



Transfer and share: semi-supervised learning from long-tailed data

Tong Wei¹ · Qian-Yu Liu² · Jiang-Xin Shi² · Wei-Wei Tu³ · Lan-Zhe Guo²

Received: 26 May 2022 / Revised: 11 August 2022 / Accepted: 12 September 2022

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2022

Abstract

Long-Tailed Semi-Supervised Learning (LTSSL) aims to learn from class-imbalanced data where only a few samples are annotated. Existing solutions typically require substantial cost to solve complex optimization problems, or class-balanced undersampling which can result in information loss. In this paper, we present the TRAS (TRANSFER and Share) to effectively utilize long-tailed semi-supervised data. TRAS transforms the imbalanced pseudo-label distribution of a traditional SSL model via a delicate function to enhance the supervisory signals for minority classes. It then transfers the distribution to a target model such that the minority class will receive significant attention. Interestingly, TRAS shows that more balanced pseudo-label distribution can substantially benefit minority-class training, instead of seeking to generate accurate pseudo-labels as in previous works. To simplify the approach, TRAS merges the training of the traditional SSL model and the target model into a single procedure by sharing the feature extractor, where both classifiers help improve the representation learning. According to extensive experiments, TRAS delivers much higher accuracy than state-of-the-art methods in the entire set of classes as well as minority classes. Code for TRAS is available at <https://github.com/Stomach-ache/TRAS>.

Keywords Long-tailed learning · Semi-supervised learning · Pseudo-label distribution · Logit transformation

Editors: Yu-Feng Li and Prateek Jain.

Tong Wei and Qian-Yu Liu have contributed equally to this work.

✉ Tong Wei
weit@seu.edu.cn

✉ Lan-Zhe Guo
guolz@lamda.nju.edu.cn

Extended author information available on the last page of the article

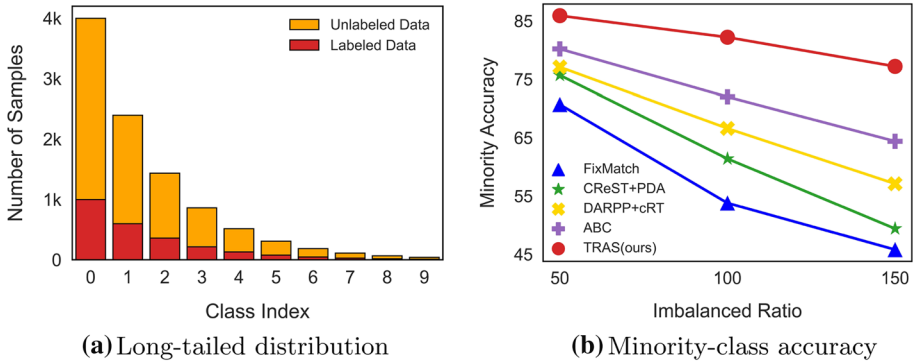


Fig. 1 **a** Long-tailed distribution of the training set under the main setting of CIFAR-10-LT. **b** Performance of minority-class accuracy(%) on CIFAR-10-LT dataset under class imbalance ratio 50, 100, and 150 with 20% of labels available. The proposed TRAS heavily improves the minority-class accuracy

1 Introduction

Deep Neural Networks (DNNs) have been successfully used in many real-world applications (Amodei et al., 2016; He et al., 2016). However, the training of DNNs relies on *large-scale & high-quality* datasets, which has become the core problem in practice. First, large-scale data annotation is highly expensive and only a small amount of labels can be accessed (Li et al., 2019; van Engelen & Hoos, 2020). Second, collected data usually follows long-tailed distribution (He & Garcia, 2009; Liu et al., 2019; Wei & Li, 2019; Wei et al., 2021b, 2022), where only some classes (the majority class) have sufficient training samples while other classes (the minority class) own a few samples as shown in Fig. 1a.

To utilize unlabeled data, Semi-Supervised Learning (SSL) emerges as an interesting solution (Miyato et al., 2018; Tarvainen & Valpola, 2017; Berthelot et al., 2019b; Sohn et al., 2020; Berthelot et al., 2019a; Zhou et al., 2021; Guo et al., 2020). It carries out model assumptions on the data distribution to build a learner to utilize unlabeled samples through selecting confident pseudo-labels. However, it is demonstrated that existing SSL methods tend to produce biased pseudo-labels towards the majority class (Kim et al., 2020), leading to undesirable performance.

Recently, Long-Tailed Semi-Supervised Learning (LTSSL) is proposed to improve the performance of SSL models on long-tailed data. The main ideas of existing LTSSL methods (Kim et al., 2020; Wei et al., 2021a; Lee et al., 2021) are two-fold. One is to improve the quality of pseudo-labels from the perspective of SSL. The other one is to employ class-balanced sampling or post-hoc classifier adjustment to alleviate class imbalance from the long tail perspective. These methods can improve the performance of conventional SSL models. However, the improvements are achieved with the cost of high computational overhead or losing information due to the undersampling of data.

How all data can be efficiently and effectively utilized is the core challenge of LTSSL and the focus of this paper. To this end, we propose a new method called TRAS (TRANSFER and SHARE) which has two key ingredients. Figure 1b showcases the effectiveness of TRAS in the minority class.

First, we compensate for the minority-class training by generating a more balanced pseudo-label distribution. Under the guidance of pseudo-label distribution, DNNs can mine the interaction information between classes to obtain richer information for minority

classes. The idea of learning from label distribution has been explored in previous literature, such as label distribution learning (Geng, 2016; Gao et al., 2017; Wang & Geng, 2019) and knowledge distillation (Xiang et al., 2020; He et al., 2021; Iscen et al., 2021), which however is still underexplored in LTSSL. To generate label distribution, knowledge distillation is a common approach via well-trained teacher models. Such a teacher model is not always available for SSL models because of limited long-tailed labeled data and high computation overhead. Alternatively, we employ a *conventional* SSL model with normally low accuracy on the minority class. This conventional SSL model is able to teach the learning of the student model after applying our proposed logit transformation. This transformation is particularly designed to enhance the minority-class *supervisory signals* without introducing extra computational cost. Subsequently, through training the student model by imitating the enhanced supervisory signals, the minority class will receive significant attention.

Second, we propose to merge the training of teacher and student models as a single procedure to reduce the computational cost. To this end, we use a double-branch neural network with a shared feature extractor and two classifiers for producing the predictions of the teacher and student. The neural network is then trained in an end-to-end way by a joint objective of these two classifiers. In addition to reduce training cost and simplify the approach, we empirically find that both classifiers can help improve the representation learning and learn clear classification boundaries between classes.

Our main contributions are summarized as follows:

1. A new LTSSL method TRAS is proposed, which significantly improves the minority-class training without introducing extra training cost.
2. TRAS transfers pseudo-label distribution from a vanilla SSL network (teacher) to another network (student) via a new logit transformation, instead of trying hard to construct a sophisticated LTSSL teacher model.
3. TRAS reveals the importance of the balancedness of pseudo-label distribution in transfer for LTSSL.
4. TRAS merges the training of teacher and student models by sharing the feature extractor, which simplifies the training procedure and benefits the representation learning.
5. TRAS achieves state-of-the-art performance in various experiments. Particularly, it improves minority-class performance by about 7% in accuracy.

2 Related work

2.1 Semi-supervised learning

Existing SSL methods aim to use unlabeled data to improve the generalization. For this purpose, consistency regularization and entropy minimization have become the most frequently used techniques and demonstrate considerable performance improvements. Specifically, Mean-Teacher (Tarvainen & Valpola, 2017) imposes consistency regularization between the prediction of the current model and the self-ensembled model obtained using exponential moving average. Virtual Adversarial Training (VAT) (Miyato et al., 2018) encourages the model to minimize the discrepancy of model predictions for unlabeled data before and after applying adversarial perturbation. MixMatch (Berthelot et al., 2019b) minimizes the entropy of model predictions by sharpening the pseudo-label distribution.

ReMixMatch (Berthelot et al., 2019a) improves MixMatch by imposing another distribution alignment regularizer and augmentation anchoring. FixMatch (Sohn et al., 2020) merges consistency regularization and entropy minimization by regularizing the prediction for weakly augmented and strongly augmented unlabeled data. However, the above-mentioned methods assume both labeled and unlabeled data is both class-balanced, leading to poor performance on the minority class when working on long-tailed datasets.

2.2 Long-tailed semi-supervised learning

To deal with long-tailed datasets, several LTSSL methods have been proposed. In a nutshell, existing methods aim to select not only confident but also more class-balanced pseudo-labels to improve the generalization for minority classes. For instance, DARP (Kim et al., 2020) proposes to estimate the underlying class distribution of unlabeled data, which is used to regularize the distribution of pseudo-labels. To this end, a convex optimization problem is solved. Additionally, CReST (Wei et al., 2021a) proposes to use class-aware confidence thresholds for selecting more pseudo-labels for the minority class. Recently, ABC (Lee et al., 2021) proposes to use an auxiliary balanced classifier built upon a conventional SSL model by class-balanced undersampling. However, these approaches either suffer from high computational cost or loss of supervisory information. In this work, we propose a new algorithm TRAS, which can fully utilize not only supervised data but also unsupervised data through efficient pseudo-label distribution transfer, and greatly improves the performance of the minority class.

3 Method: TRAS

We now introduce the problem setting in Sect. 3.1 and develop our proposed method TRAS, which consists of two key ingredients described in Sects. 3.2 and 3.3. Figure 2 shows the framework of the proposed TRAS.

3.1 Problem setting

Let $\mathcal{X} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ be a labeled dataset, where $\mathbf{x}_i \in \mathbb{R}^d$ is a training example and $y_i \in \mathbb{R}$ is the corresponding label. We introduce an unlabeled dataset $\mathcal{U} = \{\mathbf{u}_i\}_{i=1}^M$ where $\mathbf{u}_i \in \mathbb{R}^d$ is the unlabeled data point. Following ABC (Lee et al., 2021), we assume that the class distributions of \mathcal{X} and \mathcal{U} are identical. We denote the number of labeled data points of class l as N_l (notice that $\sum_{l=1}^L N_l = N$), assuming that all classes are sorted by cardinality in descending order $N_1 N_2 \cdots N_L$. In LTSSL, we set the fraction of labeled data as $\beta = \frac{N}{N+M}$ and the class imbalance ratio as $\gamma = \frac{N_1}{N_L}$. Usually, we divide the class space into the majority class and the minority class according to their frequencies in the training data. Our goal is to learn a model which generalizes well on both the majority class and the minority class.

Our proposed method, TRAS, consists of a shared feature extractor and two classifiers, providing predictions for the teacher model $P^T(y | \mathbf{x})$ and student model $P^S(y | \mathbf{x})$. There are two key ingredients to TRAS: (1) Learn through imitation, in which the student model imitates the adjusted output of the teacher model, and (2) transfer via sharing weights. In the following, we present technical details of these two ingredients.

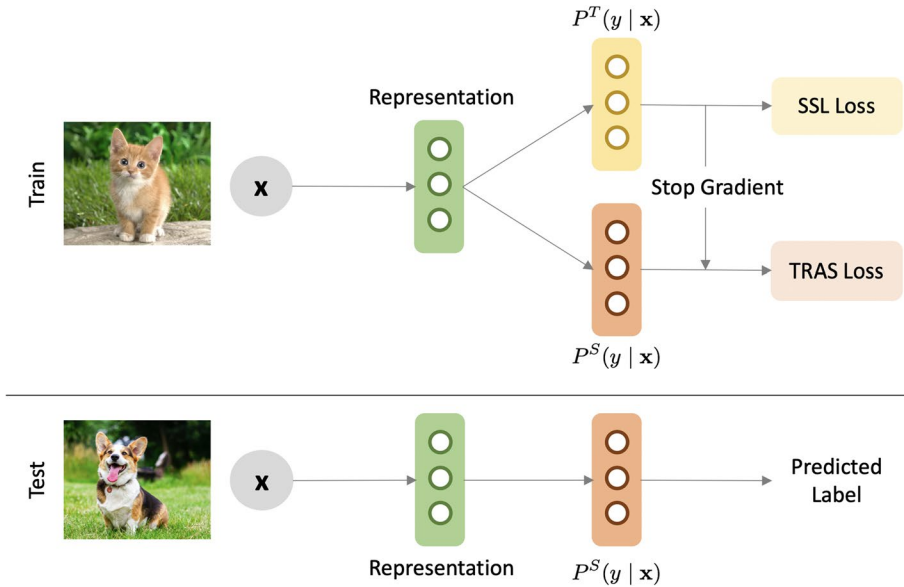


Fig. 2 The TRAS method in diagrammatic form

3.2 Ingredient #1: learn through imitation

Given labeled data, a typical approach is to train a classifier f by optimizing the softmax cross-entropy:

$$\ell_{CE}(y, f(x)) = -\log \frac{e^{f_y(x)}}{\sum_{y' \in [L]} e^{f_{y'}(x)}}. \tag{1}$$

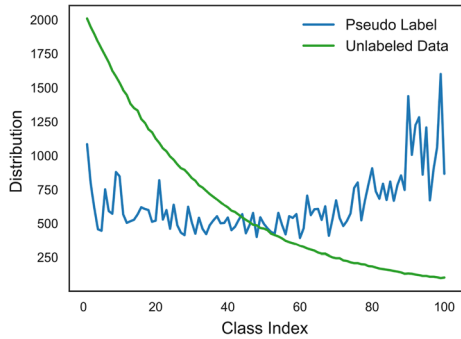
In LTSSL, however, the distribution of labeled data is heavily class-imbalanced, such that the learned classifier would be biased towards the majority class. To improve the training of the minority class, we propose to use the distribution-aware cross-entropy loss:

$$\ell_{DA-CE}(y, f(x)) = -\log \frac{e^{f_y(x) + \tau \cdot \log \pi_y}}{\sum_{y' \in [L]} e^{f_{y'}(x) + \tau \cdot \log \pi_{y'}}}, \tag{2}$$

where π_y is the estimate of class prior $\mathbb{P}(y)$ and $\tau > 0$ is a scaling parameter. By minimizing ℓ_{DA-CE} , it encourages large margins between the true label and other negative labels. Using distribution-aware cross-entropy is not a new idea in the literature of long-tailed learning, such as Logit Adjustment (Menon et al., 2020) and Balanced Softmax (Ren et al., 2020). Interestingly, existing methods show that the scaling parameter τ plays an important role in model training, but it is usually used as a constant, e.g., $\tau = 1$. In the following, we show a new instance-dependent logit scaling method.

In addition to labeled data, we can access to a large amount of unlabeled data to help improve the generalization. In LTSSL, the underlying distribution of unlabeled data is also long-tailed, and conventional SSL methods have shown impaired performance on the minority class. This paper proposes to train the model using pseudo-label

Fig. 3 Comparison of ground-truth label distribution and our generated pseudo-label distribution on CIFAR-100-LT dataset under class imbalance ratio 20 with 40% of labels available



distribution, rather than biased one-hot pseudo-labels. Intuitively, label distribution offers more supervisory signals and can benefit the minority-class training. We generate pseudo-label distribution by first training a vanilla SSL model as the teacher, and then training a student model by imitating the output distribution of teacher model. We opt for minimizing their Kullback-Leibler (KL) divergence:

$$\ell_{\text{KL}}(\tilde{\mathbf{y}}^T, \tilde{\mathbf{y}}^S) = \sum_{l=1}^L \tilde{y}_l^T \log \frac{\tilde{y}_l^T}{\tilde{y}_l^S}, \quad (3)$$

where $\tilde{\mathbf{y}}^T$ and $\tilde{\mathbf{y}}^S$ are output probabilities of the teacher and student model respectively, which illustrate the implicit information of label distribution.

Note that the teacher model is trained via a conventional SSL algorithm and the produced pseudo-label distribution is still biased towards the majority class. To further enhance supervisory signals for the minority class, we present a new logit transformation to adjust the output of the teacher model. Specifically, for sample \mathbf{x} , we transform its pseudo-label distribution as follows:

$$\tilde{\mathbf{y}}^T = \text{softmax}(\phi(\mathbf{z}^T)) = \text{softmax}(\mathbf{z}^T - \tau(\hat{y}) \cdot \log \boldsymbol{\pi}), \quad (4)$$

where \mathbf{z}^T is the output logits and \hat{y} is the pseudo-label of \mathbf{x} . In this way, the pseudo-label distribution of unlabeled data is more balanced. We demonstrate the generated label distribution in Fig. 3.

Notably, different from previous works that treat τ as a constant to scale the output logits, we use τ as a function of pseudo-labels. Concretely, given the pseudo-label \hat{y} , we define $\tau(\hat{y}) = A \cdot \alpha_{\hat{y}} + B$, where $\boldsymbol{\alpha} = \text{softmax}(-\log \boldsymbol{\pi})$ is a \hat{y} -dependent function, A and B are constants. This is because adjusting pseudo-label distribution to over-compensate the minority class can be harmful to the majority class. By employing the \hat{y} -dependent logit transformation function, we can alleviate this problem by flattening the label distribution of predicted minority-class samples more aggressively than other samples. In experiments, we simply set $A = B = 2$. Applying the proposed logit transformation generates a more balanced pseudo-label distribution to improve the training of the minority class as in Fig. 3.

Putting together the objectives for labeled and unlabeled data, we minimize the loss function for TRAS as follows:

$$\mathcal{L}_{\text{TRAS}} = \underbrace{\sum_{i=1}^N \ell_{\text{DA-CE}}(y_i, z_i^S)}_{\text{supervised loss}} + \underbrace{\sum_{j=1}^M \mathbb{I}(\max(\tilde{y}_j^S) \geq t) \ell_{\text{KL}}(\tilde{y}_j^T, \tilde{y}_j^S)}_{\text{unsupervised loss}}. \tag{5}$$

Here, $\mathbb{I}(\cdot)$ represents the indicator function, t denotes the confidence threshold and we adopt the common setup $t = 0.95$ for confident pseudo-labels from the student.

In this way, pseudo-label distribution can naturally describe the implicit information between labels. By applying the logit transformation, the distribution encodes more informative supervisory signals for the minority class, thus the student can alleviate data scarcity for minority classes.

3.3 Ingredient #2: transfer via sharing weights

Learning through imitating the teacher model can significantly compensate for the training of the minority class, however, it requires to train two separate DNNs sequentially, which is computational expensive in SSL.

To reduce the time consumption and simplify the approach, we propose to merge the training of teacher and student models into a single training procedure. In other words, the teacher and student share the feature extractor network. We further partition the parameter space into three disjoint subsets; (1) Let $\psi(\mathbf{x})$ be a feature extractor for \mathbf{x} . (2) Let $f^T(\psi(\mathbf{x}))$ denote a teacher classifier and its prediction \tilde{y}^T . (3) Similarly, let $f^S(\psi(\mathbf{x}))$ denote a student classifier and its prediction \tilde{y}^S . Subsequently, let us define:

$$z^T = \text{stop_gradient}(f^T(\psi(\mathbf{x}))), \tag{6}$$

which is the output logits of the teacher model except that its gradient will not be calculated to update the teacher model’s classifier weights. Recall that function $\phi(\cdot)$ acts as a logit transformer of z^T , we then consider:

$$\mathcal{L}_{\text{TRAS}} = \underbrace{\sum_{i=1}^N \ell_{\text{DA-CE}}(y_i, z_i^S)}_{\text{supervised loss}} + \underbrace{\sum_{j=1}^M \mathbb{I}(\max(\tilde{y}_j^S) \geq t) \ell_{\text{KL}}(\text{softmax}(\phi(z_j^T)), \tilde{y}_j^S)}_{\text{unsupervised loss}}, \tag{7}$$

as the joint objective. Note that the teacher and student share a single feature extractor, it only adds a linear classifier to the conventional SSL model, which incurs negligible training cost.

Let \mathcal{L}_{SSL} denote the loss for a conventional SSL method, the total loss function that TRAS optimizes is:

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{TRAS}} + \mathcal{L}_{\text{SSL}}. \tag{8}$$

Particularly, if FixMatch is employed as the teacher model, \mathcal{L}_{SSL} consists of a cross-entropy loss on labeled data and a consistency regularization on unlabeled data. Specifically, we have:

$$\mathcal{L}_{\text{SSL}} = \underbrace{\sum_{i=1}^N \ell_{\text{CE}}(y_i, z_i^T)}_{\text{supervised loss}} + \underbrace{\sum_{j=1}^M \mathbb{1}(\max(z_j^T) \geq t) \ell_{\text{CE}}(\hat{y}_j, \bar{z}_j^T)}_{\text{unsupervised loss}}, \quad (9)$$

where z^T and \bar{z}^T are the output logits for weak and strong data augmentation, $\hat{y} = \arg \max_i z_i^T$ represents the pseudo-label for unlabeled data. In inference, we use the student classifier f^S to predict the label.

3.4 Connection to previous work

One may note that the basic idea of our TRAS can transfer knowledge distribution from a vanilla teacher model to a student model that has good generalization for the minority class. The technique is related to knowledge distillation which has been explored in some recent long-tailed learning works. For instance, LFME (Xiang et al., 2020) proposes to train the student model via distilling multiple teachers trained on less imbalanced datasets. DiVE (He et al., 2021) shows that flattening the output distribution of the teacher model using a constant temperature parameter can help the learning of minority classes. CBD (Isen et al., 2021) distills features from the teacher to the student and shows that it can improve the learned representation of the minority class. Last but not least, xERM (Zhu et al., 2022) obtains an unbiased model by properly adjusting the weights between empirical loss and knowledge distillation loss.

In contrast to previous works that aim to solve supervised long-tailed learning, this paper studies semi-supervised long-tailed learning, where the amount of labeled data is much more limited. Moreover, previous works need to train teacher models via well-established long-tailed learning methods. However, our method TRAS only needs a vanilla SSL model as a teacher. Additionally, these methods have multiple-stage training procedures, but our method is simpler and can be trained in an end-to-end way.

4 Experiments

We conduct experiments on long-tailed version of CIFAR-10, CIFAR-100, and SVHN, in comparison with state-of-the-art LTSSL methods. We then perform hyper-parameter sensitivity studies and ablation studies to better understand our proposed TRAS.

4.1 Experimental setup

4.1.1 Datasets

We conduct experiments on common datasets long-tailed CIFAR-10 (CIFAR-10-LT), long-tailed CIFAR-100 (CIFAR-100-LT) and long-tailed SVHN (SVHN-LT) to evaluate our method. Without loss of generality, for imbalanced SSL settings, we randomly resample the datasets to meet the assumption that the distribution of labeled and unlabeled samples is consistent. We set the ratio of the class imbalance as γ ($\gamma = \frac{N_L}{N_U}$) and the number of

Table 1 Overall accuracy(%) / minority-class accuracy(%) under the main setting

Algorithm	CIFAR-10-LT	SVHN-LT	CIFAR-100-LT
	$\gamma = 100, \beta = 20\%$	$\gamma = 100, \beta = 20\%$	$\gamma = 20, \beta = 40\%$
Vanilla	55.3 _{±1.30} /33.9 _{±1.88}	77.0 _{±0.67} /63.3 _{±1.25}	40.1 _{±1.15} /25.2 _{±0.95}
VAT	55.3 _{±0.88} /28.2 _{±1.55}	81.3 _{±0.47} /68.2 _{±0.88}	40.4 _{±0.34} /24.8 _{±0.38}
BALMS	70.7 _{±0.59} /69.8 _{±1.03}	87.6 _{±0.53} /85.0 _{±0.67}	50.2 _{±0.54} /42.9 _{±1.03}
FixMatch	72.3 _{±0.33} /53.8 _{±0.63}	88.0 _{±0.30} /79.4 _{±0.54}	51.0 _{±0.20} /32.8 _{±0.41}
w/ CReST+PDA	76.6 _{±0.46} /61.4 _{±0.85}	89.1 _{±0.69} /81.7 _{±1.18}	51.6 _{±0.29} /36.4 _{±0.46}
w/ DARP	73.7 _{±0.98} /57.0 _{±2.12}	88.6 _{±0.19} /80.5 _{±0.54}	51.4 _{±0.37} /33.9 _{±0.77}
w/ DARP+cRT	78.1 _{±0.89} /66.6 _{±1.55}	89.9 _{±0.44} /83.5 _{±0.61}	54.7 _{±0.46} /41.2 _{±0.42}
w/ ABC	81.1 _{±0.82} /72.0 _{±1.77}	92.0 _{±0.38} /87.9 _{±0.73}	56.3 _{±0.19} /43.4 _{±0.42}
w/ TRAS(ours)	84.3 _{±0.23} / 82.2 _{±0.44}	93.4 _{±0.51} / 92.5 _{±0.26}	58.5 _{±0.17} / 50.3 _{±0.22}

The best results are indicated in bold

labeled data points of class l as N_l , where $N_l = N_1 * \gamma^{-\frac{l-1}{L-1}}$ and M_l for the unlabeled. Specifically, we set $N_1 + M_1 = 5000$, $L = 10$ for CIFAR-10-LT and SVHN-LT, $N_1 + M_1 = 500$, $L = 100$ for CIFAR-100-LT respectively.

Following the previous work (Lee et al., 2021), we evaluate the classification performance with imbalance ratio $\gamma = 100$ and 150 for CIFAR-10-LT and SVHN-LT and $\gamma = 20$ and 30 for CIFAR-100-LT. The ratio of labeled data β is 10%, 20% and 30% for CIFAR-10-LT and SVHN-LT, 20%, 40% and 50% for CIFAR-100-LT. Since the test set remains balanced, overall accuracy, minority-class accuracy, and Geometric Mean scores (GM) (Branco et al., 2016) with class-wise sensitivity are three main metrics to validate the proposed method.

4.1.2 Setup

We implement our method with FixMatch over the backbone of Wide ResNet-28-2 (Zagoruyko & Komodakis, 2016). Our method is compared with the supervised baseline, long-tailed supervised learning methods, and long-tailed semi-supervised learning methods, denoted by (a) Vanilla; (b) VAT (Miyato et al., 2018) and FixMatch (Sohn et al., 2020; c) BALMS (Ren et al., 2020), classifier Re-Training (cRT) (Kang et al., 2020; d) DARP (Kim et al., 2020), CReST (Wei et al., 2021a), ABC (Lee et al., 2021). We set the hyperparameters by following FixMatch and train the neural networks for 500 epochs with 500 mini-batches in each epoch, with the batch size of 64, using Adam optimizer (Kingma & Ba, 2015). The learning rate is 0.002 with a decay rate of 0.999. We start optimizing TRAS after training FixMatch for 10 epochs. For all experiments, we report the mean and standard deviation of test accuracy over multiple runs.

4.2 Experimental results

First, the performance of the algorithms compared under the main setting is in Table 1. Results of related methods are borrowed from ABC (Lee et al., 2021). It can be seen that our method achieves the best performance, and the improvement on the minority class is impressive. It is known that normal SSL methods such as VAT and FixMatch perform

Table 2 Overall accuracy(%) / minority-class accuracy(%) for CIFAR-10-LT. Two imbalance ratios γ and three labeled data ratios β are evaluated

CIFAR-10-LT			
Algorithm	$\gamma = 100, \beta = 10\%$	$\gamma = 100, \beta = 30\%$	$\gamma = 150, \beta = 20\%$
FixMatch	70.0 _{±0.59} /48.9 _{±1.04}	74.9 _{±0.63} /58.2 _{±1.28}	68.5 _{±0.60} /45.8 _{±1.15}
w/ CReST+PDA	73.9 _{±0.40} /58.9 _{±1.14}	77.6 _{±0.73} /64.0 _{±1.39}	70.0 _{±0.82} /49.4 _{±1.52}
w/ DARP+cRT	74.6 _{±0.98} /59.2 _{±2.12}	79.0 _{±0.25} /67.7 _{±0.95}	73.2 _{±0.85} /57.1 _{±1.13}
w/ ABC	77.2 _{±1.60} /65.7 _{±2.85}	81.5 _{±0.29} /72.9 _{±0.96}	77.1 _{±0.46} /64.4 _{±0.92}
w/ TRAS(ours)	82.1 _{±0.41} / 78.6 _{±0.51}	85.0 _{±0.46} / 83.0 _{±0.85}	81.7 _{±0.44} / 77.2 _{±0.92}

The best results are indicated in bold

Table 3 Overall accuracy(%) / minority-class accuracy(%) on SVHN-LT. Two imbalance ratios γ and three labeled data ratios β are evaluated

SVHN-LT			
Algorithm	$\gamma = 100, \beta = 10\%$	$\gamma = 100, \beta = 30\%$	$\gamma = 150, \beta = 20\%$
FixMatch	88.5 _{±0.23} /80.3 _{±0.42}	88.7 _{±0.36} /80.7 _{±0.65}	85.6 _{±0.17} /74.6 _{±0.43}
w/ CReST+PDA	89.2 _{±0.43} /81.7 _{±0.95}	89.9 _{±0.36} /83.0 _{±0.37}	86.7 _{±0.89} /76.7 _{±1.70}
w/ DARP+cRT	89.3 _{±0.33} /83.9 _{±0.47}	90.7 _{±0.28} /84.8 _{±0.37}	88.0 _{±0.74} /80.1 _{±1.88}
w/ ABC	92.3 _{±0.38} /88.7 _{±0.92}	92.3 _{±0.34} /88.3 _{±0.49}	91.2 _{±0.15} /86.2 _{±0.15}
w/ TRAS(ours)	93.2 _{±0.22} / 92.5 _{±0.10}	93.9 _{±0.20} / 93.4 _{±0.33}	92.1 _{±0.36} / 91.1 _{±0.86}

The best results are indicated in bold

Table 4 Overall accuracy(%) / minority-class accuracy(%) on CIFAR-100-LT. Two imbalance ratios γ and three labeled data ratios β are evaluated

CIFAR-100-LT			
Algorithm	$\gamma = 20, \beta = 20\%$	$\gamma = 20, \beta = 50\%$	$\gamma = 30, \beta = 40\%$
FixMatch	46.1 _{±0.23} /26.6 _{±0.34}	52.3 _{±0.54} /34.7 _{±0.80}	47.6 _{±0.09} /27.6 _{±0.21}
w/ CReST+PDA	46.7 _{±0.49} /29.3 _{±0.54}	52.7 _{±0.06} /37.4 _{±0.37}	48.5 _{±0.06} /30.0 _{±0.04}
w/ DARP+cRT	48.9 _{±0.11} /33.5 _{±0.17}	55.9 _{±0.43} /43.5 _{±1.28}	51.3 _{±0.29} /36.4 _{±0.50}
w/ ABC	49.7 _{±0.40} /34.6 _{±0.76}	58.3 _{±0.74} /46.7 _{±1.12}	53.6 _{±0.35} /38.8 _{±0.69}
w/ TRAS(ours)	51.6 _{±0.30} / 41.8 _{±0.61}	60.3 _{±0.75} / 53.5 _{±0.91}	55.5 _{±0.63} / 46.5 _{±0.62}

The best results are indicated in bold

unsatisfactorily on the minority class because pseudo-labels of unlabeled data are affected by the biased model thus hindering the learning of minority classes. Our method significantly improves the performance on the minority class by exploiting knowledge transfer to generate balanced label distribution, which conveys more implicit information than the one-hot pseudo-labels used in most previous LTSSL works. Moreover, our standard deviation is lower than other LTSSL methods, showing the superior stability of TRAS.

Table 5 Results of GM(%) under the main setting

Algorithm	CIFAR-10-LT	SVHN-LT	CIFAR-100-LT
	$\gamma = 100, \beta = 20\%$	$\gamma = 100, \beta = 20\%$	$\gamma = 20, \beta = 40\%$
FixMatch	62.0	87.3	38.5
w/ CReST+PDA	74.4	88.6	42.3
w/ DARP	71.5	87.6	40.4
w/ DARP+cRT	76.7	89.8	47.0
w/ ABC	80.5	91.8	49.0
w/ TRAS(ours)	81.9	93.4	54.0

The best results are indicated in bold

Table 6 Averaged class recall (%) on ImageNet-127

	ImageNet-127 32 × 32	ImageNet-127 64 × 64
FixMatch	29.7	42.3
w/ DARP	30.5	42.5
w/ DARP + cRT	39.7	51.0
w/ CReST+	32.5	44.7
w/ CoSSL	43.7	53.8
w/ TRAS(ours)	46.2	54.1

The best results are indicated in bold

To further validate the effectiveness of our method, we report the performance on various settings. The results on CIFAR-10-LT, SVHN-LT and CIFAR-100-LT are reported in Tables 2, 3 and 4. TRAS outperforms other methods in all cases with respect to both overall accuracy and minority-class accuracy. Particularly, TRAS achieves about 10%, 5%, 7% improvements in the minority class on three datasets. Moreover, TRAS is more robust to class imbalance. As the imbalance ratio increases, existing methods severely deteriorate their performance, while the accuracy of our method drops slightly.

To evaluate whether our method TRAS performs balanced prediction for all classes, we measure its performance using Geometric Mean scores (GM) of class-wise accuracy. The results in Table 5 demonstrate that the proposed algorithm yields the best and most balanced performance in all classes. Additionally, TRAS achieves more significant performance improvement on the large dataset (CIFAR-100-LT).

Additionally, we also evaluate TRAS on the more practical and challenging ImageNet127 dataset. ImageNet127 was crafted by CReST (Wei et al., 2021a) for LTSSL. It is a naturally imbalanced dataset with imbalance ratio 286 by grouping the 1000 classes of ImageNet into 127 classes based on the WordNet hierarchy. Due to limited resources, we are not able to conduct experiments on ImageNet127 with the full resolution.¹ Instead, we follow CoSSL (Fan et al., 2022) which down-samples the original images from

¹ One run of vanilla FixMatch on ImageNet127 on a single NVIDIA Tesla V100 takes 10676.5 h which is about 444 days.

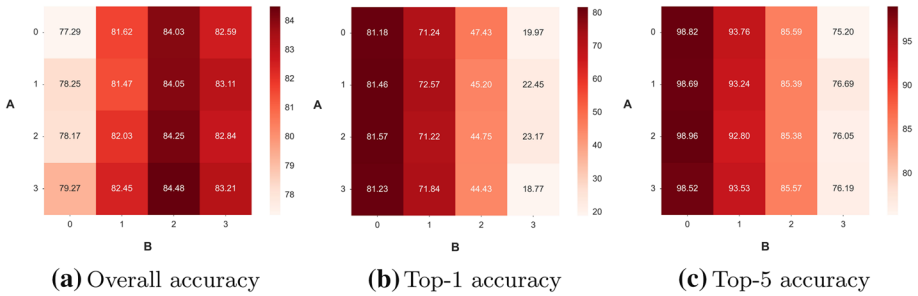


Fig. 4 The impact of values of A and B on CIFAR-10-LT under class imbalance ratio 100 with 20% of labels available. **a** Overall accuracy(%) of TRAS; **b** Top-1 accuracy(%) of pseudo-labels of the teacher after transformation; **c** Top-5 accuracy(%) of pseudo-labels of the teacher after transformation

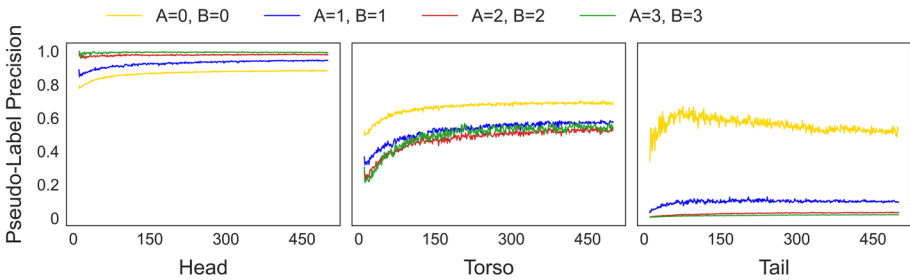


Fig. 5 Comparison of pseudo-label precision by varying the values of A and B on CIFAR-10-LT under class imbalance ratio 100 with 20% of labels available. The x-axis is the number of epochs, and the y-axis is the precision. Classes are divided into head ($\{0, 1, 2\}$), torso ($\{3, 4, 5, 6\}$) and tail ($\{7, 8, 9\}$)

ImageNet127 to smaller images of 32×32 or 64×64 pixels using the box method from Pillow library. We also randomly select 10% training samples as the labeled data. Since the test set is imbalanced as well, averaged class recall is applied as a balanced metric. We compare our method with the recent state-of-the-art approach CoSSL (Fan et al., 2022). From Table 6, we can see that TRAS achieves the best performance for both image size 32 and 64.

4.3 How does pseudo-label distribution impact the performance?

Recall that we use hyper-parameters A and B to control the distribution of pseudo-labels by $\tau(\hat{y}) = A \cdot \alpha_{\hat{y}} + B$, we now analyze their influence on the performance. The results are reported in Fig. 4. We find that B impacts the performance much larger than A , which coincides with our intuition because $\alpha_{\hat{y}} < 1$. It achieves comparable results by setting $A \in \{1, 2, 3\}$. But unlike A , $B = 2$ yields better performance than other values in our

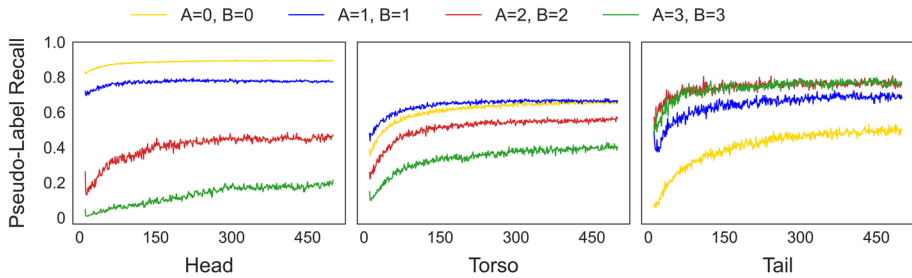


Fig. 6 Comparison of pseudo-label recall by varying the value of A and B on CIFAR-10-LT under class imbalance ratio 100 with 20% of labels. The x-axis is the number of epochs, and the y-axis is the recall

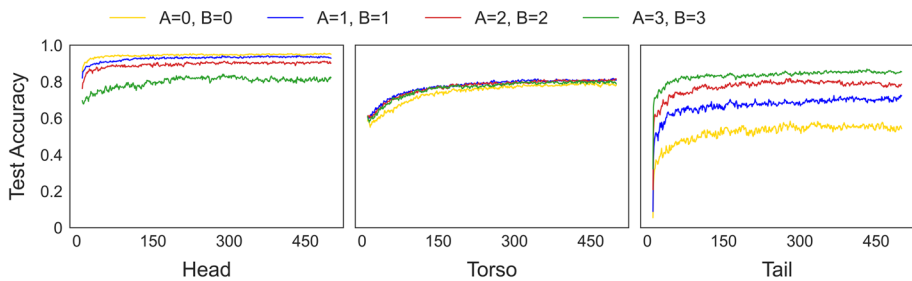


Fig. 7 Comparison of test accuracy by varying the values of A and B on CIFAR-10-LT under class imbalance ratio 100 with 20% of labels. The x-axis is the number of epochs, and the y-axis is the test accuracy

experiments. When setting $A > 3$ and $B > 3$, test accuracy severely deteriorates because of the heavy bias towards minority classes in the pseudo-label distribution.

Interestingly, we find the balancedness of pseudo-label distribution matters much more than the accuracy of pseudo-labels in transfer. As B increases, the top-5 accuracy of pseudo-labels is impaired while the overall accuracy remains competitive. This indicates that class imbalance hurts the performance more than inaccurate pseudo-labels in our approach.

To better understand this phenomenon, we investigate the impact of logit scaling parameters A and B on the quality of pseudo-labels for head, torso, and tail classes separately. As illustrated in Fig. 5, $A = 0, B = 0$ reveals superb performance with high precision. However, in Fig. 6, it shows the worst recall in the tail class. Since $A = 0, B = 0$ means that pseudo-labels are from the conventional SSL model which is biased to the head class, transferring their distribution to a target model does not help the training of tail classes, as shown in Fig. 7.

Instead, by setting $A = 2, B = 2$, it achieves the best performance in overall and tail-class accuracy as reported in Figs. 4a and 7. Notably, it produces high recall yet low precision for tail classes in Figs. 6 and 5. This observation confirms our suspicion that the balancedness of pseudo-label distribution requires more attention than the accuracy of pseudo-labels in knowledge transfer.

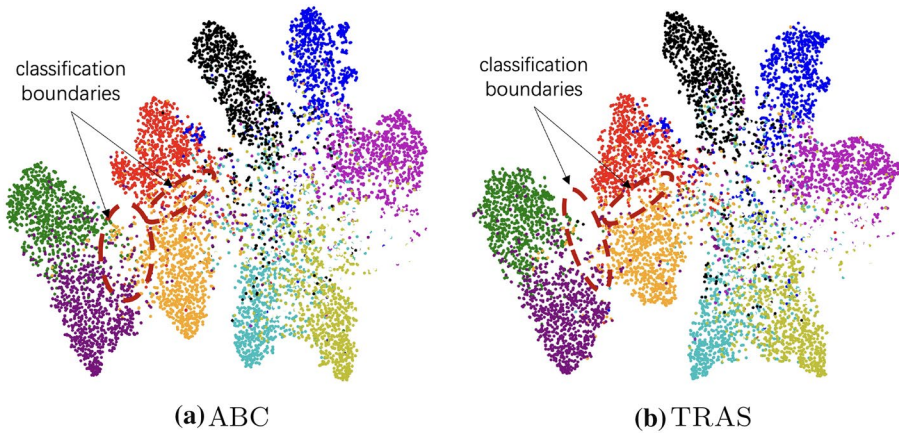


Fig. 8 Results of t-SNE for ABC (Lee et al., 2021) and our TRAS

4.4 Better understanding of TRAS

We analyze TRAS from representation and classification perspectives on CIFAR-10-LT under the main setting. First, we compare the learned representations by ABC and our TRAS via t-distributed stochastic neighbor embedding (t-SNE) (Van der Maaten & Hinton, 2008) in Fig. 8. It can be seen that TRAS has more clear classification boundaries than ABC, which demonstrates that TRAS can distinguish the difference between classes with better representation learning.

Further, to analyze the classification results, we compare the confusion matrices of the prediction on the test set in Fig. 9. Each row represents the ground-truth label and each column represents the prediction by ABC or our TRAS. The value in the i -th row and j -th column is the percentage of samples from the i -th class and predicted as the j -th class. From the results, we can see that our TRAS performs better than ABC in the minority class. Moreover, it is observed that TRAS might misclassify some majority-class samples as the minority-class ones.

Ablation studies. We conduct ablation studies on important parts of our approach under the main setting.

First, this paper applies logit transformation with $A = 0, B = 1$ to the teacher model's prediction on unlabeled data for better performance of the teacher model. By removing logit transformation, the overall accuracy and minority-class accuracy under the main setting turn out to be 83.66% (−0.64%) and 77.54% (−4.66%) on CIFAR-10-LT, respectively.

Second, we modify the distribution-aware cross-entropy for the labeled data to the common cross-entropy loss, leading to 83.41% (−0.89%) and 80.28% (−1.92%) of the overall accuracy and minority-class accuracy. The marginal decline of the performance verifies the effectiveness of the learning through imitation approach.

Finally, we remove the sample mask on unlabeled data of the student model, which means all unlabeled data is used to imitate the teacher. The experiment shows that removing the sample mask decreases the performance slightly, i.e., 83.33% (−0.97%) and 79.60% (−2.59%) for overall and minority-class accuracy respectively. This demonstrates the advantage of selecting more accurate pseudo-labels for the student model.

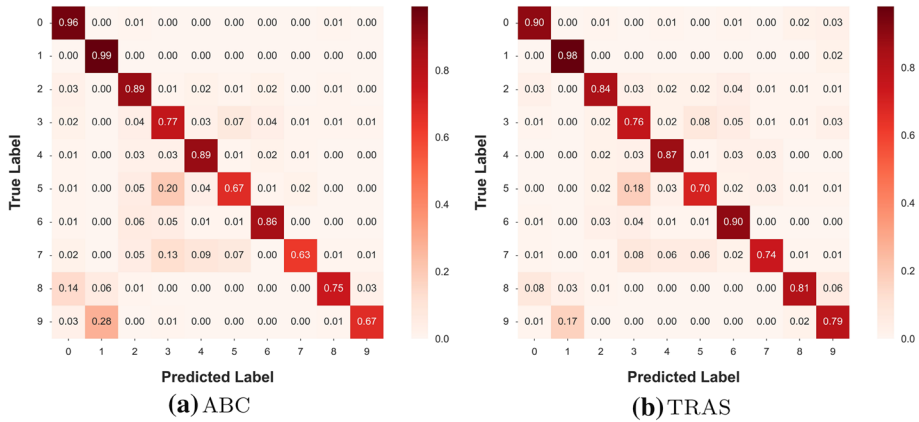


Fig. 9 Confusion matrices of the prediction on the test set of CIFAR-10-LT

4.5 Comparison with two-stage training

We further compare TRAS with two-stage training under the main setting in Table 7. In the two-stage training, we first train a FixMatch model as the teacher and then guide a student model by the teacher with our TRAS. We find that not only can our TRAS save training cost, but it is also better trained than the two-stage approach. This agrees with our expectation that, in two-stage training, the student lays more emphasis on the minority class, while fitting the pseudo-label distribution does not necessarily improve feature learning. Fortunately, two branches in TRAS share the feature learning backbone, which can improve the backbone and classifiers simultaneously.

To further show that double branches of TRAS can both improve feature learning, we stop propagating the gradients of the student branch from affecting the feature learning backbone. In this way, only the teacher model trains the feature extractor network, which is similar to the two-stage training. From the result, we see TRAS performs better, showing that the student can further enhance the feature learning by sharing the backbone with the teacher model.

Table 7 Performance comparison of overall accuracy(%)/minority-class accuracy(%) with two-stage training and TRAS-

Method	CIFAR-10-LT	SVHN-LT	CIFAR-100-LT
	$\gamma = 100, \beta = 20\%$	$\gamma = 100, \beta = 20\%$	$\gamma = 20, \beta = 40\%$
Two-stage	80.6/75.0	91.6/89.6	57.1/47.6
TRAS- ¹	83.2/80.8	92.6/92.1	57.8/50.3
TRAS	84.3/82.2	93.4/92.5	58.5/50.3

The best results are indicated in bold

¹The student does not propagate its gradients to update the feature extractor

5 Conclusion

We introduce TRAS, a new method for LTSSL. TRAS (1) learns from a more class-balanced label distribution to improve the minority-class generalization and (2) partitions the parameter space, enabling transfer via weight sharing of the transformed knowledge learned by the conventional SSL model. Extensive experiments on CIFAR-10-LT, SVHN-LT, and CIFAR-100-LT datasets show that TRAS outperforms state-of-the-art methods on the minority class by a large margin. In the sequel, it would be interesting to extend TRAS to more established SSL methods.

Author Contributions All authors contributed to the study conception and design. Formal development of the method was performed by TW and QYL. The experiments were performed by QYL. The first draft of the manuscript was written by TW. JXS, LZG, and WWT commented on previous versions of the manuscript. All authors read and approved the final manuscript. The corresponding author Tong Wei acquired the funding, and was responsible for the research activity planning and execution.

Funding This work was supported by “the Fundamental Research Funds for the Central Universities (2242022R10014)”.

Availability of data and materials The datasets used in this study are publicly available online.

Declarations

Conflict of interest All authors declare that they have no conflict of interest.

Code availability The code is available at <https://github.com/Stomach-ache/TRAS>.

Consent to participate Not applicable.

Consent for publication Not applicable.

Ethical approval Not applicable.

References

- Amodei, D., Ananthanarayanan, S., Anubhai, R., et al. (2016). Deep speech 2: End-to-end speech recognition in english and mandarin. In *International Conference on Machine Learning*, (pp. 173–182).
- Berthelot, D., Carlini, N., Cubuk, ED., et al. (2019a). Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *International Conference on Learning Representations*.
- Berthelot, D., Carlini, N., Goodfellow, I., et al. (2019). Mixmatch: A holistic approach to semi-supervised learning. *Advances in Neural Information Processing Systems*, 32, 5050–5060.
- Branco, P., Torgo, L., & Ribeiro, R. P. (2016). A survey of predictive modeling on imbalanced domains. *ACM Computing Surveys*, 49(2), 1–50.
- Fan, Y., Dai, D., Kukleva, A., et al. (2022). Cossl: Co-learning of representation and classifier for imbalanced semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Gao, B. B., Xing, C., Xie, C. W., et al. (2017). Deep label distribution learning with label ambiguity. *IEEE Transactions on Image Processing*, 26(6), 2825–2838.
- Geng, X. (2016). Label distribution learning. *IEEE Transactions on Knowledge and Data Engineering*, 28(7), 1734–1748.
- Guo, L.Z., Zhang, Z.Y., Jiang, Y., et al. (2020). Safe deep semi-supervised learning for unseen-class unlabeled data. In *International Conference on Machine Learning* (pp. 3897–3906).

- He, H., & Garcia, E. A. (2009). Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9), 1263–1284.
- He, K., Zhang, X., Ren, S., et al. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, (pp. 770–778).
- He, Y.Y., Wu, J., Wei, X.S. (2021). Distilling virtual examples for long-tailed recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 235–244).
- Iscen, A., Araujo, A., Gong, B., et al. (2021). Class-balanced distillation for long-tailed visual recognition. In *The British Machine Vision Conference*.
- Kang, B., Xie, S., Rohrbach, M. et al. (2020). Decoupling representation and classifier for long-tailed recognition. In *International Conference on Learning Representations*.
- Kim, J., Hur, Y., Park, S., et al. (2020). Distribution aligning refinery of pseudo-label for imbalanced semi-supervised learning. *Advances in Neural Information Processing Systems*, 33, 14567–14579.
- Kingma, D.P., Ba, J. (2015). Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Lee, H., Shin, S., & Kim, H. (2021). Abc: Auxiliary balanced classifier for class-imbalanced semi-supervised learning. *Advances in Neural Information Processing Systems*, 34, 7082–7094.
- Li, Y., Wang, H., Wei, T., et al. (2019). Towards automated semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence* (pp. 4237–4244).
- Liu, Z., Miao, Z., Zhan, X., et al. (2019). Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 2537–2546).
- Van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(11), 2579–2605.
- Menon, A.K., Jayasumana, S., Rawat, A.S., et al. (2020). Long-tail learning via logit adjustment. In *International Conference on Learning Representations*.
- Miyato, T., Shi, Maeda, Koyama, M., et al. (2018). Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8), 1979–1993.
- Ren, J., Yu, C., Ma, X., et al. (2020). Balanced meta-softmax for long-tailed visual recognition. *Advances in Neural Information Processing Systems*, 33, 4175–4186.
- Sohn, K., Berthelot, D., Carlini, N., et al. (2020). Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems*, 33, 596–608.
- Tarvainen, A., & Valpola, H. (2017). Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in Neural Information Processing Systems*, 30, 1195–1204.
- van Engelen, J. E., & Hoos, H. H. (2020). A survey on semi-supervised learning. *Machine Learning*, 109(2), 373–440.
- Wang, J., Geng, X. (2019). Classification with label distribution learning. In *International Joint Conference on Artificial Intelligence* (pp. 3712–3718).
- Wei, C., Sohn, K., Mellina, C., et al. (2021a). Crest: A class-rebalancing self-training framework for imbalanced semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 10857–10866).
- Wei, T., & Li, Y. F. (2019). Does tail label help for large-scale multi-label learning? *IEEE Transactions on Neural Networks and Learning Systems*, 31(7) 2315–2324.
- Wei, T., Shi, J., Li, Y., et al. (2022). Prototypical classifier for robust class-imbalanced learning. In *Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining* (pp. 44–57).
- Wei, T., Shi, J., Tu, W., et al. (2021b). Robust long-tailed learning under label noise. CoRR abs/2108.11569.
- Xiang, L., Ding, G., Han, J. (2020). Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In *European Conference on Computer Vision*, (pp. 247–263).
- Zagoruyko, S., Komodakis, N. (2016). Wide residual networks. In: *The British Machine Vision Conference*.
- Zhou, Z., Guo, L. Z., Cheng, Z., et al. (2021). Step: Out-of-distribution detection in the presence of limited in-distribution labeled data. *Advances in Neural Information Processing Systems*, 34, 29168–29180.
- Zhu, B., Niu, Y., Hua, X.S., et al. (2022). Cross-domain empirical risk minimization for unbiased long-tailed classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

Authors and Affiliations

Tong Wei¹ · Qian-Yu Liu² · Jiang-Xin Shi² · Wei-Wei Tu³ · Lan-Zhe Guo²

Qian-Yu Liu
liuqy@lamda.nju.edu.cn

Jiang-Xin Shi
shijx@lamda.nju.edu.cn

Wei-Wei Tu
tuweiwei@4paradigm.com

¹ School of Computer Science and Engineering, Southeast University, Nanjing 210096, China

² National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China

³ 4Paradigm Inc., Beijing 100000, China