

# Foster Adaptivity and Balance in Learning with Noisy Labels

Mengmeng Sheng<sup>1</sup>, Zeren Sun<sup>1</sup>(✉), Tao Chen<sup>1</sup>, Shuchao Pang<sup>1</sup>, Yucheng Wang<sup>2</sup>, and Yazhou Yao<sup>1</sup>(✉)

<sup>1</sup> Nanjing University of Science and Technology, Nanjing, China

{shengmengmeng, zeren, taochen, pangshuchao, yazhou.yao}@njust.edu.cn

<sup>2</sup> Horizon Robotics, Beijing, China


yucheng.wang@horizon.cc

**Abstract.** Label noise is ubiquitous in real-world scenarios, posing a practical challenge to supervised models due to its effect in hurting the generalization performance of deep neural networks. Existing methods primarily employ the sample selection paradigm and usually rely on dataset-dependent prior knowledge (*e.g.*, a pre-defined threshold) to cope with label noise, inevitably degrading the adaptivity. Moreover, existing methods tend to neglect the class balance in selecting samples, leading to biased model performance. To this end, we propose a simple yet effective approach named **SED** to deal with label noise in a **Self**-adaptive**E** and class-balance**D** manner. Specifically, we first design a novel sample selection strategy to empower self-adaptivity and class balance when identifying clean and noisy data. A mean-teacher model is then employed to correct labels of noisy samples. Subsequently, we propose a self-adaptive and class-balanced sample re-weighting mechanism to assign different weights to detected noisy samples. Finally, we additionally employ consistency regularization on selected clean samples to improve model generalization performance. Extensive experimental results on synthetic and real-world datasets demonstrate the effectiveness and superiority of our proposed method. The open-source project has been made available at <https://nust-machine-intelligence-laboratory.github.io/project-SED>.

**Keywords:** Noisy labels · Self-adaptive · Class-balanced · Sample selection and re-weighting

## 1 Introduction

Deep neural networks (DNNs) have witnessed remarkable achievements in many computer vision tasks, such as image classification [24, 39], object detection [42, 44], face recognition [5], and instance segmentation [8, 9]. The superior performance of DNNs is highly attributed to supervised training with large-scale and high-quality human-labeled training datasets (*e.g.*, ImageNet [12]). However, collecting large-scale datasets with accurate annotations is expensive and time-consuming, especially for tasks requiring expert annotation knowledge (*e.g.*,



**Fig. 1:** (a-b) Self-adaptive and class-balanced sample selection based on predicted probability w.r.t. given labels. The blue curve indicates the class-specific selection thresholds. (c-d) Self-adaptive and class-balanced sample re-weighting based on correction confidence. The orange curve represents the class-specific confidence threshold.

medical images [61]). To alleviate this problem, researchers start to resort to alternative methods, such as crowd-sourcing platforms [60] or web image search engines [14], for obtaining cheaper label annotations. Unfortunately, these methods usually result in unavoidable noisy labels, which tend to cause inferior model performance due to the strong learning ability of DNNs [70]. Consequently, developing robust models for learning with noisy labels is of significant importance.

Recently, a growing number of methods have been proposed for addressing the label noise problem [2, 4, 6, 17, 30, 54, 59, 62, 65]. Label correction and sample selection/re-weighting are two major strategies for tackling noisy labels. Label correction methods typically attempt to rectify labels using the noise transition matrix [15] or model predictions [29]. For example, [40] proposes to correct corrupted labels by estimating the noise transition matrix. Jo-SRC [66] uses the temporally averaged model (*i.e.*, mean-teacher model) to generate reliable pseudo-label distributions for providing supervision. However, on the one hand, the noise transition matrix is hard to estimate in real-world scenarios. On the other hand, networks tend to have better recognition capability on simple categories than hard ones. This recognition bias usually results in imbalanced label corrections (*i.e.*, samples are more likely to be corrected into simple categories) in prediction-based label correction methods, hurting the final model performance.

Another line of research focuses on the sample selection/re-weighting [19, 23, 32, 45, 49, 52, 66, 67]. Sample selection methods primarily seek to split samples into two subsets: a noisy subset and a clean subset [18, 19, 66]. Prior methods tend to regard samples with small losses as clean ones [19, 58]. For example, JoCoR [58] exploits a joint loss to select small-loss samples to encourage agreement between models. However, these methods often require proper prior knowledge (*e.g.*, a pre-defined drop rate or threshold) to achieve effective sample selection. Moreover, previous literature usually neglects class balance during sample selection, leading to biased model performance. Sample re-weighting can be deemed as a variant of sample selection, smoothing its 0/1 weighting scheme to a softer one. Samples with higher confidence are assigned larger weights, while those with lower confidence are assigned smaller weights. For example, L2RW [43] proposes to assign different sample weights based on meta-learning. However, existing

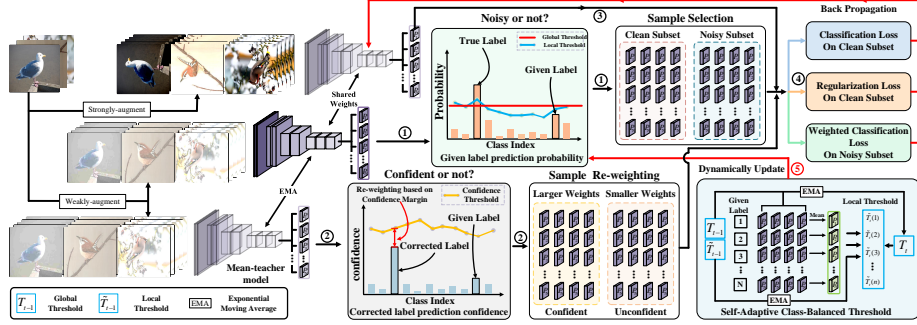
sample re-weighting methods also tend to require prior knowledge (*e.g.*, a small subset of clean samples).

To alleviate the aforementioned issues, we propose a simple yet effective method, named **SED**, to learn with noisy labels in a **Self-adaptive** and **class-balanceD** manner. Our SED integrates sample selection, label correction, and sample re-weighting. Specifically, we propose to identify clean samples based on the predicted probability w.r.t. the given labels of input samples. To promote self-adaptivity and class balance in sample selection, we propose to integrate global and local thresholds for each category when distinguishing between clean and noisy data (as shown in Fig. 1 (a) and (b)). The global and local thresholds are dynamically updated during training. Once the clean and noisy subsets are obtained, we employ a mean-teacher model to correct labels for identified noisy samples. Subsequently, we propose to re-weight label-corrected noisy samples in a self-adaptive and class-balanced fashion to alleviate the confirmation bias caused by imbalanced label correction. We impose larger/smaller weights on noisy samples with higher/lower correction confidence according to an estimated truncated normal distribution (as shown in Fig. 1 (c) and (d)). Finally, we employ an additional regularization loss term on identified clean samples to further enhance the performance and robustness of the model. Comprehensive experimental results have been provided to verify the effectiveness and superiority of our proposed SED on synthetically corrupted datasets and real-world datasets. Our contributions are summarized as follows:

- (1) We propose a simple yet effective method, named SED, to combat noisy labels. SED selects and re-weights samples in a self-adaptive and class-balanced manner, alleviating the demand for dataset-dependent prior knowledge and the negative effect caused by class imbalance.
- (2) Our proposed SED selects samples according to class-specific thresholds that are estimated in a data-driven manner, encouraging self-adaptivity and class balance in sample selection. In addition, we propose to re-weight samples based on a truncated normal distribution that is updated periodically, mitigating performance downgrade due to imbalanced label corrections.
- (3) We provide comprehensive experimental results on synthetic and real-world datasets to illustrate the superiority of our proposed SED. Extensive ablation studies are conducted to further verify the effectiveness of our method.

## 2 Related Work

**Label Correction.** The intuitive idea for handling noisy labels is to correct corrupted labels before feeding them into networks [11, 15, 16, 33, 34, 40, 57, 64, 67]. Early works propose to correct the training labels by estimating the noise transition matrix. [67] introduces an intermediate class to avoid directly estimating the noisy class posterior and then factorizes the transition matrix into the product of two sub-matrices. However, the transition matrix is hard to estimate accurately in real-world scenarios. Some other methods propose to model label noise by using predictions of DNNs [27, 55, 56, 68]. [68] proposes to directly learn label



**Fig. 2:** The overall framework of our SED. We first divide the training set into a clean subset and a noisy subset based on global and local thresholds that are dynamically updated. Our threshold design enables self-adaptivity and class balance in sample selection. We then employ a mean-teacher model to correct labels for noisy samples. Based on the correction confidence, SED adaptively assigns different weights to label-corrected noisy samples and uses them for training. Finally, SED further boosts the model performance by imposing an additional consistency regularization loss on selected clean samples. The final objective loss integrates the classification losses on clean and noisy samples and the regularization loss on clean samples.

distributions for corrupted samples in an end-to-end manner. Nevertheless, since DNNs tend to learn better on simple categories than hard ones, pseudo-labels are more likely to fall into the simple class set, leading to imbalanced label correction. In this work, we resort to the re-weighting strategy to alleviate the issue caused by imbalanced label correction.

**Sample Selection.** Another type of classical method to deal with label noise is sample selection, which divides the training set into a clean subset and a noisy subset [19, 51, 58, 66]. Previous sample selection methods primarily employ the cross-entropy loss as the selection criterion, regarding samples with small losses as clean ones. For example, Co-teaching [19] proposes to cross-update two networks using small-loss samples selected by peer networks. Some recent methods propose new selection criteria for finding clean samples [29, 66]. Jo-SRC [66] proposes to employ Jensen-Shannon Divergence for selecting clean samples globally. DISC [31] proposes to select reliable instances based on the insight of memorization strength. However, these methods usually demand pre-defined drop rates or thresholds. Furthermore, previous methods neglect the class imbalance issue in the selection process, leading to inferior and biased model performance. In this work, we employ predicted probability as the selection criterion and propose a novel threshold mechanism to enable self-adaptive and class-balanced selection.

**Sample Re-weighting.** Recently, some researchers have been devoted to re-weighting training samples to cope with noisy labels [13, 47, 53, 63]. These methods usually assign larger weights to samples that are more likely to be clean while smaller weights to others, minimizing the misleading impact of noisy sam-

ples. For example, L2RW [43] proposes a meta-learning algorithm that learns to assign weights to training examples based on their gradient directions. However, existing methods tend to require considerable prior knowledge (*e.g.*, a small subset of clean samples), posing a limit to their practicability. In this work, we design a novel re-weighting scheme to empower self-adaptivity and class balance when leveraging label-corrected noisy samples.

### 3 Method

#### 3.1 Problem Statement

Formally, considering a  $C$ -class classification problem, we denote  $D_{train} = \{(x_i, y_i) | i = 1, \dots, N\}$  as the training set with label noise, in which  $x_i$  denotes the  $i$ -th training sample and  $y_i \in \{0, 1\}^C$  is its associated label (potentially “incorrect”). We use  $y_i^*$  to represent the ground-truth label of  $x_i$  and denote  $D_{test} = \{(x_i, y_i^*) | i = 1, \dots, M\}$  as the test set with accurate labels.  $N$  and  $M$  represent the total number of samples in the training set  $D_{train}$  and test set  $D_{test}$ , respectively. The goal is to train a robust classification neural network  $\mathcal{F}(\cdot, \theta)$  ( $\theta$  denotes network parameters) on the noisy training set  $D_{train}$  to perform accurate prediction on the test set  $D_{test}$ . The conventional classification task usually hypothesizes that given labels of training samples are accurate (*i.e.*,  $y_i = y_i^*$ ), thus using the following cross-entropy loss to optimize the network.

$$\mathcal{L}_{ce} = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C y_i^c \log(p^c(x_i, \theta)), \quad (1)$$

in which  $p^c(x_i, \theta)$  denotes the predicted softmax probability of the  $i$ -th training sample  $x_i$  over its  $c$ -th class.

Due to the memorization effect [1] (*i.e.*, models tend to fit clean and simple samples first and then gradually memorize noisy ones), the network optimization based on the above loss usually leads to an ill-suited solution. One potentially useful remedy is to integrate sample selection, label correction, and sample re-weighting. In this work, we follow this paradigm to combat noisy labels by encouraging self-adaptivity and class balance.

#### 3.2 Adaptive and Balanced Sample Selection

Previous studies [19, 58, 66] usually select small-loss samples as clean ones based on pre-defined drop rates or thresholds. It should be noted that drop rates can be easily converted to thresholds during selection, thus we only discuss thresholds hereafter. The selection thresholds are usually dataset-dependent, making it challenging to adapt them to different real-world datasets. Although existing methods employ scheduling strategies (*e.g.*, a gradually increasing schedule [19]) to adjust thresholds during training for fully exploiting the model capability, these scheduling designs are rather heuristic and still require pre-defined initial

and final threshold values. Moreover, few works consider the different difficulties in learning various categories, leading to biased selection results and inferior model performance.

To this end, we propose a self-adaptive and class-balanced sample selection (SCS) strategy to address the above problems. SCS adaptively adjusts the threshold in an epoch-wise and class-wise manner to enable effective clean sample identification. Specifically, we employ global and local thresholds, which are both self-adaptive, to distinguish between clean and noisy samples in each category. Since the cross-entropy loss is unbounded, we propose to rely on predicted probability w.r.t. the given labels  $p^{y_i}(x_i, \theta)$  to determine whether the samples are clean. Samples with higher  $p^{y_i}(x_i, \theta)$  are more likely to have correct labels.

We estimate the global threshold based on the averaged predicted probability w.r.t. given labels over all training samples to reflect the overall learning state of the network. This design makes the global threshold data-driven, thus eliminating the demand for pre-defined thresholds. Moreover, we employ the exponential moving average (EMA) to further refine the global threshold, alleviating unstable training caused by large perturbation of the averaged predicted probability. By adopting an initial value of  $T_0 = \frac{1}{C}$ , our final global threshold at the  $t$ -th epoch is defined as:

$$T_t = \begin{cases} \frac{1}{C}, & t = 0 \\ mT_{t-1} + (1 - m)\frac{1}{N} \sum_{i=1}^N p^{y_i}(x_i, \theta), & t > 0 \end{cases} \quad (2)$$

Our design of the global threshold scheduling implicitly complies with the memorization effect [1]. As the training progresses, the predicted probability w.r.t. the given label gradually increases, leading to the monotonic increase of  $T_t$ . Consequently, the network can learn from more samples in the early stage but fewer samples in the later stage.

As stated above, using only a global threshold to divide the training set neglects the difference among various categories and will result in imbalanced sample selection (*i.e.*, fewer samples of complicated categories will be selected as clean data). Samples of easy categories tend to be better learned and have higher  $p^{y_i}(x_i, \theta)$ , thus requiring larger thresholds to distinguish between clean and noisy data. Therefore, we additionally propose a local threshold scheme to further adjust the global threshold. We first estimate the expectation of the model's predictions  $\tilde{E}_t(c)$  on each class  $c$  at the  $t$ -th epoch to reveal the class-specific learning status.

$$\tilde{E}_t(c) = \begin{cases} \frac{1}{C}, & t = 0 \\ m\tilde{E}_{t-1}(c) + (1 - m)\frac{1}{N} \sum_{i=1}^N p^c(x_i, \theta), & t > 0 \end{cases} \quad (3)$$

Accordingly, we obtain local threshold  $\tilde{T}_t(c)$  for each class  $c$  by normalizing  $\tilde{E}_t(c)$  and integrating it with global threshold  $T_t$  as:

$$\tilde{T}_t(c) = \frac{\tilde{E}_t(c)}{\max\{\tilde{E}_t(c) : c \in [C]\}} T_t. \quad (4)$$

On the one hand, the design of our global threshold ensures that sufficient clean samples are identified and learned by the network. On the other hand, the design of our local threshold ensures that selected clean samples are class-balanced. Finally, by unifying our proposed global and local thresholds, we divide the training set  $D_{train}$  into a clean subset  $D_c$  and a noisy subset  $D_n$  in each epoch according to Eq. (5).

$$\begin{cases} D_c = \{(x_i, y_i) | (x_i, y_i) \in D_{train}, p^{y_i}(x_i, \theta) > \tilde{T}_t(y_i)\} \\ D_n = \{(x_i, y_i) | (x_i, y_i) \in D_{train}, (x_i, y_i) \notin D_c\} \end{cases}. \quad (5)$$

### 3.3 Adaptive and Balanced Re-weighting

Recent researches propose to cope with noisy samples in a semi-supervised-learning-like (SSL-like) manner by integrating sample selection and label correction [29, 66]. Identified clean samples are used conventionally for model training, while detected noisy samples are assigned pseudo labels to correct their supervision before being used for training. However, existing methods tend to treat label-corrected noisy samples equally, neglecting their difference in reliability. Moreover, due to different learning difficulties in various categories, label correction results may be imbalanced (noisy samples are more likely to be assigned labels of simple classes), resulting in biased label correction and sub-optimal model performance.

To mitigate the above issue, we propose a self-adaptive and class-balanced re-weighting (SCR) mechanism to adaptively assign different weights to samples according to their confidence. Specifically, we use a temporally averaged model (*i.e.*, mean-teacher model  $\theta^*$ ) to generate reliable pseudo labels for detected noisy samples. By introducing the historical models, we obtain corrected labels  $y^{corr}$  using  $\theta^*$  to promote the reliability of label correction and alleviate error-propagation issues. The mean-teacher model  $\theta^*$  is not updated in the gradient back-propagation.  $\theta^*$  is updated in each training step  $t'$  as follows:

$$\theta_{t'}^* = \alpha \theta_{t'-1}^* + (1 - \alpha) \theta_{t'}, \quad (6)$$

in which  $\theta_0^*$  is initialized using the initial model parameters of  $\theta$ . Accordingly, noisy samples are assigned pseudo labels as follows:

$$y_i^{corr} = \arg \max_{j=1, \dots, C} p^j(x_i, \theta^*). \quad (7)$$

As mentioned above, the label correction results could be imbalanced due to the biased capability of the network. Consequently, we propose a re-weighting method to adaptively assign larger weights to (noisy) samples with higher correction confidence. We employ the prediction probability w.r.t. the corrected label to reveal the correction confidence. Inspired by the semi-supervised learning methods [3, 7, 10, 48], we propose to fit the underlying sample weights to a dynamic truncated normal distribution, whose mean and variance values at

the  $t$ -th epoch are  $\mu_t$  and  $\sigma_t$ . The sample weights are therefore derived in a self-adaptive fashion as:

$$\lambda(x_i) = \begin{cases} \lambda_m \exp\left(\frac{(p^{y_i^{corr}}(x_i, \theta) - \mu_t)^2}{-2\sigma_t^2}\right), & p^{y_i^{corr}}(x_i, \theta) < \mu_t, \\ \lambda_m, & otherwise \end{cases}, \quad (8)$$

in which  $\lambda_m$  is the upper bound of sample weights. Assuming sample weights to follow the dynamic truncated normal distribution is equivalent to treating the deviation of correction confidence from  $\mu_t$  as a proxy measure of the correctness of the label correction. Samples with higher confidence are less prone to be erroneously label-corrected than those with lower confidence, thus being assigned larger weights.

Moreover, to enable class-balanced re-weighting and promote training stability, we propose to estimate  $\mu_t(c)$  and  $\sigma_t^2(c)$  for each class  $c$  based on their historical estimations using EMA:

$$\mu_t(c) = \begin{cases} \frac{1}{C}, & t = 0 \\ m\mu_{t-1}(c) + (1-m)\tilde{\mu}(c), & t > 0 \end{cases}, \quad (9)$$

$$\sigma_t^2(c) = \begin{cases} 1.0, & t = 0 \\ m\sigma_{t-1}^2(c) + (1-m)\tilde{\sigma}^2(c), & t > 0 \end{cases}, \quad (10)$$

in which,

$$\tilde{\mu}(c) = \frac{1}{|D_n|} \sum_{i=1}^{|D_n|} p^{y_i^{corr}}(x_i, \theta), \quad \text{if } y_i^{corr} = c, \quad (11)$$

$$\tilde{\sigma}^2(c) = \frac{1}{|D_n|} \sum_{i=1}^{|D_n|} (p^{y_i^{corr}}(x_i, \theta) - \tilde{\mu}(c))^2, \quad \text{if } y_i^{corr} = c. \quad (12)$$

$\mu_t$  and  $\sigma_t$  of the dynamic truncated normal distribution can be adaptively estimated from the correction confidence distribution based on Eqs. (9) and (10). As the model performance improves during training,  $\mu_t$  gradually increases and  $\sigma_t$  decreases. Since the tail of the normal distribution grows exponentially tighter, the samples with lower correction confidence are given lower weights. Besides, we estimate class-specific  $\mu_t$  and  $\sigma_t$ . This effectively alleviates the class imbalance in the label correction process caused by the biased model ability.

### 3.4 Overall Framework

In summary, our proposed SED follows the paradigm that integrates sample selection, label correction, and sample re-weighting for addressing noisy labels. Details of our SED are shown in Fig. 2 and Algorithm 1.

Firstly, SED divides  $D_{train}$  into a clean subset  $D_c$  and a noisy subset  $D_n$  in a self-adaptive and class-balanced manner. For samples in the clean subset  $D_c$ , we take their given labels to calculate the classification loss  $\mathcal{L}_{D_c}$  as follow:

$$\mathcal{L}_{D_c} = -\frac{1}{|D_c|} \sum_{(x,y) \in D_c} y \log p(x, \theta). \quad (13)$$



---

**Algorithm 1** Our proposed SED algorithm

---

**Input:** The training set  $D_{train}$ , network  $\theta$ , mean-teacher network  $\theta^*$ , total epochs  $E_{total}$ , batch size  $bs$ .

```

1: for  $epoch = 1, 2, \dots, E_{total}$  do
2:   Obtain  $T_t$  and  $\tilde{T}_t$  by Eqs. (2), (3) and (4)
3:   Obtain  $D_c$  and  $D_n$  based on Eq. (5).
4:   Obtain  $y^{corr}$ ,  $\tilde{\mu}$ , and  $\tilde{\sigma}^2$  by Eqs. (7), (9) and (10) .
5:   Obtain  $\lambda(x)$  based on Eq. (8).
6:   for  $iteration = 1, 2, \dots$  do
7:     Fetch  $B = \{(x_i, y_i)\}^{bs}$  from  $D_{train}$ 
8:     Obtain  $B_{clean} \subseteq D_c$  and  $B_{noise} \subseteq D_n$ 
9:     Calculate  $\mathcal{L} = \mathcal{L}_{D_c} + \mathcal{L}_{D_n} + \mathcal{L}_{reg}$ 
10:    Update  $\theta$  by optimizing  $\mathcal{L}$ 
11:    Update  $\theta^*$  by Eq. (6)
12:  end for
13: end for

```

**Output:** Updated network  $\theta$ .

---

For samples in the noisy subset  $D_n$ , we discard their given labels and perform label correction based on a mean-teacher model using Eq. (7). Then, we calculate the loss of the noisy subset  $\mathcal{L}_{D_n}$  as

$$\mathcal{L}_{D_n} = -\frac{1}{|D_n|} \sum_{(x,y) \in D_n} \lambda(x) y^{corr} \log p(\hat{x}, \theta), \quad (14)$$

in which  $\hat{x}$  denotes the strongly-augmented view of the sample  $x$ .  $\lambda(x)$  represents the sample weight computed by Eq. (8). Finally, we incorporate an additional weighted classification loss on clean samples w.r.t. corrected labels (similar to  $\mathcal{L}_{D_n}$ ) to further enhance the robustness of the model. This loss term implicitly encourages prediction consistency between weakly- and strongly-augmented views of samples from the clean subset, regularizing the model to achieve better performance. Thus, we term this loss as the consistency regularization loss and compute it as follows:

$$\mathcal{L}_{reg} = -\frac{1}{|D_c|} \sum_{(x,y) \in D_c} \lambda(x) y^{corr} \log p(\hat{x}, \theta), \quad (15)$$

where  $\lambda(x)$  is also computed based on Eq. (8). Accordingly, the final objective loss function in our SED is:

$$\mathcal{L} = \mathcal{L}_{D_c} + \mathcal{L}_{D_n} + \mathcal{L}_{reg}. \quad (16)$$

**Table 1:** Average test accuracy (%) on CIFAR100N and CIFAR80N over the last ten epochs. Experiments are conducted under various noise conditions (“Sym” and “Asym” denote the symmetric and asymmetric label noise, respectively). Results of existing methods are mainly drawn from [50]. <sup>†</sup> means that we re-implement the method using its open-sourced code and default hyper-parameters.

Methods	Publication	CIFAR100N			CIFAR80N		
		Sym-20%	Sym-80%	Asym-40%	Sym-20%	Sym-80%	Asym-40%
Standard	-	35.14	4.41	27.29	29.37	4.20	22.25
Decoupling [37]	NeurIPS 2017	33.10	3.89	26.11	43.49	10.1	33.74
Co-teaching [19]	NeurIPS 2018	43.73	15.15	28.35	60.38	16.59	42.42
Co-teaching+ [69]	ICML 2019	49.27	13.44	33.62	53.97	12.29	43.01
JoCoR [58]	CVPR 2020	53.01	15.49	32.70	59.99	12.85	39.37
Jo-SRC [66]	CVPR 2021	58.15	23.80	38.52	65.83	29.76	53.03
SELC [36]	IJCAI 2022	55.44	23.54	45.19	57.51	22.79	47.50
DivideMix [29]	ICLR 2020	57.76	28.98	43.75	57.47	21.18	37.47
Co-LDL [50]	TMM 2022	59.73	25.12	52.28	58.81	24.22	50.69
UNICON <sup>†</sup> [25]	CVPR 2022	55.10	31.49	49.90	54.50	36.75	51.50
NCE <sup>†</sup> [28]	ECCV 2022	54.58	35.23	49.90	58.53	39.34	56.40
SOP <sup>†</sup> [35]	ICML 2022	58.63	34.23	49.87	60.17	34.05	53.34
SPRL <sup>†</sup> [46]	PR 2023	57.04	28.61	49.38	47.90	22.25	40.86
AGCE <sup>†</sup> [71]	TPAMI 2023	59.38	27.41	43.04	60.24	25.39	44.06
DISC <sup>†</sup> [31]	CVPR 2023	60.28	33.90	50.56	50.33	38.23	47.63
<b>Ours</b>	-	<b>66.50</b>	<b>38.15</b>	<b>58.29</b>	<b>69.10</b>	<b>42.57</b>	<b>60.87</b>

## 4 Experiments


In this section, we conduct experiments on two synthetically corrupted datasets (*i.e.*, CIFAR100N and CIFAR80N [66]) and three real-world datasets (*i.e.*, Web-Aircraft, Web-Car, and Web-Bird [49]). We demonstrate the superiority of our method in coping with noisy labels by comparing SED with various state-of-the-art (SOTA) methods. Moreover, we conduct extensive ablation studies to evaluate the effectiveness of each component in our SED.

### 4.1 Experiment Setup

**Synthetically Corrupted Datasets.** CIFAR100N and CIFAR80N are mainly derived from CIFAR100 [26]. CIFAR100 consists of 60,000 RGB images (50,000 for training and 10,000 for testing). We follow [66] to create the closed-set noisy dataset CIFAR100N and the open-set noisy dataset CIFAR80N. In particular, to construct the open-set noisy dataset CIFAR80N, we regard the last 20 categories in CIFAR100 as out-of-distribution samples. We adopt two classical noise structures: symmetric and asymmetric, with a noise ratio  $n \in (0, 1)$ .

**Real-World Datasets.** To further verify the effectiveness of our SED in practical scenarios, we conduct experiments on the three real-world noisy datasets (*i.e.*, Web-Aircraft, Web-Car, and Web-Bird [49]), whose training images are crawled from web image search engines. The noise rates and structures of real-world datasets are all unknown. No label verification information is provided.

**Implementation Details.** On synthetically corrupted datasets, we follow [66] to conduct experiments with a seven-layer CNN network as the backbone. The



**Fig. 3:** Comparison of different sample selection methods and the ablation results of the parameter  $\lambda_m$ . (a) The overall precision of sample selection (%) *vs.* epochs. (b) The class-wise precision of sample selection (%) *vs.* classes. (c) The test accuracy (%) *vs.* epochs. (d) The test accuracy (%) of using different  $\lambda_m$ .

network is trained using SGD with a momentum of 0.9 for 100 epochs (including 20 warm-up epochs). The batch size is 128, and the initial learning rate is 0.05. For real-world datasets, we follow [50] and leverage ResNet50 [20] pre-trained on ImageNet as our backbone. We use the SGD optimizer with a momentum of 0.9 to train the network for 110 epochs. The batch size, the initial learning rate, and the weight decay are 32, 0.005, and 0.0005. The learning rate decays in a cosine annealing manner. We train the network for 110 epochs, in which the first 10 epochs are warm-up. The EMA coefficients  $m$  and  $\alpha$  are set to 0.99 and 0.95.  $\lambda_m$  is set to 1.0 for all datasets.

**Baselines.** For CIFAR100N and CIFAR80N, we compare our method with the following SOTA methods: Decoupling [37], Co-teaching [19], Co-teaching+ [69], JoCoR [58], Jo-SRC [66], SELC [36], Co-LDL [50], UNICON [25], SOP [35], AGCE [71], and DISC [31]. For Web-Aircraft, Web-Bird, and Web-Car, besides the above methods, we additionally compare SED with other SOTA methods (*e.g.*, PENCIL [68], Hendrycks *et al.* [21], mCT-S2R [38], AFM [41], and Self-adaptive [22]). Moreover, we perform conventional training using the entire noisy dataset. The result is provided as a baseline (denoted as Standard). Results in Tables 1 and 2 are mainly obtained from [66] and [50].

## 4.2 Evaluation on Synthetic Datasets

We show the comparison results between our SED and existing SOTA methods on the synthetic datasets (*i.e.*, CIFAR100N and CIFAR80N) in Table 1.

**Results on CIFAR100N.** Table 1 shows that SED consistently achieves the best performance compared to SOTA methods on CIFAR100N. In particular, it should be noted that SED can better adapt to severely noisy situations (*i.e.*, Sym-80%), while most SOTA approaches almost fail in the most inferior case. It should be emphasized that the asymmetric noise case is often more challenging than the symmetric one. Our SED shows a significant improvement (*i.e.*,  $\geq 6.01\%$ ) on Asym-40%. Experiments on CIFAR100N show that SED can effectively deal with closed-set label noise in different noise situations.

**Results on CIFAR80N.** To simulate real-world scenarios, CIFAR80N contains both closed-set and open-set noisy labels, making it undoubtedly more challeng-

**Table 2:** The comparison with SOTA approaches in test accuracy (%) on real-world noisy datasets: Web-Aircraft, Web-Bird, Web-Car. Results of existing methods are mainly drawn from [50]. <sup>†</sup> means that we re-implement the method using its open-sourced code and default hyper-parameters.

Methods	Publication	Backbone	Performances(%)			
			Web-Aircraft	Web-Bird	Web-Car	Average
Standard	-	ResNet50	60.80	64.40	60.60	61.93
Decoupling [37]	NeurIPS 2017	ResNet50	75.91	71.61	79.41	75.64
Co-teaching [19]	NeurIPS 2018	ResNet50	79.54	76.68	84.95	80.39
Co-teaching+ [69]	ICML 2019	ResNet50	74.80	70.12	76.77	73.90
PENCIL [68]	CVPR 2019	ResNet50	78.82	75.09	81.68	78.53
Hendrycks <i>et al.</i> [21]	NeurIPS 2019	ResNet50	73.24	70.03	73.81	72.36
mCT-S2R [38]	WACV 2020	ResNet50	79.33	77.67	82.92	79.97
JoCoR [58]	CVPR 2020	ResNet50	80.11	79.19	85.10	81.47
AFM [41]	ECCV 2020	ResNet50	81.04	76.35	83.48	80.29
DivideMix [29]	ICLR 2020	ResNet50	82.48	74.40	84.27	80.38
Self-adaptive [22]	NeurIPS 2020	ResNet50	77.92	78.49	78.19	78.20
Co-LDL [50]	TMM 2022	ResNet50	81.97	80.11	86.95	83.01
UNICON <sup>†</sup> [25]	CVPR 2022	ResNet50	85.18	81.20	88.15	84.84
NCE <sup>†</sup> [28]	ECCV 2022	ResNet50	84.94	80.22	86.38	83.85
SOP <sup>†</sup> [35]	ICML 2022	ResNet50	84.06	79.40	85.71	83.06
SPRL <sup>†</sup> [46]	PR 2023	ResNet50	84.40	76.36	86.84	82.53
AGCE <sup>†</sup> [71]	TPAMI 2023	ResNet50	84.22	75.60	85.16	81.66
DISC <sup>†</sup> [31]	CVPR 2023	ResNet50	85.27	81.08	88.31	84.89
<b>Ours</b>	-	ResNet50	<b>86.62</b>	<b>82.00</b>	<b>88.88</b>	<b>85.83</b>

ing. Results shown in Table 1 illustrate: (1) in the case of Sym-20%, our SED can achieve a 3.27% performance improvement. (2) in the case of Sym-80%, while most SOTA approaches fail to tackle the massive noisy labels, SED achieves the best result. (3) when the noise scenario becomes harder (*i.e.*, Asym-40%), our SED consistently obtains the best performance, outperforming the second-best result by 7.53%. Table 1 proves that SED performs consistently better than existing methods when coping with open-set noisy datasets.

### 4.3 Evaluation on Real-world Datasets

Table 2 shows the experimental results of existing methods and SED on Web-Aircraft, Web-Bird, and Web-Car, which contain open-set and closed-set noise simultaneously. From this table, we can find that SED can achieve better (or comparable) performance against SOTA approaches in different datasets. SED achieves performances of 86.62%, 82.00%, and 88.88% on test sets of Web-Aircraft, Web-Bird, and Web-Car, respectively. The average test accuracy outperforms existing SOTA methods by 0.94%. It should be noted that the second and third-best methods (*i.e.*, DISC and UNICON) involve the Mixup training trick and two simultaneously trained networks respectively, while SED trains only one network without Mixup. Compared to existing methods, our SED eliminates the demand for dataset-dependent prior knowledge (*e.g.*, pre-defined drop rate/threshold), making it easier to adapt to different datasets.

**Table 3:** Effect of each component in the test accuracy (%) on CIFAR100N.

Model	Test Accuracy
Standard	34.10
Standard+SCS w/o local threshold	53.36
Standard+SCS w/o global threshold	55.64
Standard+SCS w/o EMA	54.72
Standard+SCS	58.21
Standard+SCS+SCR w/o re-weighting	59.75
Standard+SCS+SCR w/o EMA	60.08
Standard+SCS+SCR	60.43
Standard+SCS+SCR+CR	62.65

#### 4.4 Ablation Studies

In this section, we demonstrate the effectiveness of each component in our SED (*i.e.*, SCS, SCR, and CR). Besides, we investigate the effect of the hyper-parameter  $\lambda_m$  in Eq. (8). Unless otherwise stated, ablation experiments are conducted on CIFAR100N (Sym-50%). Table 3 and Table 4 show the impact of each component.

**Effects of Self-adaptive and Class-balanced Sample Selection.** As analyzed above, existing sample selection methods tend to struggle with the demand for dataset-dependent prior knowledge, such as pre-defined drop rate/threshold. However, these hyper-parameters are usually unknown and hard to estimate in real-world datasets. The proposed SCS strategy in our method allows adaptive sample selection in a class-balanced manner, making our SED have better generalization performance on different datasets. As shown in Table 3, employing SCS achieves a 24.11% performance gain compared to the baseline Standard. We also provide the result of using SCS without local thresholds and global thresholds. This proves that our threshold design is crucial for improving the robustness of the model.

To further demonstrate the superiority of our SCS over previous sample selection strategies, we compare our SCS with two commonly-used methods (*i.e.*, small-loss [19], and GMM [29]) in Fig. 3. As shown in Fig. 3 (a), our SCS is shown to be more effective in selecting clean samples accurately compared with the other two strategies. Additionally, we compare the sample selection accuracy for each category in the selected clean subset and present the comparison in Fig. 3 (b). It illustrates that the selection results of SCS are more balanced. The curves of test accuracy are shown in Fig. 3 (c), revealing the leading performance of our SCS compared with the other two methods and the baseline.

**Effects of Self-adaptive and Class-balanced Sample Re-weighting.** Our SED follows an SSL-like paradigm. Selected clean samples are learned conventionally, while detected noisy samples are also fed into the network for training after label correction. However, the biased model capability tends to result in imbalanced label correction, hurting the model performance. We accordingly propose SCR to re-weight detected noisy samples in a self-adaptive and class-

**Table 4:** Effect of promoting class balance on CIFAR100N (left) and CIFAR80N (right). Test accuracy (%) of SED with and without the class-balanced design is compared under different settings.

Class-balanced?	✗	✓
Sym-20%	64.16	66.59
Sym-80%	38.08	39.32
Asym-40%	52.78	58.80

Class-balanced?	✗	✓
Sym-20%	67.20	68.75
Sym-80%	39.74	42.90
Asym-40%	57.00	61.51

balanced manner when using their corrected labels for training. Table 3 shows a performance gain of 2.19% by employing our proposed SCR. The only involved hyper-parameter in the SCR is the  $\lambda_m$  in Eq. (8). Fig. 3 (d) exhibits the influence of different  $\lambda_m$  values on the test accuracy when experimenting with CIFAR100N (Sym-50%) and CIFAR80N (Asym-40%). It can be observed that the best performance is achieved when  $\lambda_m = 1.0$  on CIFAR100N (Sym-50%) and CIFAR80N (Asym-40%).

**Effects of Consistency Regularization.** Although clean samples selected by SED are more accurate and balanced than previous methods, it is inevitable that some noisy data will be mistakenly selected into the clean subset. Therefore, we impose an additional CR on the selected clean samples to enhance the model’s robustness. Table 3 shows that CR successfully boosts model performance by 2.02%, revealing the benefits that CR brings to our model.

**Effects of Promoting Class Balance.** As stated in SCS and SCR, our SED favors the class-balanced design. Specifically, SCS estimates local thresholds on each class to avoid imbalanced sample selection, while SCR also estimates  $\mu_t$  and  $\sigma_t^2$  of the dynamic truncated normal distribution for each class to encourage balanced re-weighting. As shown in Table 4, we investigate the effect of the class-balanced design in SED. We can find that our method consistently achieves better performance when incorporated with the class-balanced design, especially in harder scenarios. Table 4 effectively demonstrates that the class-balanced design in our SED is beneficial for model performance.

## 5 Conclusion

In this paper, we proposed a simple yet effective approach named SED to address the inferior model performance caused by noisy labels. We designed a self-adaptive and class-balanced sample selection strategy to distinguish between clean and noisy samples. Clean samples were learned conventionally. A mean-teacher model was employed to correct the labels of detected noisy samples. Subsequently, SED re-weighted noisy samples in a self-adaptive and class-balanced fashion based on the correction confidence when leveraging them for model training. Finally, we additionally imposed consistency regularization on the clean subset to further improve model performance. Comprehensive experiments and ablation analysis on synthetic and real-world noisy datasets validated the superiority of our SED.

## Acknowledgements

This work was supported by the National Natural Science Foundation of China (No. 62202227 and 62102182), Natural Science Foundation of Jiangsu Province (No. BK20220938 and BK20220934), China Postdoctoral Science Foundation (No. 2022M711635), Jiangsu Funding Program for Excellent Postdoctoral Talent (No. 2022ZB267), Fundamental Research Funds for the Central Universities (No.30923010303).

## References

1. Arpit, D., Jastrzebski, S., Ballas, N., Krueger, D., Bengio, E., Kanwal, M.S., Maharaaj, T., Fischer, A., Courville, A.C., Bengio, Y., Lacoste-Julien, S.: A closer look at memorization in deep networks. In: *Int. Conf. Mach. Learn.* pp. 233–242 (2017)
2. Bai, Y., Yang, E., Han, B., Yang, Y., Li, J., Mao, Y., Niu, G., Liu, T.: Understanding and improving early stopping for learning with noisy labels. In: *Adv. Neural Inform. Process. Syst.* pp. 24392–24403 (2021)
3. Berthelot, D., Carlini, N., Goodfellow, I.J., Papernot, N., Oliver, A., Raffel, C.: Mixmatch: A holistic approach to semi-supervised learning. In: *Adv. Neural Inform. Process. Syst.* pp. 5050–5060 (2019)
4. Berthon, A., Han, B., Niu, G., Liu, T., Sugiyama, M.: Confidence scores make instance-dependent label-noise learning possible. In: *Int. Conf. Mach. Learn.* pp. 825–836 (2021)
5. Boutros, F., Damer, N., Kirchbuchner, F., Kuijper, A.: Elasticface: Elastic margin loss for deep face recognition. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 1577–1586 (2022)
6. Bucarelli, M.S., Cassano, L., Siciliano, F., Mantrach, A., Silvestri, F.: Leveraging inter-rater agreement for classification in the presence of noisy labels. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 3439–3448 (2023)
7. Chen, H., Tao, R., Fan, Y., Wang, Y., Wang, J., Schiele, B., Xie, X., Raj, B., Savvides, M.: Softmatch: Addressing the quantity-quality trade-off in semi-supervised learning. In: *Int. Conf. Learn. Represent.* (2023)
8. Chen, T., Yao, Y., Tang, J.: Multi-granularity denoising and bidirectional alignment for weakly supervised semantic segmentation. *IEEE Trans. Multimedia* **32**, 2960–2971 (2023)
9. Chen, T., Yao, Y., Zhang, L., Wang, Q., Xie, G., Shen, F.: Saliency guided inter- and intra-class relation constraints for weakly supervised semantic segmentation. *IEEE Trans. Multimedia* **25**, 1727–1737 (2023)
10. Chen, Y., Tan, X., Zhao, B., Chen, Z., Song, R., Liang, J., Lu, X.: Boosting semi-supervised learning by exploiting all unlabeled data. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 7548–7557. *IEEE* (2023)
11. Cheng, D., Liu, T., Ning, Y., Wang, N., Han, B., Niu, G., Gao, X., Sugiyama, M.: Instance-dependent label-noise learning with manifold-regularized transition matrix estimation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 16609–16618 (2022)
12. Deng, J., Dong, W., Socher, R., Li, L., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 248–255 (2009)

13. Fang, T., Lu, N., Niu, G., Sugiyama, M.: Rethinking importance weighting for deep learning under distribution shift. In: *Adv. Neural Inform. Process. Syst.* (2020)
14. Fergus, R., Fei-Fei, L., Perona, P., Zisserman, A.: Learning object categories from internet image searches. *Proc. IEEE* pp. 1453–1466 (2010)
15. Goldberger, J., Ben-Reuven, E.: Training deep neural-networks using a noise adaptation layer. In: *Int. Conf. Learn. Represent.* (2017)
16. Goldberger, J., Ben-Reuven, E.: Training deep neural-networks using a noise adaptation layer. In: *Int. Conf. Learn. Represent.* (2017)
17. Gong, C., Ding, Y., Han, B., Niu, G., Yang, J., You, J., Tao, D., Sugiyama, M.: Class-wise denoising for robust learning under label noise. *IEEE Trans. Pattern Anal. Mach. Intell.* pp. 2835–2848 (2023)
18. Gui, X., Wang, W., Tian, Z.: Towards understanding deep learning from noisy labels with small-loss criterion. In: *IJCAI*. pp. 2469–2475 (2021)
19. 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: *Adv. Neural Inform. Process. Syst.* pp. 8536–8546 (2018)
20. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE Conference on Computer Vision and Pattern Recognition*. pp. 770–778 (2016)
21. Hendrycks, D., Mazeika, M., Kadavath, S., Song, D.: Using self-supervised learning can improve model robustness and uncertainty. In: *Adv. Neural Inform. Process. Syst.* pp. 15637–15648 (2019)
22. Huang, L., Zhang, C., Zhang, H.: Self-adaptive training: beyond empirical risk minimization. In: *Adv. Neural Inform. Process. Syst.* (2020)
23. 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: *Int. Conf. Mach. Learn.* pp. 2309–2318 (2018)
24. Jiang, X., Liu, S., Dai, X., Hu, G., Huang, X., Yao, Y., Xie, G.S., Shao, L.: Deep metric learning based on meta-mining strategy with semiglobal information. *IEEE Trans. Neural. Netw. Learn. Syst.* **35**(4), 5103–5116 (2024)
25. Karim, N., Rizve, M.N., Rahnavard, N., Mian, A., Shah, M.: UNICON: combating label noise through uniform selection and contrastive learning. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 9666–9676 (2022)
26. Krizhevsky, A.: Learning multiple layers of features from tiny images (2009)
27. Lee, K.H., He, X., Zhang, L., Yang, L.: Cleannet: Transfer learning for scalable image classifier training with label noise. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 5447–5456 (2018)
28. Li, J., Li, G., Liu, F., Yu, Y.: Neighborhood collective estimation for noisy label identification and correction. In: *Eur. Conf. Comput. Vis.* pp. 128–145 (2022)
29. Li, J., Socher, R., Hoi, S.C.: Dividemix: Learning with noisy labels as semi-supervised learning. In: *Int. Conf. Learn. Represent.* (2020)
30. Li, S., Xia, X., Ge, S., Liu, T.: Selective-supervised contrastive learning with noisy labels. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 316–325 (2022)
31. Li, Y., Han, H., Shan, S., Chen, X.: DISC: learning from noisy labels via dynamic instance-specific selection and correction. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 24070–24079 (2023)
32. Liu, H., Sheng, M., Sun, Z., Yao, Y., Hua, X.S., Shen, H.T.: Learning with imbalanced noisy data by preventing bias in sample selection. *IEEE Trans. Multimedia* **26**, 7426–7437 (2024)



33. Liu, H., Zhang, C., Yao, Y., Wei, X., Shen, F., Tang, Z., Zhang, J.: Exploiting web images for fine-grained visual recognition by eliminating open-set noise and utilizing hard examples. *IEEE Trans. Multimedia* **24**, 546–557 (2022)
34. Liu, H., Zhang, H., Lu, J., Tang, Z.: Exploiting web images for fine-grained visual recognition via dynamic loss correction and global sample selection. *IEEE Trans. Multimedia* **24**, 1105–1115 (2022)
35. Liu, S., Zhu, Z., Qu, Q., You, C.: Robust training under label noise by over-parameterization. In: *Int. Conf. Mach. Learn.* pp. 14153–14172 (2022)
36. Lu, Y., He, W.: SELC: self-ensemble label correction improves learning with noisy labels. In: *IJCAI*. pp. 3278–3284 (2022)
37. Malach, E., Shalev-Shwartz, S.: Decoupling "when to update" from "how to update". In: *Adv. Neural Inform. Process. Syst.* pp. 960–970 (2017)
38. Mandal, D., Bharadwaj, S., Biswas, S.: A novel self-supervised re-labeling approach for training with noisy labels. In: *IEEE Winter Conference on Applications of Computer Vision*. pp. 1370–1379 (2020)
39. Mao, J., Yao, Y., Sun, Z., Huang, X., Shen, F., Shen, H.T.: Attention map guided transformer pruning for occluded person re-identification on edge device. *IEEE Trans. Multimedia* **25**, 1592–1599 (2023)
40. Patrini, G., Rozza, A., Krishna Menon, A., Nock, R., Qu, L.: Making deep neural networks robust to label noise: A loss correction approach. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 1944–1952 (2017)
41. Peng, X., Wang, K., Zeng, Z., Li, Q., Yang, J., Qiao, Y.: Suppressing mislabeled data via grouping and self-attention. In: *Eur. Conf. Comput. Vis.* pp. 786–802 (2020)
42. Redmon, J., Farhadi, A.: YOLO9000: better, faster, stronger. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 6517–6525 (2017)
43. Ren, M., Zeng, W., Yang, B., Urtasun, R.: Learning to reweight examples for robust deep learning. In: *Int. Conf. Mach. Learn.* pp. 4331–4340 (2018)
44. Ren, S., He, K., Girshick, R.B., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. In: *Adv. Neural Inform. Process. Syst.* pp. 91–99 (2017)
45. Sheng, M., Sun, Z., Cai, Z., Chen, T., Zhou, Y., Yao, Y.: Adaptive integration of partial label learning and negative learning for enhanced noisy label learning. In: *AAAI*. pp. 4820–4828 (2024)
46. Shi, X., Guo, Z., Li, K., Liang, Y., Zhu, X.: Self-paced resistance learning against overfitting on noisy labels. *Pattern Recognition* **134**, 109080 (2023)
47. 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: *Adv. Neural Inform. Process. Syst.* pp. 1917–1928 (2019)
48. Sosea, T., Caragea, C.: Marginmatch: Improving semi-supervised learning with pseudo-margins. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 15773–15782 (2023)
49. Sun, Z., Hua, X.S., Yao, Y., Wei, X.S., Hu, G., Zhang, J.: Crssc: salvage reusable samples from noisy data for robust learning. In: *ACM Int. Conf. Multimedia*. pp. 92–101 (2020)
50. Sun, Z., Liu, H., Wang, Q., Zhou, T., Wu, Q., Tang, Z.: Co-ldl: A co-training-based label distribution learning method for tackling label noise. *IEEE Trans. Multimedia* pp. 1093–1104 (2022)
51. Sun, Z., Shen, F., Huang, D., Wang, Q., Shu, X., Yao, Y., Tang, J.: Pnp: Robust learning from noisy labels by probabilistic noise prediction. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 5311–5320 (2022)

52. Sun, Z., Yao, Y., Wei, X.S., Zhang, Y., Shen, F., Wu, J., Zhang, J., Shen, H.T.: Webly supervised fine-grained recognition: Benchmark datasets and an approach. In: *Int. Conf. Comput. Vis.* pp. 10602–10611 (2021)
53. Tu, Y., Zhang, B., Li, Y., Liu, L., Li, J., Wang, Y., Wang, C., Zhao, C.: Learning from noisy labels with decoupled meta label purifier. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 19934–19943 (2023)
54. Tu, Y., Zhang, B., Li, Y., Liu, L., Li, J., Zhang, J., Wang, Y., Wang, C., Zhao, C.: Learning with noisy labels via self-supervised adversarial noisy masking. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 16186–16195 (2023)
55. Vahdat, A.: Toward robustness against label noise in training deep discriminative neural networks. In: *Adv. Neural Inform. Process. Syst.* pp. 5596–5605 (2017)
56. Veit, A., Alldrin, N., Chechik, G., Krasin, I., Gupta, A., Belongie, S.: Learning from noisy large-scale datasets with minimal supervision. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 6575–6583 (2017)
57. Wang, X., Hua, Y., Kodirov, E., Clifton, D.A., Robertson, N.M.: Proselflc: Progressive self label correction for training robust deep neural networks. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 752–761 (2021)
58. Wei, H., Feng, L., Chen, X., An, B.: Combating noisy labels by agreement: A joint training method with co-regularization. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 13723–13732 (2020)
59. Wei, J., Liu, H., Liu, T., Niu, G., Sugiyama, M., Liu, Y.: To smooth or not? when label smoothing meets noisy labels. In: *Int. Conf. Mach. Learn.* pp. 23589–23614 (2022)
60. Welinder, P., Branson, S., Belongie, S.J., Perona, P.: The multidimensional wisdom of crowds. In: *Adv. Neural Inform. Process. Syst.* pp. 2424–2432 (2010)
61. Wu, T., Dai, B., Chen, S., Qu, Y., Xie, Y.: Meta segmentation network for ultra-resolution medical images. In: *IJCAI* pp. 544–550 (2020)
62. Xia, X., Han, B., Wang, N., Deng, J., Li, J., Mao, Y., Liu, T.: Extended  $\$t\$$ : Learning with mixed closed-set and open-set noisy labels. *IEEE Trans. Pattern Anal. Mach. Intell.* pp. 3047–3058 (2023)
63. Xia, X., Liu, T., Han, B., Gong, M., Yu, J., Niu, G., Sugiyama, M.: Sample selection with uncertainty of losses for learning with noisy labels. In: *Int. Conf. Learn. Represent.* (2022)
64. Xia, X., Liu, T., Wang, N., Han, B., Gong, C., Niu, G., Sugiyama, M.: Are anchor points really indispensable in label-noise learning? In: *Adv. Neural Inform. Process. Syst.* pp. 6835–6846 (2019)
65. Yang, E., Yao, D., Liu, T., Deng, C.: Mutual quantization for cross-modal search with noisy labels. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 7541–7550 (2022)
66. Yao, Y., Sun, Z., Zhang, C., Shen, F., Wu, Q., Zhang, J., Tang, Z.: Jo-src: A contrastive approach for combating noisy labels. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 5192–5201 (2021)
67. 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: *Adv. Neural Inform. Process. Syst.* (2021)
68. Yi, K., Wu, J.: Probabilistic end-to-end noise correction for learning with noisy labels. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 7017–7025 (2019)
69. Yu, X., Han, B., Yao, J., Niu, G., Tsang, I.W., Sugiyama, M.: How does disagreement help generalization against label corruption? In: *Int. Conf. Mach. Learn.* pp. 7164–7173 (2019)

- 70. Zhang, C., Bengio, S., Hardt, M., Recht, B., Vinyals, O.: Understanding deep learning requires rethinking generalization. In: Int. Conf. Learn. Represent. (2017)
- 71. Zhou, X., Liu, X., Zhai, D., Jiang, J., Ji, X.: Asymmetric loss functions for noise-tolerant learning: Theory and applications. *IEEE Trans. Pattern Anal. Mach. Intell.* **45**(7), 8094–8109 (2023)