

ProMix: Combating Label Noise via Maximizing Clean Sample Utility

Ruixuan Xiao¹, Yiwen Dong¹, Haobo Wang^{1*}, Lei Feng²,
Runze Wu³, Gang Chen¹ and Junbo Zhao¹

¹Zhejiang University, Hangzhou, China

²Nanyang Technological University, Singapore

³NetEase Fuxi AI Lab, Hangzhou, China

{xiaoruixuan, dyw424, wanghaobo, cg, j.zhao}@zju.edu.cn,
lfengqq@gmail.com, wurunze1@corp.netease.com

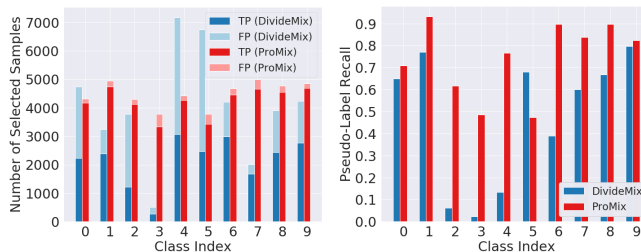
Abstract

Learning with Noisy Labels (LNL) has become an appealing topic, as imperfectly annotated data are relatively cheaper to obtain. Recent state-of-the-art approaches employ specific selection mechanisms to separate clean and noisy samples and then apply Semi-Supervised Learning (SSL) techniques for improved performance. However, the selection step mostly provides a medium-sized and decent-enough clean subset, which overlooks a rich set of clean samples. To fulfill this, we propose a novel LNL framework ProMix that attempts to maximize the utility of clean samples for boosted performance. Key to our method, we propose a matched high confidence selection technique that selects those examples with high confidence scores and matched predictions with given labels to dynamically expand a base clean sample set. To overcome the potential side effect of excessive clean set selection procedure, we further devise a novel SSL framework that is able to train balanced and unbiased classifiers on the separated clean and noisy samples. Extensive experiments demonstrate that ProMix significantly advances the current state-of-the-art results on multiple benchmarks with different types and levels of noise. It achieves an average improvement of 2.48% on the CIFAR-N dataset.

1 Introduction

The great success of deep neural networks (DNNs) attributes to massive and accurately-labeled training data, which is notoriously expensive and time-consuming to obtain. As an alternative, online queries [Blum *et al.*, 2003] and crowdsourcing [Yan *et al.*, 2014] have been widely used to enhance the labeling efficiency. However, due to the lack of expertise, such annotation processes inevitably introduce wrong labels to models and result in performance degradation. To this end, learning with noisy labels (LNL) [Zhang *et al.*, 2017], which aims to mitigate the side-effect of label noise, has attracted huge attention from the community.

*Corresponding Author.



(a) Number of selected samples. (b) Recall of pseudo-labels.

Figure 1: (a) Comparison of selected samples for each class at the end of training on CIFAR-10 dataset with 90% symmetric label noise. TP/FP indicates true/false positive. (b) Comparison of pseudo-label recall for unchosen samples for each class at the end of training on CIFAR-10 dataset with 90% symmetric label noise.

A plethora of works has been developed for the LNL problem. Some early approaches design noise-robust loss [Wang *et al.*, 2019; Wei *et al.*, 2022a]. While some other classical methods utilize the latent noise transition matrix to calibrate the loss function [Patrini *et al.*, 2017; Xia *et al.*, 2019; Zhu *et al.*, 2022]. These methods either disregard specific information about label noise or require strong statistical assumptions and tend to underperform under heavy noise and a large number of classes. Although these methods are theoretically grounded, their performance largely lags behind the fully-supervised counterpart.

Another active line of research [Wei *et al.*, 2020; Li *et al.*, 2020a; Yao *et al.*, 2021] attempts to filter out clean samples for training. For example, the pioneering Co-Teaching approach [Han *et al.*, 2018] figures out that clean samples mostly pose smaller loss values due to the memorization effect of DNNs [Zhang *et al.*, 2017]. Later, DivideMix [Li *et al.*, 2020a] further regards the remaining noisy samples as unlabeled and integrates well-developed semi-supervised learning algorithms. To date, this learning paradigm has been the de-facto standard for the LNL task [Yao *et al.*, 2021; Karim *et al.*, 2022; Cordeiro *et al.*, 2023] and shown to excel compared with early approaches. Despite the tremendous success, current selection-based methods solely ask for a decent-enough clean set with a restricted budget to ensure the high-precision of selected samples. By then, an abundance of valuable clean examples gets overlooked, while they possess

great potential for performance enhancement.

Motivated by this, we present a novel framework ProMix which comprehensively investigates the opportunities and challenges of maximizing the utility of clean samples in LNL. To achieve this, we first propose a progressive selection mechanism that dynamically expands a base class-wise selection set by collecting those having high prediction scores on the observed labels. Empirically, such a simple modification allows us to reach an extraordinary trade-off between the quality and the quantity of the clean set. As shown in Figure 1a, compared with DivideMix [Li *et al.*, 2020a], our progressive selection strategy selects much more clean samples while rarely introducing detested noise samples. Moreover, a simple baseline that is trained merely on such selected clean samples establishes the state-of-the-art performance on the real-world noisy CIFAR-N [Wei *et al.*, 2022b] benchmark, and incorporating semi-supervised learning can further boost the performance. These results clearly validate the importance of well-utilizing clean samples.

Despite the promising results, this maximal selection procedure inevitably causes side effects for the SSL as well. First, the clean samples may not be uniformly distributed across different labels, leading to a biased label distribution. Second, the selection and joint semi-supervised learning procedures are tightly interdependent, which may suffer from the confirmation bias stemming from its own hard-to-correct errors. To mitigate these problems, we develop a debiased semi-supervised training framework that comprises two key components. The first one is an auxiliary pseudo head that decouples the generation and utilization of unreliable pseudo-labels so as to alleviate the model from confirming its error recursively. Second, we incorporate a calibration algorithm that simultaneously rectifies the skewed pseudo-labels and the biased cross-entropy loss. As depicted in Figure 1b, integrated with debiased SSL, ProMix is able to generate pseudo-labels with higher quality on the unselected samples, thus facilitating the model training. By integrating the two key components into a whole framework, our proposed ProMix framework establishes state-of-the-art performance over various LNL benchmarks, even in the *imbalanced LNL scenario* (see Appendix B¹). For example, on the CIFAR-10N datasets with real-world noisy labels, ProMix outperforms the best baseline by 2.04%, 2.11%, and 3.10% under Aggregate, Rand1, and Worst label annotations respectively.

2 Related Work

Learning with Noisy Labels. Existing methods for LNL can be roughly categorized into two types. Earlier LNL approaches focus on loss correction. Since the vanilla cross-entropy (CE) loss has been proved to easily overfit corrupted data [Zhang *et al.*, 2017], a vast amount of robust losses has been designed to achieve a small and unbiased risk for unobservable clean data, including robust variants of the CE loss [Wang *et al.*, 2019; Xu *et al.*, 2019; Ma *et al.*, 2020] or sample reweighted losses [Shu *et al.*, 2019; Zheng *et al.*, 2021]. Some other methods [Patrini *et al.*, 2017; Xia *et al.*, 2019; Zhu *et al.*, 2022] consider the noise transition

matrix and depend on the estimated noise transition matrix to explicitly correct the loss function, but the transition matrix could be hard to estimate in the presence of heavy noise and a large number of classes.

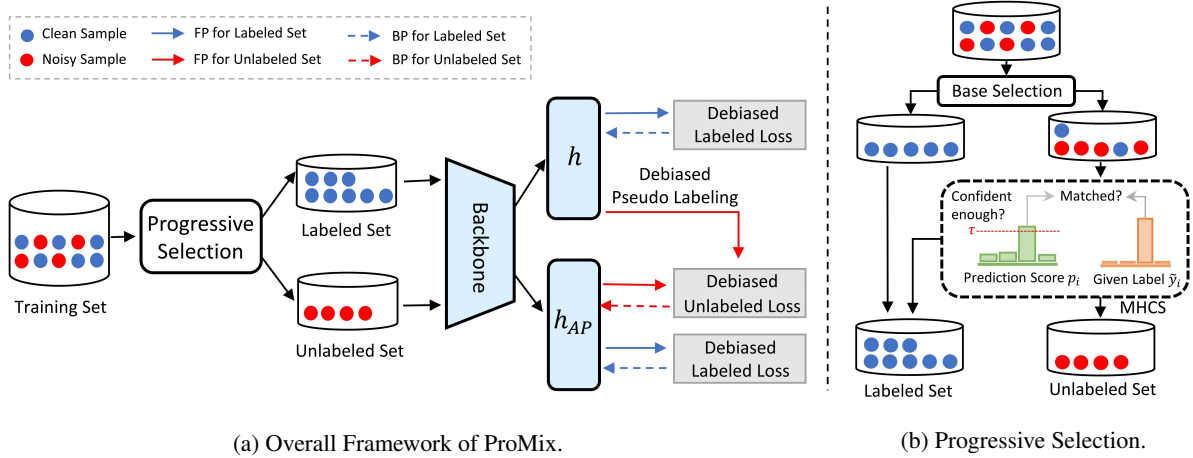
The latter strand of LNL methods tries to filter out a clean subset for robust training [Li *et al.*, 2020a; Karim *et al.*, 2022; Cordeiro *et al.*, 2023]. Early methods [Han *et al.*, 2018; Jiang *et al.*, 2018; Wei *et al.*, 2020] typically adopt a peer network and perform small loss selection that is motivated by the memorization effect of DNNs [Zhang *et al.*, 2017]. The seminal work DivideMix [Li *et al.*, 2020a] further incorporates MixMatch [Berthelot *et al.*, 2019] to leverage the unselected examples and obtain extraordinary performance. This paradigm is then followed by the most recent LNL approaches [Li *et al.*, 2020b; Liu *et al.*, 2020; Bai *et al.*, 2021; Karim *et al.*, 2022]. Recent LNL works also incorporate contrastive learning methods [Ortego *et al.*, 2021; Yao *et al.*, 2021]. Despite the promise, existing selection-based methods mostly select a mediocre clean set with medium size and restricted budget, which overlooks a flurry of clean samples and naturally exhibits a quality-quantity trade-off. A too-small budget leads to inadequate clean sample utility, while a too-large budget enrolls many unexpected noisy examples. To cope with this, our work targets to maximize the utility of clean samples to unleash their great potential.

Debiased Learning. Debiased learning has gained huge attention since the source dataset and algorithms can be naturally biased [Kim *et al.*, 2019; Nam *et al.*, 2020]. For example, in long-tailed learning, models can be easily biased towards dominant classes due to the distribution bias in source data [Zhang *et al.*, 2021a]. As for semi-supervised learning, besides the distribution bias in source data and pseudo-labels [Wang *et al.*, 2022], confirmation bias caused by inappropriate training with unreliable pseudo-labels may jeopardize generalization performance [Chen *et al.*, 2022]. Debiased learning has also been a prevailing topic discussed in some other research topics like contrastive sampling [Chuang *et al.*, 2020] and OOD detection [Ming *et al.*, 2022]. Despite the surge in other areas, the bias in LNL still remains underexplored. In this work, we attempt to alleviate the biases in our selection and pseudo-labeling procedures. We hope our work will inspire future works to integrate this important component in advanced LNL frameworks.

3 Proposed Method

Problem Formulation. For a multi-class classification task with C classes, we denote the noisy training dataset with n examples by $\mathcal{D} = \{(\mathbf{x}_i, \tilde{y}_i)\}_{i=1}^n$, where each tuple comprises of an image $\mathbf{x}_i \in \mathcal{X} \subset \mathbb{R}^d$ and the corresponding *noisy label* $\tilde{y}_i \in \mathcal{Y} = \{1, \dots, C\}$. In contrast to the supervised learning setup, the ground-truth labels are not observable. The goal of LNL is to obtain a functional mapping $f_{\theta, \phi} : \mathcal{X} \mapsto \mathcal{Y}$ that takes an example \mathbf{x}_i as the input and outputs the logits $f_{\theta, \phi}(\mathbf{x}_i)$ for predicting the true labels. This classification model $f_{\theta, \phi} = h_{\phi} \circ g_{\theta}$ consists of a feature extractor g_{θ} and a classification head h_{ϕ} , which are parameterized by θ and ϕ respectively. In our work, we adopt the cross-entropy loss $\mathcal{H} = -\sum_{j=1}^C \tilde{y}_i^j \log p_i^j$ where $\mathbf{p}_i = \text{softmax}(f(\mathbf{x}_i))$ is the

¹<https://arxiv.org/abs/2207.10276>



(a) Overall Framework of ProMix.

(b) Progressive Selection.

Figure 2: **The left:** Overall framework of ProMix. After progressive selection, the source dataset is divided into the labeled set and unlabeled set for debiased SSL. The labeled data are fed forward to both primary head h and auxiliary pseudo h_{AP} , while the unlabeled data are provided with debiased pseudo-labels by h and then propagated forward through h_{AP} . Finally, the standard cross entropy in classification loss is replaced by debiased margin-based loss. FP/BP indicates forward/backward propagation. **The right:** Details of progressive selection. The MHCS selects those examples with high confidence and matched predictions with given labels to dynamically expand a base selection set.

probabilistic output of $f(x_i)$. p_i^j is the j -th entry of \mathbf{p}_i and \tilde{y}_i^j is the j -th entry of one-hot encoded label vector $\tilde{\mathbf{y}}_i$.

Similar to previous works, our framework also divides the samples into a clean set \mathcal{D}_l and a noisy set \mathcal{D}_u , which are then applied for semi-supervised learning. The vanilla classification losses for these two parts are formulated:

$$\begin{aligned} \mathcal{L}_x(\mathcal{D}_l) &= \frac{1}{|\mathcal{D}_l|} \sum_{\mathbf{x}_i \in \mathcal{D}_l} \mathcal{H}(\tilde{\mathbf{y}}_i, \mathbf{p}_i) \\ \mathcal{L}_u(\mathcal{D}_u) &= \frac{1}{|\mathcal{D}_u|} \sum_{\mathbf{x}_i \in \mathcal{D}_u} \mathcal{H}(\text{Sharpen}(\mathbf{p}_i, T), \mathbf{p}_i) \end{aligned} \quad (1)$$

where $\text{Sharpen}(\mathbf{p}_i, T) = \frac{(p_i^c)^{\frac{1}{T}}}{\sum_{c=1}^C (p_i^c)^{\frac{1}{T}}}$ for $c = 1, 2, \dots, C$. An overview of the proposed method is illustrated in Figure 2.

3.1 Progressive Sample Selection

As aforementioned, we would like to excessively collect clean samples. To achieve this, we first draw inspiration from the classical learning theory,

Theorem 1. [Shalev-Shwartz and Ben-David, 2014] Let \mathcal{F} be the hypothesis set of f . Denote the population risk by $R(f) = \mathbb{E}[\mathcal{H}]$ and the corresponding empirical risk by $\hat{R}(f)$. We define the true hypothesis by $f^* = \arg \min R(f)$ and empirical risk minimizer by $\hat{f} = \arg \min \hat{R}(f)$. Given a dataset containing n examples, with probability at least $1 - \delta$, the following inequality holds:

$$R(\hat{f}) \leq R(f^*) + 2\hat{\mathfrak{R}}(\mathcal{H}) + 5\sqrt{\frac{2 \log(8/\delta)}{n}} \quad (2)$$

where $\hat{\mathfrak{R}}(\mathcal{F})$ is the empirical Rademacher of \mathcal{F} .

Considering a simple case in which f consists of multiple ReLU-activated fully-connected layers, we can apply the techniques outlined in [Allen-Zhu *et al.*, 2019] to derive an

upper bound of the Rademacher complexity $O(\sqrt{\frac{1}{n}})$. If we assume a perfectly clean set with n_c samples, the model has larger than a probability of $1 - O(e^{-n_c \epsilon^2 / 50})$ to give an ϵ -wrong prediction on the remaining dirty training data. Hence, this model is prone to produce a much larger probability on the observed labels for the remaining correctly-labeled data than noisy ones. Motivated by this, we present the following selection procedure for enlarging the base selection set.

Class-wise Small-loss Selection (CSS). Based on the small loss criterion, we aim to first filter out a base clean set as the cornerstone. In the training, the patterns of easy classes are usually fitted earlier and the loss values of examples with different observed labels may not be comparable [Gui *et al.*, 2021]. So it would be more appropriate to select examples class by class with the same observed labels. At each epoch, we first split the whole training dataset \mathcal{D} to C sets according to the noisy labels, i.e. $\mathcal{S}_j = \{(\mathbf{x}_i, \tilde{\mathbf{y}}_i) \in \mathcal{D} | \tilde{\mathbf{y}}_i = j\}$. Given the j -th set \mathcal{S}_j , we calculate the standard CE loss l_i for each example and select $k = \min(\lceil \frac{n}{C} \times R \rceil, |\mathcal{S}_j|)$ examples with smallest l_i to constitute the clean set \mathcal{C}_j , where R is the filter rate that is identical for all classes. Finally, the integral clean set is merged as $\mathcal{D}_{CSS} = \cup_{j=1}^C \mathcal{C}_j$. Compared to the original small-loss selection, we relate k to the average number of examples n/C such that produce a roughly balanced base set.

Matched High Confidence Selection (MHCS). After getting the base clean set, we also introduce another selection strategy to make capital out of the potentially clean samples missed by CSS. Specifically, we calculate the confidence scores $e_i = \max_j p_i^j$ for remaining samples. Then, we select those examples whose predictions match the given labels, while their confidence is high enough,

$$\mathcal{D}_{MHCS} = \{(\mathbf{x}_i, \tilde{\mathbf{y}}_i) \in \mathcal{D} | e_i \geq \tau, y'_i = \tilde{\mathbf{y}}_i\} \quad (3)$$

where $y'_i = \arg \max_j p'_j$ is the predicted label of x_i . In practice, we set a high threshold τ such that the selected samples have a high probability of being clean. According to our theoretical motivation, the proposed MHCS mechanism is able to automatically explore new clean data to dynamically enlarge the clean set $\mathcal{D}_l = \mathcal{D}_{CSS} \cup \mathcal{D}_{MHCS}$ by trusting those examples with matched high confidence.

Remark. Our MHCS selection strategy draws inspiration from the pseudo-labeling procedure of FixMatch [Sohn *et al.*, 2020]. But, in contrast to the SSL setup, our clean set is dynamically generated. With the naive high confidence selection, the DNNs can assign an arbitrary label with high confidence and select them as clean, which results in a vicious cycle. Thus, we choose those samples to have high confidence in their given labels, which tend to be originally clean.

3.2 Debiased Semi-Supervised Training

After sample selection, we would also like to integrate semi-supervised learning that utilizes the remaining noisy samples to boost performance. Nevertheless, our excessive selection procedure inevitably induces unexpected characteristics. First, the selected clean samples may inherently be imbalanced amongst different labels, since some labels are typically more ambiguous than others. Hence, the model may be biased towards the classes that possess more clean samples and in turn pollutes the quality of selection (as shown in Figure 1a). Second, the selection and pseudo-labeling are tightly coupled, which naturally brings about confirmation bias, so we should be particularly careful to train the model with the generated pseudo-labels. To escape such a dilemma, we introduce our debiased SSL framework in what follows.

Mitigating Confirmation Bias. In most SSL pipelines, the generation and utilization of pseudo-labels are usually achieved by the same network blocks or between interdependent peer networks. As a result, the mistakes in pseudo-labels are nearly impossible to self-correct and can accumulate amplifying the existing confirmation bias. In order to alleviate this, we aim to disentangle the generation and utilization of pseudo-labels, which is achieved via introducing an extra classifier termed Auxiliary Pseudo Head (APH) h_{AP} . The primary classifier h is the task-specific head that generates pseudo-labels for the unlabeled data but is merely optimized with labeled samples in \mathcal{D}_l . The newly devised classifier h_{AP} share the same representations with h but have independent parameters. It receives the pseudo-labels from h and is trained with both labeled set \mathcal{D}_l and unlabeled set \mathcal{D}_u . Based on Eq. (1), the classification loss is formulated,

$$\mathcal{L}_{cls} = \underbrace{\mathcal{L}_x(\{\tilde{y}_i, \mathbf{p}_i\})}_{\text{original head}} + \underbrace{\mathcal{L}_x(\{\tilde{y}_i, \mathbf{p}'_i\}) + \lambda_u \mathcal{L}_u(\{\text{Sharpen}(\mathbf{p}_i, T), \mathbf{p}'_i\})}_{\text{auxiliary pseudo head}} \quad (4)$$

where \mathbf{p}_i is the softmax score from h and \mathbf{p}'_i is that from h_{AP} . λ_u controls the strength of the unsupervised loss. Through this procedure, the backbone benefits from the unlabeled data for enhanced representation via h_{AP} . At the same time, the primary classifier h remains reliable without seeing unlabeled samples, which reduces the confirmation bias. It is worth noting that the auxiliary pseudo head only participates in training and is discarded at inference time.

Mitigating Distribution Bias. In addition to confirmation bias, there may still be distribution bias whose causes are two-fold. First, the selected samples can falsely calibrate the classifier due to skewed label distribution. Second, previous work [Wang *et al.*, 2022] shows that pseudo-labels can naturally bias towards some easy classes even if the training data are curated and balanced, which can iteratively dominate the training process. To learn a well-calibrated classifier, we incorporate a Debiased Margin-based Loss (DML) l_{DML} to encourage a larger margin between the sample-rich classes and the sample-deficient classes. l_{DML} can be defined as:

$$l_{DML} = - \sum_{j=1}^C \tilde{y}_i^j \log \frac{e^{f^j(\mathbf{x}_i) + \alpha \cdot \log \pi_j}}{\sum_{k=1}^C e^{f^k(\mathbf{x}_i) + \alpha \cdot \log \pi_k}} \quad (5)$$

where α is a tuning parameter that controls debiasing strength and π is the underlying class distribution $\mathbb{P}(y | \mathbf{x})$ indicating which label possesses more samples on the training set. Here we assume π is known and will show how to estimate it later. The standard cross entropy loss \mathcal{H} is then replaced by l_{DML} .

Similarly, we also resort to Debiased Pseudo Labeling (DPL) following [Menon *et al.*, 2021] to calibrate the pseudo-labels. The refined logit \tilde{f}_i which is then fed into the softmax layer for subsequent pseudo-label generation is formulated:

$$\tilde{f}_i = f(\mathbf{x}_i) - \alpha \log \pi \quad (6)$$

That is, we suppress the logits on those easy classes and ensure other classes are fairly learned. The refined pseudo-labels are then used for calculating unlabeled loss in Eq. (4).

Next, we elaborate on our dynamical estimation of π . Notably, we employ two different class distribution vectors for the labeled and unlabeled data respectively. Formally, given a mini-batch of samples \mathcal{B} from either \mathcal{D}_l or \mathcal{D}_u , we update the estimation of class distribution π in a moving-average style:

$$\pi = m\pi + (1 - m) \frac{1}{|\mathcal{B}|} \sum_{\mathbf{x}_i \in \mathcal{B}} \mathbf{p}_i \quad (7)$$

where $m \in [0, 1)$ is a momentum coefficient. When calculating the loss in Eq. (4), we use different estimation π_l and π_u for the two sets. For better understanding of the whole procedure, we refer the readers to Appendix A for more details.

To summarize the pipeline of debiased SSL, the labeled data are fed forward to both h and h_{AP} while the unlabeled data are provided with the debiased pseudo-labels by h as Eq. (6) and then propagated forward through h_{AP} . Finally, the standard cross entropy in Eq. (4) is replaced by l_{DML} . Through the above debiasing procedure, the refined pseudo-labels with higher quality and enhanced tolerance for incorrect pseudo-labels assist in milking unlabeled samples and in return enhance the purity of sample selection.

3.3 Practical Implementation

Label Guessing by Agreement (LGA). While we still anticipate selecting matched high confidence instances, there are no free labels to be matched on falsely-labeled data whose predictions mismatch the observed labels. To further distill potential ‘‘clean’’ samples from unlabeled data, we perform label guessing to correct noisy labels on the unlabeled set \mathcal{D}_u .

Dataset	CIFAR-10					CIFAR-100			
	Sym.				Asym.	Sym.			
Methods\Noise Ratio	20%	50%	80%	90%	40%	20%	50%	80%	90%
CE	86.8	79.4	62.9	42.7	85.0	62.0	46.7	19.9	10.1
Co-Teaching+	89.5	85.7	67.4	47.9	-	65.6	51.8	27.9	13.7
JoCoR	85.7	79.4	27.8	-	76.4	53.0	43.5	15.5	-
M-correction	94.0	92.0	86.8	69.1	87.4	73.9	66.1	48.2	24.3
PENCIL	92.4	89.1	77.5	58.9	88.5	69.4	57.5	31.1	15.3
DivideMix	96.1	94.6	93.2	76.0	93.4	77.3	74.6	60.2	31.5
ELR+	95.8	94.8	93.3	78.7	93.0	77.6	73.6	60.8	33.4
LongReMix	96.2	95.0	93.9	82.0	94.7	77.8	75.6	62.9	33.8
MOIT	94.1	91.1	75.8	70.1	93.2	75.9	70.1	51.4	24.5
SOP+	96.3	95.5	94.0	-	93.8	78.8	75.9	63.3	-
PES(semi)	95.9	95.1	93.1	-	-	77.4	74.3	61.6	-
ULC	96.1	95.2	94.0	86.4	94.6	77.3	74.9	61.2	34.5
ProMix(last)	97.59	97.30	95.05	91.13	96.51	82.39	79.72	68.95	42.74
ProMix(best)	97.69	97.40	95.49	93.36	96.59	82.64	80.06	69.37	42.93

Table 1: Accuracy comparisons on CIFAR-10/100 with symmetric (20%-90%) and asymmetric noise (40%). We report both the averaged test accuracy over last 10 epochs and the best accuracy of ProMix. Results of previous methods are the best test accuracies cited from their original papers, where the blank ones indicate that the corresponding results are not provided. **Bold entries** indicate superior results.

We train two peer networks f_1 and f_2 in parallel to revise the noisy labels by agreement. Then we select a new subset \mathcal{D}_c to enlarge the clean set $\mathcal{D}_l = \mathcal{D}_l \cup \mathcal{D}_c$ by the condition,

$$\begin{aligned} & (\max_j f_1^j(\mathbf{x}) \geq \tau) \wedge (\max_j f_2^j(\mathbf{x}) \geq \tau) \wedge \\ & (\arg \max_j f_1^j(\mathbf{x}) == \arg \max_j f_2^j(\mathbf{x})) \end{aligned} \quad (8)$$

In other words, the two networks should both hold high confidence on the same label. After that, we take the shared prediction $y' = \arg \max_j f_1^j(\mathbf{x})$ as the guessed labels for \mathcal{D}_c . With such reliable refurbishment, more purified clean samples could be involved in \mathcal{D}_l to facilitate the SSL training.

Training Objective. To improve the representation ability of ProMix, we involve consistency training along with mixup augmentation on the reliable labeled set \mathcal{D}_l . See Appendix A for more details. The losses for consistency regularization and mixup augmentation are denoted as \mathcal{L}_{cr} and \mathcal{L}_{mix} respectively. The overall training loss is given by,

$$\mathcal{L}_{total} = \mathcal{L}_{cls} + \gamma(\mathcal{L}_{cr} + \mathcal{L}_{mix}) \quad (9)$$

where γ is a weighting factor. Notably, we ensemble the outputs of the two peer networks at the inference time.

4 Experiments

In this section, we present the main results and part of the ablation results to demonstrate the effectiveness of the ProMix framework. We put more experimental results, including results on imbalanced dataset with synthetic noise, results on dataset with instance-dependent label noise and more results for ablation experiment in Appendix B.

4.1 Setup

Datasets. We first evaluate the performance of ProMix on CIFAR-10, CIFAR-100 [Krizhevsky *et al.*, 2009], Clothing1M [Xiao *et al.*, 2015] and ANIMAL-10N [Song *et al.*,

2019] dataset. For CIFAR-10/100, we conduct experiments with synthetic symmetric and asymmetric label noise following the previous protocol [Tanaka *et al.*, 2018]. Symmetric noise is introduced by randomly flipping each ground-truth label to all possible labels with probability r and asymmetric noise mimics structure of realistic noise taking place in similar classes. We also verify the effectiveness of ProMix on the recently proposed real-world noisy CIFAR-N [Wei *et al.*, 2022b] dataset. The CIFAR-N dataset is composed of CIFAR-10N and CIFAR-100N, which equips the training samples of CIFAR-10 and CIFAR100 with human-annotated real-world noisy labels collected from Amazon Mechanical Turk. We adopt three noise types, Aggregate, Rand1 and Worst with a noise rate of 9.03%, 17.23% and 40.21% respectively for CIFAR-10N and one noise type Noisy-Fine whose noise rate is 40.20% for CIFAR-100N. Clothing1M dataset is a large-scale real-world noisy dataset containing 1 million images crawled from online shopping websites whose labels are collected by extracting tags from the surrounding texts. ANIMAL-10N is also a real-world noisy dataset consisting of human-annotated online images of ten confusable animals.

Implementation details. For CIFAR experiments, we use two ResNet-18 as the backbone of the peer networks. These networks are trained for 600 epochs, with a warm-up period of 10 epochs for CIFAR-10 and 30 epochs for CIFAR-100. We employ SGD as the optimizer with a momentum of 0.9 and weight decay of $5e^{-4}$. The batch size is fixed as 256 and the initial learning rate is 0.05, which decays by a cosine scheduler. The threshold τ is set as 0.99 and 0.95 for CIFAR-10 and CIFAR-100. The debiasing factor α is set as 0.8 and 0.5 for CIFAR-10 and CIFAR-100. m is fixed as 0.9999. We leverage RandAugment to generate a strongly augmented view for consistency training. The filter rate R is set as $\{0.7, 0.5, 0.2, 0.1\}$ for symmetric noise with noise ratio $r = \{0.2, 0.5, 0.8, 0.9\}$ respectively and 0.5 for asym-

Method	CE	PENCIL	JoCoR	DivideMix	ELR+	LongReMix	CORES*	SOP	PES	ULC	ProMix
Accuracy	69.21	73.49	70.30	74.76	74.81	74.38	73.24	73.50	74.64	74.90	74.94

 Table 2: Accuracy comparisons on Clothing1M dataset. **Bold entries** indicate superior results.

Methods	CIFAR-10N			CIFAR-100N
	Aggre	Rand1	Worst	Noisy Fine
CE	87.77	85.02	77.69	55.50
Co-Teaching+	91.20	90.33	83.83	60.37
JoCoR	91.44	90.30	83.37	59.97
DivideMix	95.01	95.16	92.56	71.13
ELR+	94.83	94.43	91.09	66.72
CORES*	95.25	94.45	91.66	55.72
SOP+	95.61	95.28	93.24	67.81
PES(Semi)	94.66	95.06	92.68	70.36
ProMix	97.65	97.39	96.34	73.79

 Table 3: Accuracy comparisons on CIFAR-10N and CIFAR-100N under different noise types. **Bold entries** indicate superior results.

metric noise and all settings On CIFAR-N. LGA is performed at the last $K = 250$ epochs such that the networks are accurate enough. We also set a maximum ratio threshold of total selected samples $\rho = 0.9$ for all settings to restrict excessive correction. We linearly ramp up γ from 0 to 1 and λ_u from 0 to 0.1 to avoid overfitting false labels at the beginning.

For Clothing1M, following previous work [Li *et al.*, 2020a], the backbone is ResNet-50 pre-trained on ImageNet and the training is conducted for 80 epochs. We set the learning rate as 0.002 and weight decay at 0.001. The filter rate R is set as 0.7, K is fixed as 40 and α is set at 0.8. We leverage AutoAugment [Cubuk *et al.*, 2019] with ImageNet-Policy for augmentation. For ANIMAL-10N, we adopt VGG-19 [Simonyan and Zisserman, 2015] with batch-normalization as the backbone following [Song *et al.*, 2019]. We use SGD with the same settings as CIFAR for 300 epochs of training. The filter rate R is set as 0.7. See Appendix for more details.

4.2 Main Results

We compare ProMix with multiple state-of-the-art methods. Specifically, they are Co-Teaching+ [Yu *et al.*, 2019], M-correction [Arazo *et al.*, 2019], PENCIL [Yi and Wu, 2019], SELFIE [Song *et al.*, 2019], JoCoR [Wei *et al.*, 2020], DivideMix [Li *et al.*, 2020a], MOIT [Ortego *et al.*, 2021], PLC [Zhang *et al.*, 2021b], NCT [Chen *et al.*, 2021], ELR+ [Liu *et al.*, 2020], JoCoR [Wei *et al.*, 2020], CORES* [Cheng *et al.*, 2021], PES [Bai *et al.*, 2021], SOP+ [Liu *et al.*, 2022], SELC [Lu and He, 2022], ULC [Huang *et al.*, 2022], LongReMix [Cordeiro *et al.*, 2023] and the naive cross-entropy (CE) method. For their performance, we directly adopt the reported results from their receptive papers. For fair comparisons, we follow the standard LNL evaluation protocols [Li *et al.*, 2020a; Xia *et al.*, 2019] and adopt the same backbone as in previous works for all benchmarks.

Comparison with synthetic noisy labels. Table 1 shows the results on CIFAR-10/100 dataset with different types and

Method	CE	SELFIE	PLC	NCT	SELC	ProMix
Accuracy	79.4	81.8	83.4	84.1	83.7	89.98

Table 4: Accuracy comparisons on ANIMAL-10N dataset.

levels of synthetic label noise. ProMix significantly outperforms all the rivals by a notable margin with different settings. Specifically, on the CIFAR-10 dataset with symmetric noise, ProMix exceeds the best baseline by 1.39%, 1.90%, 1.49% and 6.96% under the noise ratio of 20%, 50%, 80%, 90% respectively. Moreover, most baselines exhibit significant performance degradation under heavier noise, whereas ProMix displays robustness and remains competitive. These observations clearly verify the effectiveness of ProMix.

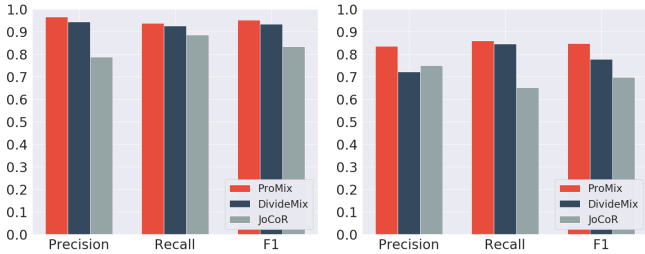
Comparison with real-world noisy labels. As shown in Table 3, on the CIFAR-N dataset with real-world noisy labels, our proposed ProMix achieves the best performance compared with other existing methods on all four settings. In specific, on the CIFAR-10N dataset, the absolute accuracy gains for ProMix are 2.04%, 2.11% and 3.10% under Aggregate, Rand1, and Worst label annotations. Additionally, ProMix further expand the lead and reach a performance gain at 2.66% on the even more challenging CIFAR-100N dataset. Table 2 and Table 4 show that ProMix surpasses all the previous methods on the large-scale Clothing1M and ANIMAL-10N dataset, particularly achieves a lead of 5.88% on ANIMAL-10N. These results clarify that our method can be well applied in real-world scenarios with large-scale data.

4.3 Ablation Study

Effect of sample selection. We first dismiss the label guessing in variant *ProMix w/o LGA*, where we note that LGA could boost performance, especially under severe noise. Based on this variant, we further equip ProMix with different sample selection strategies: 1) *ProMix w/o MHCS* which performs CSS and overlooks the remaining matched high confidence samples; 2) *ProMix w/o Base Selection* which discards the base selection and only selects matched high confidence samples. From Table 5, we can observe that ProMix w/o MHCS largely underperforms ProMix, which verifies the efficacy of MHCS mechanism. ProMix w/o Base Selection cannot establish a competitive performance because of inadequate labeled samples. This also proves that a clean-enough base selection set is the foundation of further progressive selection and debiased SSL. Moreover, We test a variant *ProMix with Only Clean* which merely employs the selected samples and disregards the unselected samples. This variant suffers from perceptible performance degradation, but still beats all the baselines on CIFAR-N by a large margin. Such result indicates that clean samples are the real treasures in LNL, well-utilizing which is the key to successful learning.

Ablation	LGA	Selection Strategy	CIFAR-10 Sym.80%	CIFAR-100 Sym.80%	CIFAR-10N Worst	CIFAR-100N Noisy Fine
ProMix	✓	CSS+MHCS	95.05	68.95	96.34	73.79
w/o LGA	✗	CSS+MHCS	94.32	61.19	95.57	73.10
w/o MHCS	✗	CSS	93.17	60.35	95.11	70.40
w/o Base Selection	✗	MHCS	39.59	21.01	74.09	40.75
w/o CBR	✗	CSS+MHCS	93.96	60.72	95.26	72.61
w/o DBR	✗	CSS+MHCS	94.07	60.91	95.07	72.51
with Only Clean	✗	CSS+MHCS	93.82	60.35	95.18	72.06

Table 5: Ablation study of ProMix on CIFAR-10-Symmetric 80%, CIFAR-100-Symmetric 80%, CIFAR-10N-Worst and CIFAR-100N-Noisy. ProMix with Only Clean denotes training with merely selected samples. CBR/DBR indicates Confirmation/Distribution Bias Removal.

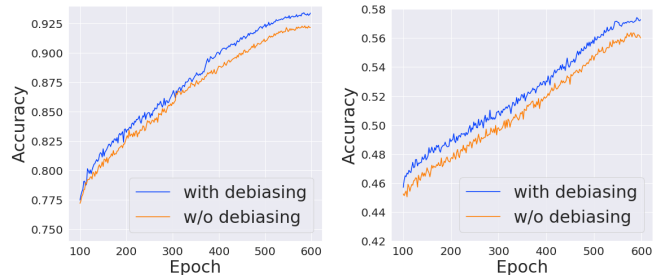


(a) CIFAR-10-Symmetric 80%. (b) CIFAR-100-Symmetric 80%.

Figure 3: Comparison of clean sample selection on CIFAR-10/100 dataset with 80% symmetric noise. The threshold of DivideMix and forget rate of JoCoR have been re-tuned to get the best F1 score.

Effect of bias mitigation. Next, we ablate the contributions of the bias mitigation components in the debiased SSL. Specifically, we compare ProMix with two variants: 1) *ProMix w/o Confirmation Bias Removal (CBR)* which discards h_{AP} , generating and utilizing the pseudo-labels on the same classification head; 2) *ProMix w/o Distribution Bias Removal (DBR)* which sticks to the vanilla pseudo-labeling and adopts the standard cross entropy loss. From Table 5, we can observe that the lack of any one of these two components will bring about performance degradation, which indicates the effectiveness of bias mitigation. The performance drop grows more severe in imbalance scenarios as shown in Appendix. Moreover, to reveal the impact of debiasing more intuitively, we compare the accuracy of the pseudo-labels of unlabeled data during training with and without incurring the bias mitigation strategies. As shown in Figure 4, we can see that integrated with the debiasing strategies, the pseudo-labels on unlabeled data grow more accurate, hence facilitating the training process and reciprocating the sample selection process.

ProMix improves the clean sample utility. To further prove that ProMix indeed improves the clean sample utility, we compare the precision and recall of filtered clean samples of ProMix and the other two baselines, DivideMix [Li *et al.*, 2020a] and JoCoR [Wei *et al.*, 2020]. From Figure 3, we observe that ProMix significantly outperforms its competitors in both precision and recall, that is, ProMix selects much more clean samples while introducing fewer noisy samples. In specific, ProMix achieves a precision of 96.59% and a recall of 93.78% on CIFAR-10-Sym 80%, with a lead of 2.23%



(a) CIFAR-10-Symmetric 80%. (b) CIFAR-100-Symmetric 80%.

Figure 4: Accuracy of pseudo-labels for the unlabeled data with and without debiasing on CIFAR-10/100 with 80% Symmetric noise. LGA is disabled in order to avoid unexpected interference.

in precision and 1.15% in recall compared to DivideMix. In more straightforward terms, ProMix picks 450 more clean samples and 891 fewer false positive noise samples than DivideMix, most of which are precious and meaningful hard samples, greatly facilitating training. Besides, on CIFAR-100-Sym 80% where performances of competitors decline, ProMix leads by an even wider margin, especially in precision. This indicates that ProMix is based on a pipeline that first selects a clean-enough set with high precision, and then tries to expand the clean set with MHCS while still ensuring high precision, gradually maximizing clean sample utility. See Appendix B for more results of clean sample utility.

5 Conclusion

In this work, we propose a novel framework called ProMix for learning with noisy labels with the principle of maximizing the utility of clean examples. ProMix dynamically expands a balanced base selection set by progressively collecting those samples with matched high confidence for the subsequent SSL. We also incorporate an auxiliary pseudo head and a calibration mechanism to mitigate the confirmation bias and the distribution bias in the process of training. The proposed framework is simple but achieves significantly superior performance to the state-of-the-art LNL techniques on different datasets with various kinds and levels of noise. We hope our work can inspire future research along this direction to promote the LNL methods by well utilizing the clean samples, which are the true treasures of the data.

Acknowledgements

Junbo Zhao wants to thank the support by the NSFC Grants (No. 62206247) and the Fundamental Research Funds for the Central Universities. Lei Feng is supported by the Joint NTU-WeBank Research Centre on Fintech (Award No: NWJ-2021-005), Nanyang Technological University, Singapore.

References

- [Allen-Zhu *et al.*, 2019] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *ICML*, volume 97, pages 242–252. PMLR, 2019.
- [Arazo *et al.*, 2019] Eric Arazo, Diego Ortego, Paul Albert, Noel E. O’Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. In *ICML*, volume 97, pages 312–321. PMLR, 2019.
- [Bai *et al.*, 2021] Yingbin Bai, Erkun Yang, Bo Han, Yanhua Yang, Jiatong Li, Yinian Mao, Gang Niu, and Tongliang Liu. Understanding and improving early stopping for learning with noisy labels. In *NeurIPS*, pages 24392–24403, 2021.
- [Berthelot *et al.*, 2019] David Berthelot, Nicholas Carlini, Ian J. Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *NeurIPS*, pages 5050–5060, 2019.
- [Blum *et al.*, 2003] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50(4):506–519, 2003.
- [Chen *et al.*, 2021] Yingyi Chen, Xi Shen, Shell Xu Hu, and Johan A. K. Suykens. Boosting co-teaching with compression regularization for label noise. In *CVPR*, pages 2688–2692. IEEE, 2021.
- [Chen *et al.*, 2022] Baixu Chen, Janguang Jiang, Ximei Wang, Pengfei Wan, Jianmin Wang, and Mingsheng Long. Debaised self-training for semi-supervised learning. In *NeurIPS*, 2022.
- [Cheng *et al.*, 2021] Hao Cheng, Zhaowei Zhu, Xingyu Li, Yifei Gong, Xing Sun, and Yang Liu. Learning with instance-dependent label noise: A sample sieve approach. In *ICLR*. OpenReview.net, 2021.
- [Chuang *et al.*, 2020] Ching-Yao Chuang, Joshua Robinson, Yen-Chen Lin, Antonio Torralba, and Stefanie Jegelka. Debaised contrastive learning. In *NeurIPS*, 2020.
- [Cordeiro *et al.*, 2023] Filipe R. Cordeiro, Ragav Sachdeva, Vasileios Belagiannis, Ian D. Reid, and Gustavo Carneiro. Longremix: Robust learning with high confidence samples in a noisy label environment. *Pattern Recognit.*, 133:109013, 2023.
- [Cubuk *et al.*, 2019] Ekin D. Cubuk, Barret Zoph, Dandelion Mané, Vijay Vasudevan, and Quoc V. Le. Autoaugment: Learning augmentation strategies from data. In *CVPR*, pages 113–123. IEEE, 2019.
- [Gui *et al.*, 2021] Xian-Jin Gui, Wei Wang, and Zhang-Hao Tian. Towards understanding deep learning from noisy labels with small-loss criterion. In *IJCAI*, pages 2469–2475. ijcai.org, 2021.
- [Han *et al.*, 2018] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor W. Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, pages 8536–8546, 2018.
- [Huang *et al.*, 2022] Yingsong Huang, Bing Bai, Shengwei Zhao, Kun Bai, and Fei Wang. Uncertainty-aware learning against label noise on imbalanced datasets. In *AAAI*, pages 6960–6969. AAAI Press, 2022.
- [Jiang *et al.*, 2018] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, volume 80, pages 2309–2318. PMLR, 2018.
- [Karim *et al.*, 2022] Nazmul Karim, Mamshad Nayeem Rizve, Nazanin Rahnavard, Ajmal Mian, and Mubarak Shah. UNICON: combating label noise through uniform selection and contrastive learning. In *CVPR*, pages 9666–9676. IEEE, 2022.
- [Kim *et al.*, 2019] Byungju Kim, Hyunwoo Kim, Kyungsu Kim, Sungjin Kim, and Junmo Kim. Learning not to learn: Training deep neural networks with biased data. In *CVPR*, pages 9012–9020. IEEE, 2019.
- [Krizhevsky *et al.*, 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Master’s thesis, Department of Computer Science, University of Toronto*, 2009.
- [Li *et al.*, 2020a] Junnan Li, Richard Socher, and Steven C. H. Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *ICLR*. OpenReview.net, 2020.
- [Li *et al.*, 2020b] Mingchen Li, Mahdi Soltanolkotabi, and Samet Oymak. Gradient descent with early stopping is provably robust to label noise for overparameterized neural networks. In *AISTATS*, volume 108, pages 4313–4324. PMLR, 2020.
- [Liu *et al.*, 2020] Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. In *NeurIPS*, 2020.
- [Liu *et al.*, 2022] Sheng Liu, Zihui Zhu, Qing Qu, and Chong You. Robust training under label noise by over-parameterization. *CoRR*, abs/2202.14026, 2022.
- [Lu and He, 2022] Yangdi Lu and Wenbo He. SELC: self-ensemble label correction improves learning with noisy labels. In *IJCAI*, pages 3278–3284. ijcai.org, 2022.
- [Ma *et al.*, 2020] Xingjun Ma, Hanxun Huang, Yisen Wang, Simone Romano, Sarah M. Erfani, and James Bailey. Normalized loss functions for deep learning with noisy labels. In *ICML*, volume 119, pages 6543–6553. PMLR, 2020.
- [Menon *et al.*, 2021] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas

- Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. In *ICLR*. OpenReview.net, 2021.
- [Ming *et al.*, 2022] Yifei Ming, Hang Yin, and Yixuan Li. On the impact of spurious correlation for out-of-distribution detection. In *AAAI*, pages 10051–10059. AAAI Press, 2022.
- [Nam *et al.*, 2020] Jun Hyun Nam, Hyuntak Cha, Sungsoo Ahn, Jaeho Lee, and Jinwoo Shin. Learning from failure: Debiasing classifier from biased classifier. In *NeurIPS*, 2020.
- [Ortego *et al.*, 2021] Diego Ortego, Eric Arazo, Paul Albert, Noel E. O’Connor, and Kevin McGuinness. Multi-objective interpolation training for robustness to label noise. In *CVPR*, pages 6606–6615. IEEE, 2021.
- [Patrini *et al.*, 2017] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *CVPR*, pages 2233–2241. IEEE Computer Society, 2017.
- [Shalev-Shwartz and Ben-David, 2014] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning - From Theory to Algorithms*. Cambridge University Press, 2014.
- [Shu *et al.*, 2019] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. In *NeurIPS*, pages 1917–1928, 2019.
- [Simonyan and Zisserman, 2015] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [Sohn *et al.*, 2020] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fix-match: Simplifying semi-supervised learning with consistency and confidence. In *NeurIPS*, 2020.
- [Song *et al.*, 2019] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. SELFIE: refurbishing unclean samples for robust deep learning. In *ICML*, volume 97, pages 5907–5915. PMLR, 2019.
- [Tanaka *et al.*, 2018] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *CVPR*, pages 5552–5560. IEEE, 2018.
- [Wang *et al.*, 2019] Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *ICCV*, pages 322–330. IEEE, 2019.
- [Wang *et al.*, 2022] Xudong Wang, Zhirong Wu, Long Lian, and Stella X. Yu. Debaised learning from naturally imbalanced pseudo-labels. In *CVPR*, pages 14627–14637. IEEE, 2022.
- [Wei *et al.*, 2020] Hongxin Wei, Lei Feng, Xiangyu Chen, and Bo An. Combating noisy labels by agreement: A joint training method with co-regularization. In *CVPR*, pages 13723–13732. IEEE, 2020.
- [Wei *et al.*, 2022a] Hongxin Wei, Huiping Zhuang, Renchunzi Xie, Lei Feng, Gang Niu, Bo An, and Yixuan Li. Logit clipping for robust learning against label noise. *CoRR*, abs/2212.04055, 2022.
- [Wei *et al.*, 2022b] Jiaheng Wei, Zhaowei Zhu, Hao Cheng, Tongliang Liu, Gang Niu, and Yang Liu. Learning with noisy labels revisited: A study using real-world human annotations. In *ICLR*. OpenReview.net, 2022.
- [Xia *et al.*, 2019] Xiaobo Xia, Tongliang Liu, Nannan Wang, Bo Han, Chen Gong, Gang Niu, and Masashi Sugiyama. Are anchor points really indispensable in label-noise learning? In *NeurIPS*, pages 6835–6846, 2019.
- [Xiao *et al.*, 2015] Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *CVPR*, pages 2691–2699. IEEE Computer Society, 2015.
- [Xu *et al.*, 2019] Yilun Xu, Peng Cao, Yuqing Kong, and Yizhou Wang. L_{dmi}: A novel information-theoretic loss function for training deep nets robust to label noise. In *NeurIPS*, pages 6222–6233, 2019.
- [Yan *et al.*, 2014] Yan Yan, Rómer Rosales, Glenn Fung, Subramanian Ramanathan, and Jennifer G. Dy. Learning from multiple annotators with varying expertise. *Mach. Learn.*, 95(3):291–327, 2014.
- [Yao *et al.*, 2021] Yazhou Yao, Zeren Sun, Chuanyi Zhang, Fumin Shen, Qi Wu, Jian Zhang, and Zhenmin Tang. Jorsrc: A contrastive approach for combating noisy labels. In *CVPR*, pages 5192–5201, 2021.
- [Yi and Wu, 2019] Kun Yi and Jianxin Wu. Probabilistic end-to-end noise correction for learning with noisy labels. In *CVPR*, pages 7017–7025. IEEE, 2019.
- [Yu *et al.*, 2019] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor W. Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *ICML*, volume 97, pages 7164–7173. PMLR, 2019.
- [Zhang *et al.*, 2017] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR*. OpenReview.net, 2017.
- [Zhang *et al.*, 2021a] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep long-tailed learning: A survey. *CoRR*, abs/2110.04596, 2021.
- [Zhang *et al.*, 2021b] Yikai Zhang, Songzhu Zheng, Pengxiang Wu, Mayank Goswami, and Chao Chen. Learning with feature-dependent label noise: A progressive approach. In *ICLR*. OpenReview.net, 2021.
- [Zheng *et al.*, 2021] Guoqing Zheng, Ahmed Hassan Awadallah, and Susan T. Dumais. Meta label correction for noisy label learning. In *AAAI*, pages 11053–11061. AAAI Press, 2021.
- [Zhu *et al.*, 2022] Zhaowei Zhu, Jialu Wang, and Yang Liu. Beyond images: Label noise transition matrix estimation for tasks with lower-quality features. *CoRR*, abs/2202.01273, 2022.