Distribution-Aware Robust Learning from Long-Tailed Data with Noisy Labels

Jae Soon Baik¹, In Young Yoon¹, Kun Hoon Kim¹, and Jun Won Choi²

¹ Hanyang University, Korea {jsbaik, inyoungyoon, khkim}@spa.hanyang.ac.kr, ² Seoul National University, Korea junwchoi@snu.ac.kr

Abstract. Deep neural networks have demonstrated remarkable advancements in various fields using large, well-annotated datasets. However, real-world data often exhibit long-tailed distributions and label noise, significantly degrading generalization performance. Recent studies addressing these issues have focused on noisy sample selection methods that estimate the centroid of each class based on high-confidence samples within each target class. The performance of these methods is limited because they use only the training samples within each class for class centroid estimation, making the quality of centroids susceptible to longtailed distributions and noisy labels. In this study, we present a robust training framework called Distribution-aware Sample Selection and Contrastive Learning (DaSC). Specifically, DaSC introduces a Distributionaware Class Centroid Estimation (DaCC) to generate enhanced class centroids. DaCC performs weighted averaging of the features from all samples, with weights determined based on model predictions. Additionally, we propose a confidence-aware contrastive learning strategy to obtain balanced and robust representations. The training samples are categorized into high-confidence and low-confidence samples. Our method then applies Semi-supervised Balanced Contrastive Loss (SBCL) using high-confidence samples, leveraging reliable label information to mitigate class bias. For the low-confidence samples, our method computes Mixup-enhanced Instance Discrimination Loss (MIDL) to improve their representations in a self-supervised manner. Our experimental results on CIFAR and real-world noisy-label datasets demonstrate the superior performance of the proposed DaSC compared to previous approaches. Our code is available at https://github.com/JaesoonBaik1213/DaSC.

Keywords: Noisy label · Long-tailed learning · Contrastive learning

1 Introduction

In recent years, notable achievements have been made in various machine-learning applications owing to large-scale datasets with high-quality annotated labels

^{*}Corresponding author



 $\mathbf{2}$

J.S. Baik et al.

Fig. 1: (a) Class distribution on long-tailed CIFAR-10 with symmetric noise of 0.6. (b) and (c) show the illustration of previous class centroid estimation and our distributionaware class centroid estimation, respectively. In the previous methods, samples with low confidence, depicted as grey dots, are excluded from the class centroid estimation. In contrast, our approach leverages all samples for centroid estimation, using their prediction values as weights. The colors of the samples indicate the prediction values.

[11,16]. However, in practice, obtaining a large-scale dataset with accurately annotated labels is challenging and expensive. A realistic alternative is to collect data from the Internet through web crawling, but this is typically affected by long-tail distributions as well as mislabeled data [55]. Training a model on such datasets results in biases toward head classes and suffers from memorization of label noise [12], consequently leading to poor generalization performance.

Numerous methods have been explored to address the challenges of longtailed distributions and noisy labels. In long-tailed (LT) learning, re-balancing techniques are the most straightforward method to mitigate bias toward head classes. These techniques include oversampling, undersampling [5, 60], and adjusting instance-level or class-level weights in loss functions [4, 10, 41, 44]. Recently, inspired by the success of supervised contrastive learning (SCL) [25], several studies [9, 22, 30, 46, 61] have aimed to generate balanced representations by incorporating class distribution information into SCL. For learning with noisy labels (NL), various strategies have been proposed. These include robust loss functions [31, 41], sample selection methods for distinguishing clean from noisy samples [1, 12, 18, 45], and correcting noisy labels using model predictions [28,39,45]. Semi-supervised learning (SSL) has also been applied to enhance learning performance by utilizing noisy data in training [24, 28]. These methods often employ suitable distance metrics to identify noisy samples, which are then utilized as unlabeled data for SSL. Recent studies have leveraged SCL with reliable label information [18, 29, 36, 59] and unsupervised contrastive learning [6, 53]to produce representations robust to noisy labels.

Despite the success of individual approaches, simultaneously addressing longtailed distributions and label noise remains a significant challenge. The combined issue of noisy labels and long-tailed distributions (NL-LT) complicates the accurate identification of true data distributions, significantly degrading the effectiveness of these methods. Recent studies [27,33,51,52] have employed feature-based noisy sample selection methods to identify noisy labels. As illustrated in Fig. 1b, these methods rely on the distance of input features from the class centroids, which are obtained by averaging the features of high-confidence samples within each class.

However, this approach has several limitations. First, it uses only a portion of the training data within each class for estimating class centroids. As shown in Fig. 1a, samples mislabeled as other classes are excluded from the centroid estimation process. This exclusion results in insufficient data samples to reliably determine class centroids, which is particularly problematic for tail classes. Second, this method assigns equal weight to all samples when calculating class centroids. In the presence of labeling errors, this uniform weighting overlooks the potential for certain samples to be erroneous. Last, this approach does not actively seek to enhance representation quality. This limitation diminishes its effectiveness for tail classes, where obtaining high-quality representations is more critical than for head classes.

We aim to address these limitations by introducing a novel feature-based noisy sample selection method called Distribution-aware Class Centroid Estimation (DaCC). As shown in Fig. 1c, DaCC estimates class centroids by weighted averaging the features of samples from various classes. DaCC determines these weights based on model predictions, without relying on noisy labels, allowing for the adaptive inclusion of data samples from different classes in the centroid estimation process. Consequently, our method effectively utilizes informative samples with high-confidence scores when estimating class centroids.

Furthermore, we propose a Confidence-aware Contrastive Learning strategy to achieve balanced and robust representations in noisy, long-tailed classification. Specifically, we categorize samples into high-confidence and low-confidence groups based on the confidence scores derived from their pseudo-labels. For the high-confidence group, we employ Semi-supervised Balanced Contrastive Loss (SBCL), which mitigates class bias using confident label information. Simultaneously, for the low-confidence group, we introduce Mixup-enhanced Instance Discrimination Loss (MIDL), which enhances data representation in a self-supervised manner. MIDL utilizes mixup augmentation [57] to construct diverse negative keys, which leads to more robust representations for the lowconfidence group.

We integrate all the aforementioned components, presenting a novel training framework, referred to as *Distribution-aware Sample Selection and Contrastive Learning (DaSC)*. We conduct experiments on both modified CIFAR datasets and real-world datasets that exhibit NL-LT. The experiments demonstrate that DaSC achieves significant performance improvements over existing methods. The key contributions of this study are summarized as follows:

- We present a novel class centroid estimation method called DaCC. This method utilizes the entire set of training samples as compared to the existing methods. Class centroids are generated by averaging samples with weights obtained by temperature-scaled model predictions. This strategy leads to more reliable class centroid estimation and efficient use of training data.

- We also introduce a confidence-aware contrastive learning strategy to improve the quality of feature representations. This method applies two different contrastive losses, SBCL and MIDL, to low-confidence and high-confidence sample groups, respectively. SBCL uses label information to achieve a balanced representation of high-confidence samples, while MIDL employs a mixup augmentation to learn a robust representation for low-confidence samples.
- Our experimental results confirm that DaSC achieves state-of-the-art performance across various configurations.

2 Related Work

2.1 Long-Tailed Learning

Numerous approaches have been proposed to address the challenges of learning from long-tailed data. Re-balancing is a classical strategy for addressing class bias, including data re-sampling [5, 43] and loss re-weighting [4, 10, 34, 44]. Recently, long-tailed data issues have been studied from the perspectives of both representation learning and classifier learning [23, 34]. Ensemble-based methods have also been proposed to achieve robust performance under long-tailed data distributions [49, 60].

2.2 Learning with Noisy Labels

Various methods have been explored to achieve robust learning against noisy labels. These methods can be broadly classified into three approaches: 1) estimation of the noise transition matrix [7,38,54,56]; 2) robust loss functions [31,41,48]; and 3) noisy sample selection [1, 12, 18, 20, 28, 45]. Recently, some studies have explored applying SSL to address a noisy label issue [24, 28, 58]. These methods treated clean samples as labeled data and samples with noisy labels as unlabeled data after detecting correctly labeled samples.

2.3 Long-Tailed Learning with Noisy Labels

A series of studies have been conducted to resolve both noisy labels and longtailed distributions. Meta-learning methods [41,44] have been proposed to prioritize clean data by adaptively learning explicit instance weights from metadata. HAR [3] introduced an adaptive Lipschitz regularization that takes into account the uncertainty inherent in a particular data sample. To better differentiate mislabeled examples, RoLT [52] and PCL [51] calculated a centroid of each class by taking the average of the features with high confidence in the corresponding target class. The obtained class centroids were subsequently used for noisy sample selection in RoLT and as part of the classifier in PCL. Recently, SFA [27] adopted a Gaussian distribution to estimate the distribution of class centroids, addressing uncertainty in the representation space. TABASCO [33] employed two different

⁴ J.S. Baik et al.

metrics, one derived from model predictions and the other from the estimated class centroids. It then utilized a Gaussian Mixture Model (GMM) to determine the most appropriate metric for detecting correctly labeled samples with different noise types.

2.4 Contrastive Learning for Addressing Long-Tailed Distribution and Noisy Labels

Contrastive learning aims to enhance the representation by comparing sample pairs, pulling semantically similar samples closer together while pushing semantically dissimilar samples further apart. Contrastive learning has been successfully applied to various tasks [6, 13, 15, 47]. Recently, SCL [25] introduced fully supervised contrastive learning to encourage the samples of the same class to lie closer in feature space compared to those of different classes. Several LT methods [9,22,30,46,61] have enhanced SCL by utilizing supervision provided by long-tailed training data. PaCo [9] introduced parametric learnable class centers to learn a balanced representation. BCL [61] proposed a balanced contrastive loss that uses logit compensation to induce all samples to form a regular simplex. In NL learning, several methods [18,29,36,59] have employed contrastive learning [6, 15, 25] to better distinguish between noisy and clean data. MOIT [36] proposed interpolated contrastive learning that combines mixup with SCL to learn robust representations. Sel-CL [29] identified confidence pairs and modeled their pair-wise relations to obtain a robust representation.

3 Proposed Method

3.1 Overview of the Proposed Method

Problem Setup. Consider a K-class classification task. We train the model using the training dataset $\mathcal{D} = \{(x(i), \tilde{y}(i))\}_{i=1}^N$, where N denotes the total number of training samples, x(i) denotes the *i*-th training sample, and $\tilde{y}(i) \in \{1, ..., K\} = \mathcal{Y}$ is its associated label. As a result of labeling errors, a fraction of the training samples are incorrectly labeled, i.e., $\exists i \in [1, N], \tilde{y}(i) \neq y^*(i)$, where $y^*(i)$ indicates the true label for the sample x(i). We denote the training sample set of class k as $\mathcal{D}_k = \{(x(i), \tilde{y}(i)) | \tilde{y}(i) = k\}$. The training dataset \mathcal{D} follows a long-tailed distribution with an imbalance ratio of ρ , where $\rho = \min_k |\mathcal{D}_k| / \max_k |\mathcal{D}_k|$. Our goal is to achieve robust classification performance on unseen data, given a long-tailed noisy training dataset \mathcal{D} .

Overview. Figure 2 illustrates the overview of the proposed DaSC framework. Our architecture comprises DaCC, a sample selection module, a shared feature extractor f, a conventional classifier g^c , a balanced classifier g^b , a pseudo-label generator, and a Multi-Layer Perceptron (MLP) projector q. Initially, DaCC computes the class centroids. Then, correctly labeled samples are identified based on their distance from the class centroids. The next step is to train the model using noisy samples as unlabeled samples within the SSL framework [2]. For each



Fig. 2: Illustration of the proposed DaSC framework: Given a dataset, the DaSC detects correctly labeled samples using class centroids obtained from DaCC. Subsequently, the model is trained under a semi-supervised learning (SSL) framework, where confidence-aware contrastive learning is employed to enhance the representations. Cat and sg denotes concatenation and stop-gradient operations, respectively.

training sample $(x(i), \tilde{y}(i))$, two strongly augmented samples $x^1(i)$ and $x^2(i)$ and a weakly augmented sample $x^{w}(i)$ are generated [28]. Additionally, a mixup sample $(x^{mix}(i), y^{mix}(i))$ is created from $x^1(i)$ and $x^2(i)$ using mixup augmentation [57]. The conventional classifier g^c is applied to $x^{mix}(i)$ to produce a model prediction $p^{(mix,c)}(i) = \left[p_1^{(mix,c)}(i), \dots, p_K^{(mix,c)}(i)\right]^T$, where $p_k^{(mix,c)}(i)$ denotes the predicted probability of the k-th class. The balanced classifier g^b is then applied to $x^{mix}(i)$ to produce a model prediction $p^{(mix,b)}(i)$. This balanced classifier uses Balanced Softmax [40] to mitigate class bias. Following the DivideMix [28], we employ a pseudo-label generator to produce a pseudo-label $\hat{y}(i)$, based on the provided labels and model predictions for labeled samples or solely on the model predictions for unlabeled samples. The projector q is applied to both $x^{mix}(i)$ and the concatenation of $x^1(i)$ and $x^2(i)$, producing lower-dimensional representations $z^{mix}(i)$ and z(i), respectively. These representations are used in confidence-aware contrastive learning to enhance representation quality. Specifically, we classify an input sample x(i) into a high-confidence sample or a lowconfidence sample by comparing the maximum confidence score in $\hat{y}(i)$ with a threshold τ_c . Subsequently, we employ SBCL for the high-confidence sample group and MIDL for the low-confidence sample group.

3.2 Distribution-Aware Class Centroid Estimation

Conventional sample selection methods proposed in [27, 33, 52] used only target class samples to estimate their class centroids. In contrast, the proposed DaCC method estimates class centroids using samples from all classes, weighting each sample based on its confidence score. However, noisy labels can adversely affect the class centroids. To mitigate this, we apply temperature scaling (TS) to the model predictions, reducing the weights for unreliable samples and increasing them for more reliable ones. DaCC takes the *i*-th image x(i) as input and produces the low-dimensional features z'(i) using the backbone network followed by the MLP projector. Then, it applies the conventional classifier and the balanced classifier to produce the prediction outputs $p^c(i) = [p_1^c(i), ..., p_K^c(i)]$ and $p^b(i) = [p_1^b(i), ..., p_K^b(i)]$, respectively, where k is the class index. The class centroid c_k for the k-th class is calculated as

$$c_k = Norm\left(\sum_{x(i)\in\mathcal{D}^I} \hat{p}_k^c(i) z'(i)\right),\tag{1}$$

where

$$\hat{p}_{k}^{c}(i) = \frac{\exp\left(p_{k}^{c}(i)/\tau_{T}\right)}{\sum_{k=1}^{K}\exp\left(p_{k}^{c}(i)/\tau_{T}\right)},\tag{2}$$

$$\mathcal{D}^{I} = \{x(i) | i \in [1, N], \max_{k} p_{k}^{c}(i) > \tau, \max_{k} p_{k}^{b}(i) > \tau\},$$
(3)

where τ_T denotes the temperature parameter, τ denotes the confidence threshold, and Norm denotes l_2 normalization. As training goes on, we gradually increase the confidence threshold as $\tau = \phi^t \hat{\tau}$, where t is the epoch index. Following the setting from the previous works [27,51,52], ϕ is set to 1.005 and $\hat{\tau}$ is set to 1/K. Note that the samples used to obtain the centroid for the k-th class belong to any classes.

After calculating the class centroids, we identify correctly labeled samples. The assignment probability of a sample x(i) belonging to class centroid c_k is given by

$$\Gamma(y(i) = k | z'(i), c) = \frac{\exp(z'(i) \cdot c_k)}{\sum_k \exp(z'(i) \cdot c_k)},\tag{4}$$

where $x \cdot y$ denotes the dot-product operation. Subsequently, we fit a twocomponent GMM to maximize the likelihood of the assignment probability, *i.e.*, $\Gamma(y(i) = k|x(i), c) \sim \sum_{j=1}^{2} w_{k,j} \mathcal{N}(\mu_{k,j}, \sigma_{k,j})$, where $\mu_{k,j}$ and $\sigma_{k,j}$ are the mean and variance of the *j*-th Gaussian component. We expect that correctly labeled samples will cluster close to the class centroids, while noisy samples will be more dispersed. Therefore, samples with a posterior probability exceeding 0.5 are classified as clean, and the remaining samples are identified as noisy.

3.3 Confidence-Aware Contrastive Learning

Confidence-aware contrastive learning applies different strategies to high-confidence and low-confidence samples. We classify a sample x(i) by comparing the maximum confidence score in the pseudo-label $\hat{y}(i)$ with a threshold τ_c . If the score exceeds the threshold, the sample is classified as a high-confidence one; otherwise, it is classified as a low-confidence one. 8 J.S. Baik et al.

Semi-Supervised Balanced Contrastive Loss. Given a set of high-confidence samples B^{high} within a mini-batch B, the SBCL is expressed as

$$\mathcal{L}_{SBCL} = \frac{1}{|B^{high}|} \sum_{i \in B^{high}} \mathcal{L}_{SBCL}(i), \tag{5}$$

where

$$\mathcal{L}_{SBCL}(i) = -\frac{1}{|B_{y_i}^{high}|} \sum_{p \in B_{y_i}^{high} \setminus \{i\}} \log \frac{\exp(z(i) \cdot z(p)/\tau_s)}{\sum_{j=1}^{K} \frac{1}{|B_j^{high}|} \sum_{k \in B_j^{high}} \exp(z(i) \cdot z(k)/\tau_s)},$$
(6)

where $A \setminus x$ denotes the set subtraction operation, y_i is the index of the largest entry of the pseudo-label $\hat{y}(i)$, B_j^{high} represents the set of sample indices for high-confidence samples whose predictions correspond to the class j within a mini-batch B, and τ_s is the temperature parameter. This method effectively leverages reliable label information from high-confidence samples to penalize head classes, aiding in the learning of balanced representations.

Mixup Enhanced Instance Discrimination Loss. Building on a simple instance discrimination task [53], we propose mixup-enhanced instance discrimination loss for low-confidence samples. For each sample x(i) within the mini-batch B^{low} of low-confidence samples, we generate two augmented images, $x^1(i)$ and $x^2(i)$, serving as the query and positive key, respectively. We obtain the feature representations $z^1(i)$ and $z^2(i)$ from $x^1(i)$ and $x^2(i)$ by applying the backbone network, followed by the MLP projector. Additionally, we devise a variant that generates a negative key from a mixup image x^{mix} [57]. Note that mixup images are used to enhance the diversity of negative samples. The mixup image is obtained by taking a convex combination between $x^1(i)$ and $x^2(i)$ as

$$x^{mix}(i) = \lambda x^{1}(i) + (1 - \lambda)x^{2}(i),$$
(7)

where λ denotes the mixing ratio obtained from Beta distribution, e.g., $\lambda \sim Beta(\alpha, \alpha)$. We extract the key representations $z^{mix}(i)$ from $x^{mix}(i)$ and feed this negative key into the queue (i.e., the memory bank) with the fixed size |M|. Mixup samples serve as hard negative samples [21, 35], effectively improving the performance of instance discrimination loss. The proposed mixup-enhanced instance discrimination loss (MIDL) can be expressed by

$$\mathcal{L}_{MIDL} = \frac{1}{|B^{low}|} \sum_{i \in B^{low}} \mathcal{L}_{MIDL}(i), \tag{8}$$

where

$$\mathcal{L}_{MIDL}(i) = -\log \frac{\exp(z^1(i) \cdot z^2_+(i)/\tau_m)}{\exp(z^1(i) \cdot z^2_+(i)/\tau_m) + \sum_{z^{mix}_- \in M} \exp(z^1(i) \cdot z^{mix}_-/\tau_m)}, \quad (9)$$

where M denotes the memory bank, τ_m denotes the temperature parameter, $z_+^2(i)$ denotes the representation of a positive key from a different view of the

query representation $z^1(i)$ and z_{-}^{mix} denotes the representation of a negative key sourced from the memory bank. This loss significantly enhances discrimination between low-confidence labeled samples in a self-supervised manner.

3.4 Training with Semi-supervised Learning

Balanced Classifier. Various methods [27, 49, 60] have implemented balanced classifiers to mitigate class bias towards head classes. Our method also employs a balanced classifier trained with Balanced Softmax [40]. We input the mixup image $x^{mix}(i)$ into the feature extractor f, followed by the balanced classifier, to produce the logit $\eta^{(mix,b)}(i) = \left[\eta_1^{(mix,b)}(i), ..., \eta_K^{(mix,b)}(i)\right]$, where k is the class index. The loss function for Balanced Softmax is expressed as

$$\mathcal{L}_{BS}(i) = -\log \frac{n_k \exp\left(\eta_k^{(mix,b)}(i)\right)}{\sum_{k=1}^K n_k \exp\left(\eta_k^{(mix,b)}(i)\right)},\tag{10}$$

where n_k denotes the number of samples of class k among the data samples selected by the proposed DaCC. During the warmup phase, we determine n_k based on the distribution of the training dataset \mathcal{D} .

Total Loss. To effectively utilize noisy samples, we consider them as unlabeled data and train the model under a semi-supervised learning (SSL) framework [2]. The total training loss is

$$\mathcal{L} = \mathcal{L}_{MixMatch} + \mathcal{L}_{BMixMatch} + \lambda_{SBCL} \mathcal{L}_{SBCL} + \lambda_{MIDL} \mathcal{L}_{MIDL}, \qquad (11)$$

where $\mathcal{L}_{MixMatch}$ denotes the MixMatch loss for the conventional classifier, $\mathcal{L}_{BMixMatch}$ denotes the MixMatch loss with \mathcal{L}_{BS} for the balanced classifier, and λ_{SBCL} and λ_{MIDL} are the coefficients for SBCL and MIDL, respectively. Algorithm 1 presents the summary of our DaSC training framework. During the test time, we use the average of outputs from the classifiers g^c and g^b as the final model output. Furthermore, following previous studies [27, 28, 52], we employ a co-training framework against confirmation bias.

4 Experiments

4.1 Setup

Datasets. The proposed method was evaluated on long-tailed version of the CIFAR-10 and CIFAR-100 datasets [26], with synthetic noise introduced to the labels. Additionally, we used real-world noise datasets including CIFAR-10N, CIFAR-100N [50], and Red mini-ImageNet [19] for evaluation.

For the synthetic CIFAR-10 and CIFAR-100 datasets, we followed the experimental setup in [33]. We initially created long-tailed version of the CI-FAR datasets by setting the number of samples for the k-th class as $N_k = \max_k |\mathcal{D}_k| \cdot \rho^{\frac{k-1}{K-1}}$, where ρ is the imbalance ratio and K indicates the number 10 J.S. Baik et al.

Algorithm 1 DaSC framework

Input: Training dataset $\mathcal{D} = \{(x_i, \tilde{y}_i)\}_{i=1}^N$, model parameters θ , warmup epochs T_0 , warmup epochs for semi-supervised balanced contrastive loss T_{SBCL}

1: for $t = 1, ..., T_0$ do $\mathcal{L} = \mathcal{L}_{CE} + \mathcal{L}_{BS} + \mathbb{I}(t > T_{SBCL})\lambda_{SBCL}\mathcal{L}_{SBCL}$ 2:3: Update θ 4: end for 5: for $t = T_0 + 1, \dots, T$ do Compute the confidence threshold $\tau = \Phi^t \hat{\tau}$ 6: Compute the class centroids c by Eqs. (1)-(3) 7: 8: for i = 1, ..., N do Calculate the assignment probability Γ by Eq. (4) 9: 10: end for $\mathcal{D}^{clean}, \mathcal{D}^{noisy} = \mathrm{GMM}(\Gamma)$ 11: $\mathcal{L} = \mathcal{L}_{MixMatch} + \mathcal{L}_{BMixMatch} + \lambda_{SBCL} \mathcal{L}_{SBCL} + \lambda_{MIDL} \mathcal{L}_{MIDL}$ 12:13:Update θ 14: end for

of classes. Subsequently, we introduced two types of synthetic noise into the datasets: symmetric (Sym.) and asymmetric (Asym.) label noise. Symmetric noise involves randomly assigning labels to any other class based on a specified noise ratio. Asymmetric noise, on the other hand, changes the labels to semantically related classes (e.g., truck \rightarrow automobile), taking class information into account. In our experiments, we used noise ratios of 0.4 and 0.6 for symmetric noise and 0.2 and 0.4 for asymmetric noise.

For CIFAR-10N (10N) and CIFAR-100N (100N) [50], we modified the CIFAR datasets to the aforementioned long-tailed version and then introduced noisy labels created by human annotations. Red mini-ImageNet (Red) [19] is a smaller version of ImageNet [42] that includes real-world noisy web labels with specific ratios. In this study, we used web labels with noise ratios of 0.2 and 0.4 for Red mini-ImageNet.

Baselines. The following baseline methods were compared with our method. Methods addressing long-tailed distribution include LA [34], LDAM [4], IB Loss [37], and BBN [60]. Methods for dealing with noisy labels include DivideMix [28], UNICON [24], and TCL [18]. The latest methods addressing both noisy labels and long-tailed distribution include Meta-weight Net [44], RoLT [52], HAR [3], ULC [17], SFA [27], and TABASCO [33].

4.2 Implementation Details

We used PreAct ResNet18 and ResNet18 [14] as backbone networks for the experiments on the CIFAR and Red mini-ImageNet datasets, respectively. The model was trained for 100 epochs using a stochastic gradient descent optimizer with an initial learning rate of 0.02, a momentum of 0.9, and a batch size of 64. Weight decay of 1×10^{-4} and 5×10^{-5} were applied for the CIFAR and Red mini-ImageNet datasets, respectively. Contrastive loss was introduced after

the first 10 epochs, and correctly labeled samples were detected after the first 30 epochs. Data augmentation techniques included random horizontal flips, random cropping, and RandAugment [8]. Table 1 provides the configuration parameters of our DaSC.

Table 1: The configuration parameters used for DaSC.

Hyperparameter	Value
The coefficient λ_{SBCL} for SBCL	0.5
The temperature τ_s for SBCL	0.1
The coefficient λ_{MIDL} for MIDL	0.3
The temperature τ_m for MIDL	0.5
The temperature parameter τ_T	0.1
The confidence threshold τ_c for sample grouping	0.9

4.3 Main Results

Results on CIFAR Datasets. Tables 2a and 2b display the performance of the proposed DaSC on long-tailed noisy versions of the CIFAR datasets. We examine both symmetric and asymmetric noise in Tables 2a and 2b, respectively. The proposed DaSC consistently outperforms other baseline methods in all experimental setups. For CIFAR-10 with a symmetric noise ratio of 0.4, DaSC achieves performance gains of 2.23% and 3.51% over the current best methods, SFA and TABASCO, respectively. These performance gains tend to increase as the noise ratio and the number of classes rise. On CIFAR-100 with a symmetric noise ratio of 0.6, DaSC achieves substantial performance gains of 6.69% and 8.42% over SFA and TABASCO, respectively.

Results on Real-World Datasets. Table 3 shows the performance of DaSC on real-world datasets with long-tailed distributions. These datasets are more challenging than those with synthetic labeling noise because samples are more prone to being mislabeled into semantically similar classes, and noisy labels are unevenly distributed. Notably, our DaSC method achieves significant performance gains on real-world datasets created by human annotations, aligning with the experimental results observed in the previous experiments. Our DaSC achieves a performance gain of 2.04% over the existing best method, TABASCO [33] with a noise ratio of 0.2 on the Red mini-ImageNet dataset. Additionally, DaSC outperforms TABASCO by 5.40% on CIFAR-10N and by 4.27% on CIFAR-100N.

4.4 Performance Analysis

Ablation Study. We conducted an ablation study to analyze the impact of the main components in the proposed DaSC. Table 4 presents the experimental results on CIFAR datasets with various noise types and ratios. DaCC offers significant performance gains over the baseline, confirming the effectiveness of the

12 J.S. Baik et al.

Dataset CIFAR-10 CIFAR-100 Dataset CIFAR-10 CIFAR-100 Imbalance Ratio 0.1 Imbalance Ratio 0.1 Noise Ratio 0.40.6 0.4 0.6 Noise Ratio 0.2 04 0.20.4 Baseline CE 71.67 61.16 34.53 23.63 Baseline CE 79.90 62.88 44.45 32.05 LA [34] 70.56 54.92 29.07 23.21 71.49 59.88 39.34 28.49 LA [34] 74.58 62.29 40.06 33.26 LDAM [4] 70.53 61.97 31.30 23.13 LDAM [4] LT LT IB [37] 73.24 62.62 32.40 25.84 IB [37] 73.49 58.36 45.0235.25BBN [60] 71.16 56.38 30.48 23.00 BBN [60] 75.57 60.91 39.16 31.03 DivideMix [28] DivideMix [28] 82.67 80.17 54.71 44.98 80.92 69.35 58.09 41.99 UNICON [24] NL 84.25 82.29 52.34 45.87 NL UNICON [24] 72.81 69.04 55.99 44.70 TCL [18] 79.55 72.13 51.49 41.19 TCL [18] 84.55 72.97 58.61 44.75 MW-Net [44] 70.90 59.8532.03 21.71MW-Net [44] 79.34 65.49 42.5230.42RoLT [52] RoLT [52] 81.62 76.58 42.95 32.59 73.30 58.29 48.19 39.32 HAR [3] 77.44 63.75 38.17 26.09 HAR [3] 82.85 69.19 48.50 33.20 NL-LT NL-LT ULC [17] 84.46 83.25 54.91 44.66 ULC [17] 74.07 73.19 54.45 43.20 SFA [27] SFA [27] 86.81 82.89 56.70 47.71 87.70 78.13 57.16 46.09 TABASCO [33] 85.53 84.83 56.52 45.98 TABASCO [33] 82.10 80.57 59.39 50.51 DaSC 89.89 88.85 63.22 53.66 DaSC 89.04 87.12 61.85 54.40 (a) Symmetric noise (b) Asymmetric noise

Table 2: Performance of the proposed DaSC on the long-tailed version of CIFAR datasets with (a) a symmetric noise and (b) an asymmetric noise.

proposed centroid estimation method. Specifically, it achieves a 1.85% performance gain on CIFAR-10 with a symmetric noise ratio of 0.4. Incorporating TS into DaCC leads to further performance improvement of 0.4% gain under the same setup. Adding SBCL further enhances performance by 0.76%. The combination of SBCL and MIDL provides a 1.51% performance gain, resulting in a total performance gain of 3.76% over the baseline. These performance gains are consistent across various datasets and noise ratios, demonstrating the robustness and versatility of the proposed DaSC.

Hyperparameters. The proposed DaSC performs consistently well across various hyperparameter configurations. The experimental results with different hyperparameter configurations are provided in the supplementary material.

Generation of Negative Keys for MIDL. Table 5 illustrates the effect of using mixup samples to generate negative keys for MIDL. In the absence of mixup samples, low-confidence samples are used as negative keys. Notably, using mixup samples results in better performance compared to not using them, demonstrating that mixup samples provide more effective negative keys for contrastive learning.

Analysis of Confidence-Aware Contrastive Learning. First, Figure 3 shows the accuracy of pseudo-labels relative to true labels for different sample groups. The terms 'All', 'High', and 'Low' refer to using all samples, a high-confidence sample group, and a low-confidence sample group, respectively. As expected, the high-confidence samples align more closely with the true labels, while low-confidence samples tend to be less consistent with the true labels.

Next, Table 6 evaluates the effect of using different sample groups for SBCL and MIDL. Our results confirm that using the high-confidence sample group for SBCL and the low-confidence sample group for MIDL yields the best perfor-

 Table 3: Performance of the proposed Table 4: Ablation study on the long-tailed

 DaSC on the long-tailed version of CIFAR version of CIFAR.

 datasets and Red mini-ImageNet dataset.

Dataset Red 10N 10NN Imbalance Ratio ≈ 0.1 0.1 0.1 Noise Ratio 0.2 0.4 ≈ 0.4 Baseline CE 40.42 31.46 60.44 38.10 Haseline CE 40.42 31.46 60.44 38.10 LA [34] 26.62 25.88 65.74 36.50 LDAM [4] 26.64 23.46 62.50 38.48 IB [37] 23.80 20.08 65.91 42.48 BBN [60] 42.18 42.64 68.42 37.22 NL UNICON [24] 40.18 41.64 69.54 51.74 ROLT [52] 22.56 24.22 75.46 51.74 MW-Net [44] 42.66 40.26 69.73 44.50 NL-LT [52] 22.56 24.22 75.44 65.17 SFA [27] 48.12 47.06 75.71 51.72 SFA [27] 42.50 43.62 45.51						
$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	Dataset		Red		10N	100N
$\begin{array}{c c c c c c c c c c c c c c c c c c c $	Imbalance Ratio		≈ 0.1		0.1	
$ \begin{array}{c c c c c c c c c c c c c c c c c c c $	Noise Ratio		0.2 0.4		≈ 0.4	
LA [34] 26.82 25.88 65.74 36.50 LDAM [4] 26.64 23.46 62.50 38.48 IB [37] 23.80 22.08 65.91 42.48 BBN [60] 42.18 42.64 68.42 37.22 DivideMix [28] 48.76 48.96 67.85 44.25 NL DivideMix [28] 40.18 41.64 69.54 51.93 TCL [18] 37.12 37.76 73.46 51.74 MW-Net [44] 42.66 40.26 69.73 44.20 RoLT [52] 22.56 24.22 75.24 46.61 HAR [3] 46.61 38.71 74.97 44.54 NL-LT ULC [17] 48.12 47.06 75.71 51.72 SFA [27] 42.50 43.62 84.57 53.68 TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 86.01 53.81	Baseline	CE	40.42	31.46	60.44	38.10
LDAM [4] 26.64 23.46 62.50 38.48 IB [37] 23.80 22.08 65.91 42.48 BBN [60] 42.18 42.64 68.42 37.22 MU DivideMix [28] 48.76 48.96 67.85 44.25 NL DivideMix [28] 48.76 40.76 69.54 51.93 TCL [18] 37.12 37.76 73.46 51.74 MW-Net [44] 42.66 40.26 60.73 44.20 RoLT [52] 22.56 24.22 75.24 46.61 HAR [3] 46.61 38.71 74.97 44.54 NL-LT ULC [17] 48.12 47.06 75.71 51.72 SFA [27] 42.50 43.62 84.57 53.68 TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 50.15 51.70	LT	LA [34]	26.82	25.88	65.74	36.50
III IB [37] 23.80 22.08 65.91 42.48 BBN [60] 42.18 42.64 68.42 37.22 DivideMix [28] 48.76 48.96 67.85 44.25 NL DivideMix [28] 48.76 48.96 67.85 44.25 NL UNICON [24] 40.18 41.64 69.54 51.93 TCL [18] 37.12 37.76 73.46 51.74 MW-Net [44] 42.66 40.26 60.73 44.20 RoLT [52] 22.56 24.22 75.24 46.61 HAR [3] 46.61 38.71 74.97 44.54 NL-LT ULC [17] 48.12 47.06 75.71 51.72 SFA [27] 42.50 43.62 84.57 53.68 TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 86.01 53.81		LDAM [4]	26.64	23.46	62.50	38.48
BBN [60] 42.18 42.64 68.42 37.22 DivideMix [28] 48.76 48.96 67.85 44.25 NL UNICON [24] 40.18 41.64 69.54 51.93 TCL [18] 37.12 37.76 73.46 51.74 MW-Net [44] 42.66 40.26 60.73 44.20 RoLT [52] 22.56 24.22 75.24 46.61 HAR [3] 46.61 38.71 74.94 44.54 NL-LT ULC [17] 48.12 47.06 75.71 51.72 SFA [27] 42.50 43.62 84.57 53.68 TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 86.01 53.81		IB [37]	23.80	22.08	65.91	42.48
DivideMix [28] 48.76 48.96 67.85 44.25 NL UNICON [24] 40.18 41.64 69.54 51.93 TCL [18] 37.12 37.76 73.46 51.74 MW-Net [44] 42.66 40.26 69.53 44.20 RoLT [52] 22.56 24.22 75.24 46.61 HAR [3] 46.61 38.71 74.97 44.54 NL-LT ULC [17] 48.12 47.06 75.71 51.72 SFA [27] 42.50 43.62 84.57 53.68 TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 86.01 58.10		BBN [60]	42.18	42.64	68.42	37.22
NL UNICON [24] 40.18 41.64 69.54 51.93 TCL [18] 37.12 37.76 73.46 51.74 MW-Net [44] 42.66 40.26 69.73 44.20 RoLT [52] 22.56 24.22 75.24 46.61 HAR [3] 46.61 38.71 74.97 44.54 NL-LT ULC [17] 48.12 47.06 75.71 51.72 SFA [27] 42.50 43.62 84.57 53.68 TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 86.01 58.10	NL	DivideMix [28]	48.76	48.96	67.85	44.25
TCL [18] 37.12 37.76 73.46 51.74 MW-Net [44] 42.66 40.26 69.73 44.20 RoLT [52] 22.56 24.22 75.24 46.61 HAR [3] 46.61 38.71 74.97 44.54 NL-LT ULC [17] 48.12 47.06 75.71 51.72 SFA [27] 42.50 43.62 84.57 53.68 TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 86.01 58.10		UNICON [24]	40.18	41.64	69.54	51.93
MW-Net [44] 42.66 40.26 69.73 44.20 RoLT [52] 22.56 24.22 75.24 46.61 HAR [3] 46.61 38.71 74.97 44.54 NL-LT ULC [17] 48.12 47.06 75.71 51.72 SFA [27] 42.50 43.62 84.57 53.68 TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 86.01 58.10		TCL [18]	37.12	37.76	73.46	51.74
RoLT [52] 22.56 24.22 75.24 46.61 HAR [3] 46.61 38.71 74.97 44.54 NL-LT ULC [17] 48.12 47.06 75.71 51.72 SFA [27] 42.50 43.62 84.57 53.68 TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 86.01 58.10	NL-LT	MW-Net [44]	42.66	40.26	69.73	44.20
HAR [3] 46.61 38.71 74.97 44.54 NL-LT ULC [17] 48.12 47.06 75.71 51.72 SFA [27] 42.50 43.62 84.57 53.68 TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 86.01 58.10		RoLT [52]	22.56	24.22	75.24	46.61
NL-LT ULC [17] 48.12 47.06 75.71 51.72 SFA [27] 42.50 43.62 84.57 53.68 TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 86.01 58.10		HAR [3]	46.61	38.71	74.97	44.54
SFA [27] 42.50 43.62 84.57 53.68 TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 86.01 53.10		ULC [17]	48.12	47.06	75.71	51.72
TABASCO [33] 50.20 49.68 80.61 53.83 DaSC 52.24 50.58 86.01 58.10		SFA [27]	42.50	43.62	84.57	53.68
DaSC 52.24 50.58 86.01 58.10		TABASCO [33]	50.20	49.68	80.61	53.83
		DaSC	52.24	50.58	86.01	58.10

Main Component CIFAR-10 CIFAR-100 Sample Selection Contrastive Learning Sym. Asym. Sym. Asym. DaCC TS SBCL MIDL 0.40.2 0.40.285.28 84.7654.7658.35 87.13 87.98 57.81 59.93 87.53 88.48 58.23 61.28 √ 88.29 89.14 60.99 61.94 √ 89.04 89.89 61.85 63.22

Table 5: Performance comparison withdifferent types of negative keys.

	CIFAR-10		CIFAR-100	
Negative Key	Sym.	Asym.	Sym.	Asym.
	0.4	0.2	0.4	0.2
Without Mixup	88.25	89.35	61.22	62.50
With Mixup	89.04	89.89	61.85	63.22



Fig. 3: Accuracy of pseudo-labels for different sample groups.

Table 6: Performance of DaSC when dif-
ferent sample groups are used for SBCL
and MIDL.

Samp	Sample Set CIFAR-10		nple Set CIFAR-10 CIFAR-1		R-100
SBCL	MIDI	Sym.	Asym.	Sym.	Asym.
	MIDL	0.4	0.2	0.4	0.2
All	All	88.19	89.20	60.86	62.20
High	All	88.45	89.52	61.30	62.69
All	Low	88.76	89.46	61.25	62.77
High	Low	89.04	89.89	61.85	63.22

mance among the configurations tested in the experiments. This highlights the effectiveness of our contrastive learning strategy, which efficiently utilizes each sample group for representation learning.

Performance Evaluation on Different Class Sets. In Fig. 4, we compare the performance of DaSC with that of the baseline across different class sets: 'Many', 'Medium', and 'Few'. This categorization is based on the framework presented in previous work [32]. We evaluate both methods on the long-tailed version of CIFAR-10 with symmetric noise of 0.6 and asymmetric noise of 0.4. The proposed DaSC method significantly improves performance across all class sets. We observe that DaSC enhances performance not only for the Few and Medium class sets but also for the Many class set. These improvements suggest that DaSC effectively enhances data representation using confidence-aware contrastive learning, leading to better recognition of samples across the Many, Medium, and Few class sets.



Fig. 4: Performance of the DaSC com- Fig. 5: AUC of the proposed DaSC in pared with the baseline method for differ- comparison of SFA for noisy label detecent class sets including 'Many', 'Medium', tion. We present the performance for 'All' and 'Few'. and 'Few' class sets.

Accuracy of Detecting Noisy Labels. Figure 5 illustrates the area under the curve (AUC) for DaSC in detecting noisy labels using centroids for each epoch. We compare our method with the SFA, with both methods trained on the long-tailed version of CIFAR-10, incorporating symmetric label noise of 0.6 and asymmetric label noise of 0.4. Additionally, we present the AUC results for both methods specifically for the 'Few' class set. Experimental results demonstrate that our method consistently outperforms SFA throughout the entire training process. Notably, while the accuracy of SFA fluctuates across epochs, our method shows a smooth improvement. In the Few class set, SFA struggles to converge until the end of the training, whereas our method maintains consistently strong detection performance. These results demonstrate the effectiveness of our method in detecting correctly labeled samples, even within the Few class set.

5 Conclusion

In this study, we proposed a novel robust learning framework, DaSC, to address the challenges of long-tailed class distribution and noisy label simultaneously. Recent approaches have employed feature-based noisy sample selection methods that rely on class centroids computed using high-confidence samples within each target class. In contrast, we introduced DaCC, which estimates class centroids using samples from all other classes. To obtain high-quality class centroids, we adaptively weight the samples based on a confidence score refined by temperature scaling. Additionally, we introduced a confidence-aware contrastive learning framework to enhance representation quality. We first categorized each training sample into a high-confidence sample and a low-confidence sample. Then, we computed SBCL for high-confidence samples and MIDL for low-confidence samples. SBCL incorporates class-balanced contrastive learning loss using class information from high-confidence samples, while MIDL employs mixup samples derived from low-confidence samples as negative keys for contrastive learning. Our extensive experiments demonstrated that the proposed DaSC achieves stateof-the-art performance in widely used benchmarks.

Acknowledge

This work was partly supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No.2020R1A2C2012146), Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) [NO.RS-2021-II211343, Artificial Intelligence Graduate School Program (Seoul National University)], and the National Research Foundation (NRF) funded by the Korean government (MSIT) (No. RS-2024-00421129).

References

- Arazo, E., Ortego, D., Albert, P., O'Connor, N., McGuinness, K.: Unsupervised label noise modeling and loss correction. In: ICML. pp. 312–321. PMLR (2019)
- Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. In: NeurIPS. vol. 32, pp. 5050–5060 (2019)
- 3. Cao, K., Chen, Y., Lu, J., Aréchiga, N., Gaidon, A., Ma, T.: Heteroskedastic and imbalanced deep learning with adaptive regularization. In: ICLR (2021)
- Cao, K., Wei, C., Gaidon, A., Aréchiga, N., Ma, T.: Learning imbalanced datasets with label-distribution-aware margin loss. In: NeurIPS. pp. 1565–1576 (2019)
- Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P.: Smote: Synthetic minority over-sampling technique. JAIR 16, 321–357 (2002)
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML. pp. 1597–1607. PMLR (2020)
- Cheng, D., Ning, Y., Wang, N., Gao, X., Yang, H., Du, Y., Han, B., Liu, T.: Classdependent label-noise learning with cycle-consistency regularization. In: NeurIPS. vol. 35, pp. 11104–11116 (2022)
- 8. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.: Randaugment: Practical automated data augmentation with a reduced search space. In: NeurIPS (2020)
- Cui, J., Zhong, Z., Liu, S., Yu, B., Jia, J.: Parametric contrastive learning. In: ICCV. pp. 715–724 (2021)
- Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S.: Class-balanced loss based on effective number of samples. In: CVPR. pp. 9268–9277 (2019)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021)
- Han, B., Yao, Q., Yu, X., Niu, G., Xu, M., Hu, W., Tsang, I.W., Sugiyama, M.: Co-teaching: Robust training of deep neural networks with extremely noisy labels. In: NeurIPS. pp. 8536–8546 (2018)
- Han, T., Xie, W., Zisserman, A.: Self-supervised co-training for video representation learning. In: NeurIPS. vol. 33, pp. 5679–5690 (2020)
- He, H., Garcia, E.A.: Learning from imbalanced data. IEEE Trans. KDE 21(9), 1263–1284 (2009)
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR. pp. 9729–9738 (2020)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)

- 16 J.S. Baik et al.
- Huang, Y., Bai, B., Zhao, S., Bai, K., Wang, F.: Uncertainty-aware learning against label noise on imbalanced datasets. In: AAAI. pp. 6960–6969 (2022)
- Huang, Z., Zhang, J., Shan, H.: Twin contrastive learning with noisy labels. In: CVPR. pp. 11661–11670 (2023)
- Jiang, L., Huang, D., Liu, M., Yang, W.: Beyond synthetic noise: Deep learning on controlled noisy labels. In: ICML. pp. 4804–4815. PMLR (2020)
- Jiang, L., Zhou, Z., Leung, T., Li, L., Fei-Fei, L.: Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In: ICML. pp. 2309– 2318 (2018)
- Kalantidis, Y., Sariyildiz, M.B., Pion, N., Weinzaepfel, P., Larlus, D.: Hard negative mixing for contrastive learning. In: NeurIPS. vol. 33, pp. 21798–21809 (2020)
- 22. Kang, B., Li, Y., Xie, S., Yuan, Z., Feng, J.: Exploring balanced feature spaces for representation learning. In: ICLR (2020)
- Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., Kalantidis, Y.: Decoupling representation and classifier for long-tailed recognition. In: ICLR (2020)
- Karim, N., Rizve, M.N., Rahnavard, N., Mian, A., Shah, M.: Unicon: Combating label noise through uniform selection and contrastive learning. In: CVPR. pp. 9676–9686 (2022)
- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. In: NeurIPS. vol. 33, pp. 18661–18673 (2020)
- 26. Krizhevsky, A.: Learning multiple layers of features from tiny images. Technical report, Department of Computer Science, University of Toronto (2009)
- Li, H.T., Wei, T., Yang, H., Hu, K., Peng, C., Sun, L.B., Cai, X.L., Zhang, M.L.: Stochastic feature averaging for learning with long-tailed noisy labels. In: IJCAI (2023)
- 28. Li, J., Socher, R., Hoi, S.C.H.: Dividemix: Learning with noisy labels as semisupervised learning. In: ICLR (2020)
- Li, S., Xia, X., Ge, S., Liu, T.: Selective-supervised contrastive learning with noisy labels. In: CVPR. pp. 316–325 (2022)
- Li, T., Cao, P., Yuan, Y., Fan, L., Yang, Y., Feris, R.S., Indyk, P., Katabi, D.: Targeted supervised contrastive learning for long-tailed recognition. In: CVPR. pp. 6918–6928 (2022)
- Liu, S., Niles-Weed, J., Razavian, N., Fernandez-Granda, C.: Early-learning regularization prevents memorization of noisy labels. In: NeurIPS. vol. 33, pp. 20331– 20342 (2020)
- Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., Yu, S.X.: Large-scale long-tailed recognition in an open world. In: CVPR. pp. 2537–2546 (2019)
- Lu, Y., Zhang, Y., Han, B., Cheung, Y.m., Wang, H.: Label-noise learning with intrinsically long-tailed data. In: ICCV. pp. 1369–1378 (2023)
- Menon, A.K., Jayasumana, S., Rawat, A.S., Jain, H., Veit, A., Kumar, S.: Long-tail learning via logit adjustment. In: ICLR (2021)
- Oh, J., Yun, C.: Provable benefit of mixup for finding optimal decision boundaries. In: ICML. pp. 26403–26450. PMLR (2023)
- Ortego, D., Arazo, E., Albert, P., O'Connor, N.E., McGuinness, K.: Multi-objective interpolation training for robustness to label noise. In: CVPR. pp. 6606–6615 (2021)
- Park, S., Lim, J., Jeon, Y., Choi, J.Y.: Influence-balanced loss for imbalanced visual classification. In: ICCV. pp. 735–744 (2021)

- Patrini, G., Rozza, A., Krishna Menon, A., Nock, R., Qu, L.: Making deep neural networks robust to label noise: A loss correction approach. In: CVPR. pp. 1944– 1952 (2017)
- 39. Reed, S., Lee, H., Anguelov, D., Szegedy, C., Erhan, D., Rabinovich, A.: Training deep neural networks on noisy labels with bootstrapping. In: ICLRW (2015)
- Ren, J., Yu, C., Ma, X., Zhao, H., Yi, S., et al.: Balanced meta-softmax for longtailed visual recognition. In: NeurIPS. vol. 33, pp. 4175–4186 (2020)
- Ren, M., Zeng, W., Yang, B., Urtasun, R.: Learning to reweight examples for robust deep learning. In: ICML. pp. 4334–4343 (2018)
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. IJCV 115(3), 211–252 (2015)
- Shen, L., Lin, Z., Huang, Q.: Relay backpropagation for effective learning of deep convolutional neural networks. In: ECCV. pp. 467–482. Springer (2016)
- Shu, J., Xie, Q., Yi, L., Zhao, Q., Zhou, S., Xu, Z., Meng, D.: Meta-weight-net: Learning an explicit mapping for sample weighting. In: NeurIPS. pp. 1917–1928 (2019)
- Tanaka, D., Ikami, D., Yamasaki, T., Aizawa, K.: Joint optimization framework for learning with noisy labels. In: CVPR. pp. 5552–5560 (2018)
- Wang, P., Han, K., Wei, X.S., Zhang, L., Wang, L.: Contrastive learning based hybrid networks for long-tailed image classification. In: CVPR. pp. 943–952 (2021)
- Wang, R., Wu, Z., Weng, Z., Chen, J., Qi, G.J., Jiang, Y.G.: Cross-domain contrastive learning for unsupervised domain adaptation. IEEE Trans. Multimedia (2022)
- Wang, R., Liu, T., Tao, D.: Multiclass learning with partially corrupted labels. IEEE Trans. NNLS 29(6), 2568–2580 (2017)
- 49. Wang, X., Lian, L., Miao, Z., Liu, Z., Yu, S.: Long-tailed recognition by routing diverse distribution-aware experts. In: ICLR (2021)
- 50. Wei, J., Zhu, Z., Cheng, H., Liu, T., Niu, G., Liu, Y.: Learning with noisy labels revisited: A study using real-world human annotations. In: ICLR (2022)
- Wei, T., Shi, J.X., Li, Y.F., Zhang, M.L.: Prototypical classifier for robust classimbalanced learning. In: PAKDD. pp. 44–57. Springer (2022)
- 52. Wei, T., Shi, J.X., Tu, W.W., Li, Y.F.: Robust long-tailed learning under label noise. arXiv preprint arXiv:2108.11569 (2021)
- 53. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via nonparametric instance discrimination. In: CVPR. pp. 3733–3742 (2018)
- 54. Xia, X., Liu, T., Han, B., Wang, N., Gong, M., Liu, H., Niu, G., Tao, D., Sugiyama, M.: Part-dependent label noise: Towards instance-dependent label noise. In: NeurIPS. vol. 33, pp. 7597–7610 (2020)
- Xiao, T., Xia, T., Yang, Y., Huang, C., Wang, X.: Learning from massive noisy labeled data for image classification. In: CVPR. pp. 2691–2699 (2015)
- Yao, Y., Liu, T., Han, B., Gong, M., Deng, J., Niu, G., Sugiyama, M.: Dual t: Reducing estimation error for transition matrix in label-noise learning. In: NeurIPS. vol. 33, pp. 7260–7271 (2020)
- Zhang, H., Cissé, M., Dauphin, Y.N., Lopez-Paz, D.: Mixup: Beyond empirical risk minimization. In: ICLR (2018)
- Zhao, G., Li, G., Qin, Y., Liu, F., Yu, Y.: Centrality and consistency: Two-stage clean samples identification for learning with instance-dependent noisy labels. In: ECCV. pp. 21–37. Springer (2022)

- 18 J.S. Baik et al.
- Zheltonozhskii, E., Baskin, C., Mendelson, A., Bronstein, A.M., Litany, O.: Contrast to divide: Self-supervised pre-training for learning with noisy labels. In: WACV. pp. 1657–1667 (2022)
- Zhou, B., Cui, Q., Wei, X.S., Chen, Z.M.: Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In: CVPR. pp. 9719–9728 (2020)
- Zhu, J., Wang, Z., Chen, J., Chen, Y.P.P., Jiang, Y.G.: Balanced contrastive learning for long-tailed visual recognition. In: CVPR. pp. 6908–6917 (2022)