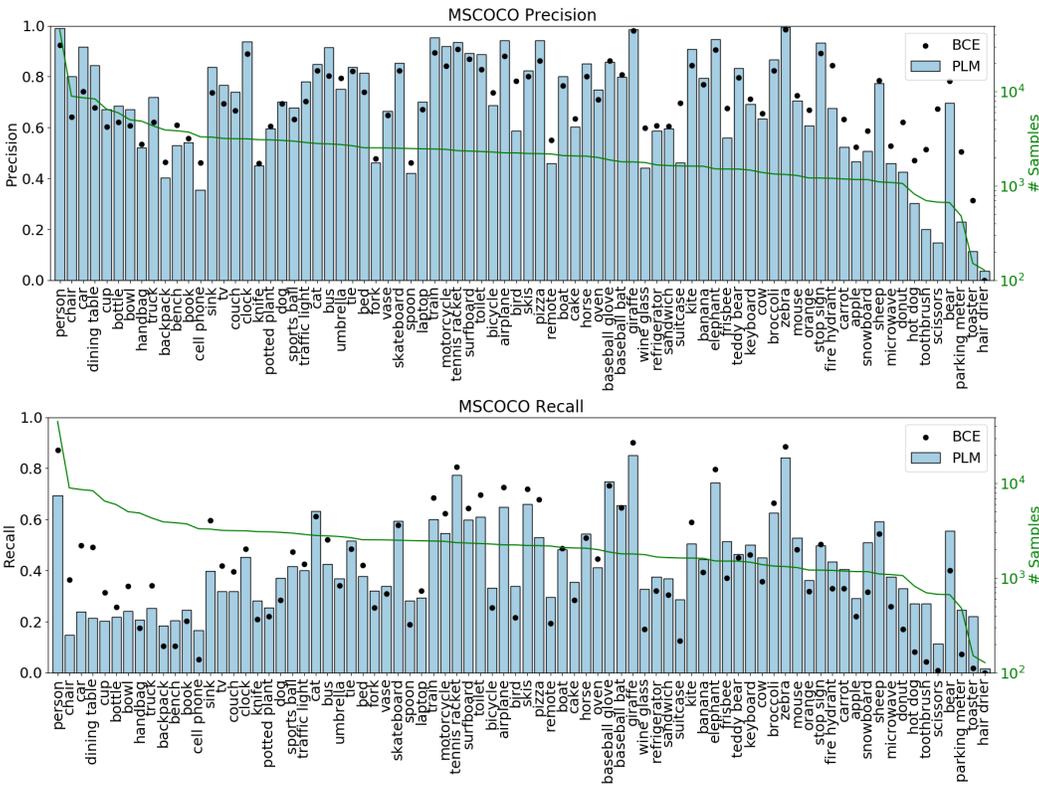


Figure 6: The class-wise precision and recall scores on the MSCOCO dataset. The classes are ordered based on the number of samples in the training set. The bars represent the results achieved by using the PLM method, and the points are the results when using BCE.

### A.3 ADDITIONAL RESULTS

**Additional MSCOCO Results** We include more results on the MSCOCO dataset. Figure 6 contains the class-wise precision and recall scores on the MSCOCO dataset. We find that using PLM leads to an large improvement in recall for the tail classes, and a slight improvement in precision for the head classes (7.86% improvement on the 10 most frequent classes). This is in line with the observation that PLM reduces over-prediction on the head classes while reducing under-prediction on the tail classes (as well as some difficult classes).

**Additional CIFAR Results** In Tables 5 and 6 we present results on CIFAR-10 and CIFAR-100 for the imbalance ratios  $\rho = 20, 50, 200$ . Once again, PLM leads to an improvement over CB for all

Imbalance	200		50		20	
	Overall	K=5	Overall	K=5	Overall	K=5
BCE	34.88	55.84	23.34	35.52	16.37	22.54
Focal	36.28	58.80	23.38	35.22	16.53	23.00
CB	31.11	-	20.73	-	15.64	-
LDAM-DRW	-	-	-	-	-	-
PLM (BCE)	<b>29.67</b>	44.65	20.50	28.66	<b>15.21</b>	19.92
PLM (Focal)	29.79	<b>44.54</b>	<b>20.19</b>	<b>27.94</b>	<b>15.21</b>	<b>19.60</b>

Table 5: Error rates for single-label image classification on imbalanced CIFAR-10. We measure the average error across all classes (Overall) as well as the error on the 5 classes with the fewest number of samples (K=5).

Imbalance	200		50		20	
	Overall	K=50	Overall	K=50	Overall	K=50
BCE	63.92	86.10	55.25	72.84	48.66	59.98
Focal	64.67	87.84	55.71	74.60	48.62	63.86
CB	63.77	-	54.68	-	47.41	-
LDAM-DRW	-	-	-	-	-	-
PLM (BCE)	<b>61.37</b>	80.79	52.49	67.07	46.15	57.26
PLM (Focal)	61.68	<b>79.43</b>	<b>52.09</b>	<b>65.39</b>	<b>46.11</b>	<b>56.28</b>

Table 6: Error Rates for single-label image classification on imbalanced CIFAR-100. We measure the average error across all classes (Overall) as well as the error on the 50 classes with the fewest number of samples (K=50).

Initial Ratio	CIFAR-10	CIFAR-100
$\min\{r_c\}$	28.82	73.46
$\max\{r_c\}$	30.74	56.73
$\text{mean}\{r_c\}$	25.65	56.08
$r_c$	<b>24.86</b>	<b>55.98</b>

Table 7: Ablations on CIFAR-10 and CIFAR-100 (both with  $\rho = 100$ ) measuring the effect of the initial ratio.  $\lambda = 1$  and  $\lambda = 0.01$  for CIFAR-10 and CIFAR-100 respectively.

levels of imbalance. One interesting result is that focal loss (with and without PLM) tends to perform best with smaller levels of imbalance ( $\rho \leq 50$ ) and underperforms when compared to BCE for larger levels of imbalance. This is consistent with our multi-label experiments, which are on datasets with relatively high levels of imbalance.

**Ablation on the initialization of the ratio** We analyse the importance of the initial ratio,  $\bar{r}_{c,1}$ , for our algorithm. We run ablations on CIFAR-10 and CIFAR-100 ( $\rho = 100$ ), where we initialize the ratio with various values. The results for this ablation can be found in Table A.3. We find that initializing the ratio with the dataset’s ratio ( $r_c$ ) or the mean of the dataset’s ratio ( $\text{mean}\{r_c\}$ ) leads to the lowest error. We find that initializing with the maximum or minimum ratio of the dataset tends to lead to poor results. Since the datasets are heavily imbalanced, initializing with either of these ratios leads to too many masked labels which adversely impacts the network’s training. We find that initializing the ratio with the dataset’s ratio is necessary to achieve good performance on the large-scale dataset MSCOCO. This may be the result of the larger amount of imbalance in MSCOCO or the presence of very difficult classes which have many samples (i.e. a difficult class with many positive samples may require a larger ratio than an easy class with fewer samples).