

## L2B: Learning to Bootstrap Robust Models for Combating Label Noise

Yuyin Zhou<sup>1\*</sup> Xianhang Li<sup>1\*</sup> Fengze Liu<sup>2</sup> Qingyue Wei<sup>5</sup>  
Xuxi Chen<sup>3</sup> Lequan Yu<sup>4</sup> Cihang Xie<sup>1</sup> Matthew P. Lungren<sup>5</sup> Lei Xing<sup>5</sup>

\*equal contribution

<sup>1</sup>University of California, Santa Cruz <sup>2</sup>Johns Hopkins University <sup>3</sup>UT Austin

<sup>4</sup>The University of Hong Kong <sup>5</sup>Stanford University

### Abstract

*Deep neural networks have shown great success in representation learning. However, when learning with noisy labels (LNL), they can easily overfit and fail to generalize to new data. This paper introduces a simple and effective method, named Learning to Bootstrap (L2B), which enables models to bootstrap themselves using their own predictions without being adversely affected by erroneous pseudo-labels. It achieves this by dynamically adjusting the importance weight between real observed and generated labels, as well as between different samples through meta-learning. Unlike existing instance reweighting methods, the key to our method lies in a new, versatile objective that enables implicit relabeling concurrently, leading to significant improvements without incurring additional costs.*

*L2B offers several benefits over the baseline methods. It yields more robust models that are less susceptible to the impact of noisy labels by guiding the bootstrapping procedure more effectively. It better exploits the valuable information contained in corrupted instances by adapting the weights of both instances and labels. Furthermore, L2B is compatible with existing LNL methods and delivers competitive results spanning natural and medical imaging tasks including classification and segmentation under both synthetic and real-world noise. Extensive experiments demonstrate that our method effectively mitigates the challenges of noisy labels, often necessitating few to no validation samples, and is well generalized to other tasks such as image segmentation. This not only positions it as a robust complement to existing LNL techniques but also underscores its practical applicability. The code and models are available at <https://github.com/yuyinzhou/l2b>.*

### 1. Introduction

In computer vision, deep learning has made significant strides, especially when provided with extensive, high-

quality datasets. However, the persistent issue of label noise in real-world datasets, which stems from factors such as inter-observer variability, human annotation errors, and adversarial rival, can significantly undermine performance [31]. As the size of datasets for deep learning continues to grow, the impact of label noise may become more significant. Understanding and addressing label noise is crucial for improving the accuracy and reliability of deep learning models [28, 49, 53, 60, 70, 71, 73].

Existing learning with noisy labels (LNL) methods, such as [9, 35], focus on loss correction by estimating a noise corruption matrix, which is often challenging and involves assumptions [13, 29, 54]. Recent research like [10, 15, 62] primarily targets identifying and utilizing clean samples within noisy datasets, frequently treating low-loss samples as clean [3]. Unlike approaches that discard noisy examples, meta-learning methods [38, 40] assign adaptive weights to each sample, with noisier ones receiving lower weights. However, this may compromise performance in high-noise scenarios by neglecting or underweighting portions of the training data. To better utilize corrupted samples, several studies have focused on using network predictions, or pseudo-labels [19], to recalibrate labels [2, 37, 43, 44, 61]. The bootstrapping loss method [37] is notable for using pseudo-labels in training targets, countering noisy sample effects. However, the static weight of pseudo-labels can lead to overfitting and inadequate label correction [2]. Addressing this, Arazo et al. [2] developed a dynamic bootstrapping method that adjusts the balance between actual and pseudo-labels using a mixture model.

In contrast to prior works that individually reweight labels or instances, our paper introduces a simple and effective approach to concurrently adjust both, elegantly unified under a meta-learning framework. We term our method as **Learning to Bootstrap (L2B)**, as our goal is to enable the network to self-boost its capabilities by harnessing its own predictions in combating label noise. Specifically, L2B introduces a new, versatile loss that allows dynamically ad-

justing the balance between true and pseudo labels, as well as the weights of individual samples. This adjustment is based on performance metrics from a separate, clean validation set within a meta-network framework. Unlike previous bootstrapping loss methods [2, 37, 68], which reallocate labels through a fixed weighted combination of pseudo and true labels, L2B offers greater flexibility. It uniquely does not limit the weights to sum to one, allowing for more nuanced reweighting across all instances and labels. Furthermore, we empirically show that meta-learning algorithms’ need for a clean validation set can be removed by dynamically creating an online meta set from the training data using a Gaussian mixture model [36]. This not only enhances our method’s practicality but also facilitates its integration with current LNL techniques like DivideMix [22], UniCon [16], and C2D [69]. Consequently, L2B attains superior results without relying on a validation set.

In addition, we theoretically prove that our formulation, which reweights different loss terms, can be reduced to the original bootstrapping loss and therefore conducts an implicit relabeling instead. Through a meta-learning process, L2B achieves significant improvements (e.g., **+8.9%** improvement on CIFAR-100 with 50% noise) compared with the instance reweighting baseline with no extra cost. This versatile bootstrapping procedure of L2B presents a simple and effective plug-in compatible with existing LNL methods. Our comprehensive tests across both natural and medical image datasets such as CIFAR-10, CIFAR-100, Clothing 1M, and ISIC2019, covering various types of label noise and recognition tasks, highlight L2B’s superiority over contemporary label correction and meta-learning techniques.

## 2. Related Works

**Learning with noisy labels.** Various approaches have been proposed to tackle the challenge of training models with noisy labeled data. Noisy detection approaches [10, 13, 58, 63] focus on identifying and reducing the influence of noisy samples to mitigate label inaccuracies. Label correction strategies [22, 37, 68] aim to refine pseudo labels to better match true labels. Within this domain, one group emphasizes robust representation learning through unsupervised contrastive learning [8, 16, 23, 69], while another significant group employs meta-learning, using a subset of clean data for optimization guidance [21, 38, 40, 52, 56, 68, 70].

**Explicit relabeling.** Existing works propose to directly identify noisy samples and relabel them through estimating the noise transition matrix [9, 35, 54, 59] or modeling noise by graph models or neural networks [20, 46, 47, 55]. Patrini *et al.* [35] and Hendrycks *et al.* [13] estimate the label corruption matrix to directly correct the loss function. However, these methods usually require assumptions about noise modeling. For instance, Hendrycks *et al.* [13] assume that

the noisy label is only dependent on the true label and independent of the data. Another line of approaches proposes to leverage the network prediction (pseudo-labels) for explicit relabeling [11, 33, 37, 44, 61]. However, using a uniform weight for all samples, as in [37], can exacerbate the influence of noisy data, impeding effective label correction. Semi-supervised LNL techniques [22, 68] segment training data into labeled “clean samples” and unlabeled noisy sets, subsequently relabeled using pseudo-labels. To bolster the reliability of these pseudo-labels, unsupervised contrastive learning approaches are employed [8, 16, 23, 69].

**Instance reweighting.** To counteract the adverse effects of corrupted examples, various strategies focus on reweighting or selecting training instances to minimize the influence of noisy samples [6, 15, 38]. Based on the observation that deep neural networks tend to learn simple patterns first before fitting label noise [3], many methods treat samples with small loss as clean ones [10, 15, 39, 50, 62]. Rather than directly selecting clean examples for training, meta-learning techniques [38, 40, 56] adjust instance weights, and curriculum learning [15] sequences them by noise levels. Such strategies enhance robustness in medical imaging [30, 57], but overlooking training subsets can affect performance in high-noise scenarios.

**Meta-learning.** Meta-learning methods [21, 38, 40, 52, 56, 68, 70] use a small clean validation set to optimize model weights and hyper-parameters. Techniques include instance reweighting [38, 40, 56], which involves bi-level optimization for determining training sample contributions. Another line of works view label correction as a separate meta-process [52, 68, 70]. Meta-learning has also been utilized to prevent overfitting to noisy labels [21]. Recently, meta-learning approaches has also been modified for other purposes. For instance, CMW-Net [41] adaptively generates sample weight based on the intrinsic bias characteristics of different sample classes. DMLP [45] combines self-supervised representation learning and a linear meta-learner for label correction.

Different from the aforementioned approaches which separately handle instance reweighting and label reweighting, we introduce a generic learning objective that concurrently meta-learns per-sample loss weights while implicitly relabeling the training data.

## 3. Methodology

### 3.1. Preliminary

Given a set of  $N$  training samples, i.e.,  $\mathcal{D}_{tra} = \{(x_i, y_i) | i = 1, \dots, N\}$ , where  $x_i \in \mathbb{R}^{W \times H}$  denotes the  $i$ -th image and  $y_i$  is the observed noisy label. In this work, we also assume that there is a small unbiased and clean validation set  $\mathcal{D}_{val} = \{(x_i^v, y_i^v) | i = 1, \dots, M\}$  and  $M \ll N$ , where the superscript  $v$  denotes the validation set. Let

$\mathcal{F}(\cdot, \theta)$  denote the neural network model parameterized by  $\theta$ . Given an input-target pair  $(x, y)$ , we consider the loss function of  $\mathcal{L}(\mathcal{F}(x, \theta), y)$  (e.g., cross-entropy loss) to minimize during the training process. Our goal in this paper is to properly utilize the small validation set  $\mathcal{D}_{val}$  to guide the model training on  $\mathcal{D}_{tra}$ , for reducing the negative effects brought by the noisy annotation.

To establish a more robust training procedure, [37] proposed the bootstrapping loss to enable the learner to “disagree” with the original training label, and effectively re-label the data during the training. Specifically, the training targets will be generated using a convex combination of training labels and predictions of the current model (i.e., pseudo-labels [19]), for purifying the training labels. Therefore, for a  $L$ -class classification problem, the loss function for optimizing  $\theta$  can be derived as follows:

$$y_i^{\text{pseudo}} = \arg \max_{l=1, \dots, L} \mathcal{P}(x_i, \theta), \quad (1)$$

$$\theta^* = \arg \min_{\theta} \sum_{i=1}^N \mathcal{L}(\mathcal{F}(x_i, \theta), \beta y_i^{\text{real}} + (1-\beta) y_i^{\text{pseudo}}), \quad (2)$$

where  $\beta$  is used for balancing the weight between the real labels and the pseudo-labels.  $\mathcal{P}(x_i, \theta)$  is the model output.  $y_i^{\text{real}}$  and  $y_i^{\text{pseudo}}$  denote the observed label and the pseudo-label respectively. However, in this method,  $\beta$  is manually selected and fixed for all training samples, which does not prevent fitting the noisy ones and can even lead to low-quality label correction [2]. Moreover, we observe that this method is quite sensitive to the selection of the hyperparameter  $\beta$ . For instance, as shown in Figure 1(a), even a similar  $\beta$  selection (i.e.,  $\beta = 0.6$  vs.  $\beta = 0.8$ ) behaves differently under disparate noise levels, making the selection of  $\beta$  even more intractable. Another limitation lies in that equation 2 treats all examples as equally important during training, which could easily cause overfitting for biased training data.

### 3.2. Learning to Bootstrap through Meta-Learning

To address these above challenges, in this paper, we aim to learn to bootstrap the model by conducting a joint label reweighting and instance reweighting. To achieve this, we propose to generate meta-learned weights for guiding our main learning objective:

$$\theta^*(\alpha, \beta) = \arg \min_{\theta} \sum_{i=1}^N \alpha_i \mathcal{L}(\mathcal{F}(x_i, \theta), y_i^{\text{real}}) + \beta_i \mathcal{L}(\mathcal{F}(x_i, \theta), y_i^{\text{pseudo}}), \quad (3)$$

with  $\{\alpha_i, \beta_i\}_{i=1}^N$  being the balance weights. Here we note that this new learning objective can be regarded as a general form of the original bootstrapping loss, as equation 3

can be reduced to equation 2 when  $\alpha_i + \beta_i = 1$  given that  $\mathcal{L}(\cdot)$  is the cross-entropy loss (see details in Appendix ??). By relaxing this constraint such that  $\alpha, \beta \geq \mathbf{0}$ , we can see that the optimization of equation 3 not only allows the main learner to explore the optimal combination between the two loss terms but also concurrently adjust the contribution of different training samples. In addition, compared with equation 2, the optimization of equation 3 does not rely on explicitly generating new training targets (i.e.,  $\beta y_i^{\text{real}} + (1-\beta) y_i^{\text{pseudo}}$ ), but rather conducts implicit relabeling during training by reweighting different loss terms. We note that the key to L2B is that the sum of  $\alpha_i$  and  $\beta_i$  need not be 1, which results in **+8.9%** improvement on CIFAR-100 with 50% noise (Section 4.3).

Note that this form is also similar to self-distillation in [25]. But different from [25] where the weights are determined by heuristics, our weights  $\alpha, \beta$  are meta-learned based on its performance on the validation set  $\mathcal{D}_{val}$ , that is

$$\alpha^*, \beta^* = \arg \min_{\alpha, \beta \geq \mathbf{0}} \frac{1}{M} \sum_{i=1}^M \mathcal{L}(\mathcal{F}(x_i^v, \theta^*(\alpha, \beta)), y_i^v). \quad (4)$$

It is necessary to constrain  $\alpha_i, \beta_i \geq 0$  for all  $i$  to avoid potential unstable training [38]. Both the meta learner (i.e., equation 4) and the main learner (i.e., equation 3) are optimized concurrently, which allows the model to maximize the performance on the clean validation set  $\mathcal{D}_{val}$  by adjusting the importance weights of the observed and the pseudo-labels in a differentiable manner.

**Online Approximation.** For each step  $t$  at training, a mini-batch of training examples  $\{(x_i, y_i), 1 \leq i \leq n\}$  with  $n \ll N$  is sampled to estimate a temporary adjustment to the parameters based on the descent direction of the loss function. For simplicity, let  $f_i(\theta)$  denote  $\mathcal{L}(\mathcal{F}(x_i, \theta), y_i^{\text{real}})$  and  $g_i(\theta)$  denote  $\mathcal{L}(\mathcal{F}(x_i, \theta), y_i^{\text{pseudo}})$  in the following sections. Given any  $\alpha, \beta$ , we use

$$\hat{\theta}_{t+1} = \theta_t - \lambda \nabla \left( \sum_{i=1}^n \alpha_i f_i(\theta) + \beta_i g_i(\theta) \right) \Big|_{\theta=\theta_t} \quad (5)$$

to approach the solution of equation 3. Here  $\lambda$  is the step size. We then estimate the corresponding optimal  $\alpha, \beta$  as

$$\alpha_t^*, \beta_t^* = \arg \min_{\alpha, \beta \geq \mathbf{0}} \frac{1}{M} \sum_{i=1}^M f_i^v(\hat{\theta}_{t+1}). \quad (6)$$

However, directly solving for equation 6 at every training step requires too much computation cost. To reduce the computational complexity, we apply one step gradient descent of  $\alpha_t, \beta_t$  on a mini-batch of validation set  $\{(x_i^v, y_i^v), 1 \leq i \leq m\}$  with  $m \leq M$  as an approximation. Specifically,

$$(\alpha_{t,i}, \beta_{t,i}) = -\eta \nabla \left( \sum_{i=1}^m f_i^v(\hat{\theta}_{t+1}) \right) \Big|_{\alpha_i=0, \beta_i=0}, \quad (7)$$

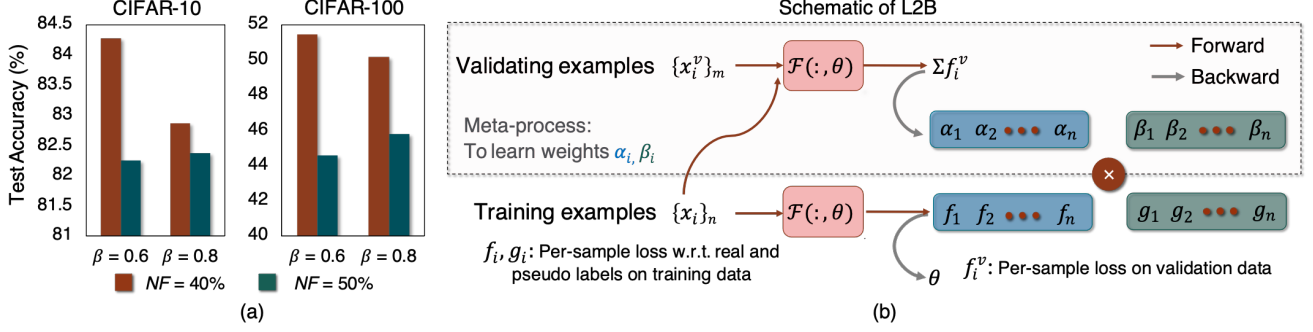


Figure 1. (a) The original bootstrapping loss [37] is sensitive to the reweighting hyper-parameter  $\beta$ . Under different noise levels, the optimal  $\beta$  is different ( $NF$  stands for noise fraction). (b) Schematic description of our Learning to Bootstrap (i.e., **L2B**) method. The reweighting hyper-parameters are learned in a meta-process.

where  $\eta$  is the step size for updating  $\alpha, \beta$ . To ensure that the weights are non-negative, we apply the following rectified function:

$$\tilde{\alpha}_{t,i} = \max(\alpha_{t,i}, 0), \tilde{\beta}_{t,i} = \max(\beta_{t,i}, 0). \quad (8)$$

To stabilize the training process, we also normalize the weights in a single training batch so that they sum up to one:

$$\tilde{\alpha}_{t,i} = \frac{\tilde{\alpha}_{t,i}}{\sum_{i=1}^n \tilde{\alpha}_{t,i} + \tilde{\beta}_{t,i}}, \tilde{\beta}_{t,i} = \frac{\tilde{\beta}_{t,i}}{\sum_{i=1}^n \tilde{\alpha}_{t,i} + \tilde{\beta}_{t,i}}. \quad (9)$$

Finally, we estimate  $\theta_{t+1}$  based on the updated  $\alpha_t, \beta_t$  so that  $\theta_{t+1}$  can consider the meta information included in  $\alpha_t, \beta_t$ :

$$\theta_{t+1} = \theta_t - \lambda \nabla \left( \sum_{i=1}^n \tilde{\alpha}_{t,i} f_i(\theta) + \tilde{\beta}_{t,i} g_i(\theta) \right) \Big|_{\theta=\theta_t}. \quad (10)$$

See Appendix ?? for detailed calculation of the gradient in equation 10. A schematic description of our Learning to Bootstrap algorithm is illustrated in Figure 1(b) and the overall optimization procedure can be found in Algorithm 1.

### 3.3. Convergence Analysis

In proposing equation 3, we show that with the first-order approximation of  $\alpha, \beta$  in equation 7 and some mild assumptions, our method guarantees to convergence to a local minimum point of the validation loss, which yields the best combination of  $\alpha, \beta$ . Details of the proof are provided in Appendix ??.

## 4. Experiments

### 4.1. Datasets

**CIFAR-10 & CIFAR-100.** Both CIFAR-10 and CIFAR-100 contain 50K training images and 10K test images of size  $32 \times 32$ . Following previous works [17, 22, 44], we

### Algorithm 1 Learning to Bootstrap

**Require:**  $\theta_0, \mathcal{D}_{tra}, \mathcal{D}_{val}, n, m, L$

**Ensure:**  $\theta_T$

- 1: **for**  $t = 0 \dots T - 1$  **do**
- 2:  $\{x_i, y_i\} \leftarrow \text{SampleMiniBatch}(\mathcal{D}_{tra}, n)$
- 3:  $\{x_i^v, y_i^v\} \leftarrow \text{SampleMiniBatch}(\mathcal{D}_{val}, m)$
- 4: For the  $i$ -th sample of  $\mathcal{D}_{tra}$ , compute  $y_i^{\text{pseudo}} = \arg \max_{l=1, \dots, L} \mathcal{P}(x_i, \theta_t)$
- 5: Learnable weights  $\alpha, \beta$
- 6: Compute training loss  $l_f \leftarrow \sum_{i=1}^n \alpha_i f_i(\theta_t) + \beta_i g_i(\theta_t)$
- 7:  $\hat{\theta}_{t+1} \leftarrow \theta_t - \lambda \nabla l_f \Big|_{\theta=\theta_t}$
- 8: Compute validation loss  $l_g \leftarrow \frac{1}{m} \sum_{i=1}^m f_i^v(\hat{\theta}_{t+1})$
- 9:  $(\alpha_t, \beta_t) \leftarrow -\eta \nabla l_g \Big|_{\alpha=0, \beta=0}$
- 10:  $\tilde{\alpha}_{t,i} \leftarrow \max(\alpha_{t,i}, 0), \tilde{\beta}_{t,i} \leftarrow \max(\beta_{t,i}, 0)$
- 11:  $\tilde{\alpha}_{t,i} \leftarrow \frac{\tilde{\alpha}_{t,i}}{\sum_{i=1}^n \tilde{\alpha}_{t,i} + \tilde{\beta}_{t,i}}, \tilde{\beta}_{t,i} \leftarrow \frac{\tilde{\beta}_{t,i}}{\sum_{i=1}^n \tilde{\alpha}_{t,i} + \tilde{\beta}_{t,i}}$
- 12: Apply learned weights  $\alpha, \beta$  to reweight the training loss as  $\hat{l}_f \leftarrow \sum_{i=1}^n \tilde{\alpha}_{t,i} f_i(\theta_t) + \tilde{\beta}_{t,i} g_i(\theta_t)$
- 13:  $\theta_{t+1} \leftarrow \theta_t - \lambda \nabla \hat{l}_f \Big|_{\theta=\theta_t}$
- 14: **end for**

experimented with both *symmetric* and *asymmetric* label noise. In our method, we used 1,000 clean images in the validation set  $\mathcal{D}_{val}$  following [13, 15, 38, 40, 70].

**ISIC2019.** Following [57], we also evaluated our algorithm on a medical image dataset, i.e., skin lesion classification data, under different symmetric noise levels. Our experiments were conducted on the 25,331 dermoscopic images of the 2019 ISIC Challenge<sup>1</sup>, where we used 20400 images as the training set  $\mathcal{D}_{tra}$ , 640 images as the validation set  $\mathcal{D}_{val}$ , and tested on 4291 images.

<sup>1</sup><https://challenge2019.isic-archive.com/data.html>



Table 1. Comparison in test accuracy (%) with the baseline methods on CIFAR-10/100 datasets with symmetric noise.

Dataset	CIFAR-10				CIFAR-100				ISIC			
	20%	30%	40%	50%	20%	30%	40%	50%	20%	30%	40%	50%
Method/Noise ratio	20%	30%	40%	50%	20%	30%	40%	50%	20%	30%	40%	50%
Cross-Entropy (CE)	86.9	84.9	83.3	81.3	59.6	52.2	49.2	44.4	79.4	77.5	75.3	73.7
Bootstrap [37]	85.2	84.8	82.9	79.2	61.8	54.2	50.2	45.8	80.8	77.7	75.7	74.8
L2RW [38]	90.6	89.0	86.6	85.3	67.8	63.8	59.7	55.6	80.1	77.7	76.3	74.1
L2B (Ours)	<b>92.2</b>	<b>90.7</b>	<b>89.9</b>	<b>88.5</b>	<b>71.8</b>	<b>69.5</b>	<b>67.3</b>	<b>64.5</b>	<b>81.1</b>	<b>80.2</b>	<b>78.6</b>	<b>76.8</b>

**Clothing 1M.** We evaluate on real-world noisy dataset, Clothing 1M [55], which has 1 million training images collected from online shopping websites with labels generated from surrounding texts. In addition, the Clothing 1M also provides an official validation set of 14,313 images and a test set of 10,526 images.

## 4.2. Implementation Details

For all CIFAR-10 and CIFAR-100 comparison experiments, we used an 18-layer PreActResNet [12] as the baseline network following the setups in [22], unless otherwise specified. The model was trained using SGD with a momentum of 0.9, a weight decay of 0.0005, and a batch size of 256 for CIFAR-100 and 512 for CIFAR-10. The network was trained from scratch for 300 epochs. We set the learning rate as 0.15 initially with a cosine annealing decay. Following [22], we set the warm up period as 10 epochs for both CIFAR-10 & CIFAR-100. The optimizer and the learning rate schedule remained the same for both the main and the meta model. Gradient clipping is applied to stabilize training. All experiments were conducted with one V100 GPU, except for the experiments on Clothing 1M which were conducted with one RTX A6000 GPU.

For ISIC2019 experiments, we used ResNet-50 with ImageNet pretrained weights. A batch size of 64 was used for training with an initial learning rate of 0.01. The network was trained for 30 epochs in total with the warmup period as 1 epoch. All other implementation details remained the same as above. For Clothing 1M experiments, we used an ImageNet pre-trained 18-layer ResNet [12] as our baseline. We finetuned the network with a learning rate of 0.005 for 300 epochs. The model was trained using SGD with a momentum of 0.9, a weight decay of 0.0005, and a batch size of 256. Following [22], to ensure the labels (noisy) were balanced, for each epoch, we sampled 250 mini-batches from the training data.

## 4.3. Performance Comparisons

**Efficacy of L2B.** We compare our method with different baselines: 1) Cross-Entropy (the standard training), 2) Bootstrap [37], which modifies the training loss by generating new training targets, and 3) L2RW [38], which reweights different instances through meta-learning under different levels of symmetric labels noise ranging from

20% ~ 50%. To ensure a fair comparison, we report the best epoch for all comparison approaches. All results are summarized in Table 1. Compared with the naive bootstrap method and the baseline meta-learning-based instance reweighting method L2RW, the performance improvement is substantial, especially under larger noise fraction, which suggests that using meta-learning to automatically bootstrap the model is more beneficial for LNL. For example, on CIFAR-100, the accuracy improvement of our proposed L2B reaches 7.6% and 8.9% under 40% and 50% noise fraction, respectively. We also demonstrate a set of qualitative examples to illustrate how our proposed L2B benefits from the joint instance and label reweighting paradigm in Figure 2. We can see that when the online estimated pseudo label is of high-quality, i.e., the pseudo label is different from the noisy label but equal to the clean label, our model will automatically assign a much higher weight to  $\beta$  for corrupted training samples. On the contrary,  $\alpha$  can be near zero in this case. This indicates that our L2B algorithm will pay more attention to the correct pseudo label than the real noisy label when computing the losses. In addition, we also show several cases where the online pseudo label has not yet been corrected and therefore is equal to the noisy label during the training process, where we can see that  $\alpha$  and  $\beta$  are almost identical under this circumstance since there will be no need to correct it. We note that by the end of the training, most noisy examples will be successfully corrected, leading to significantly different weighting of pseudo and noisy labels that will help rectify the training.

The relatively small values of  $\alpha$  and  $\beta$  are due to that we use a large batch size (i.e., 512) for CIFAR-10 experiments. By normalizing the weights in each training batch (see equation 9), the value of  $\alpha$  and  $\beta$  can be on the scale of  $10^{-4}$ .

**Comparison with the state-of-the-arts.** We compare our method with SOTA methods on CIFAR 10 and CIFAR 100 in Table 2. We demonstrate our L2B is compatible with existing LNL methods. When integrated with existing LNL methods like DivideMix [22], UniCon [16], C2D [69], L2B consistently enhances performance across varying noise ratios on both datasets. Notably, L2B-C2D surpasses all competing methods in various settings, achieving 94.4% and 60.7% accuracy under the noise ratio of 90% for CIFAR-10 and CIFAR-100. We also test our model

Table 2. Comparison in test accuracy (%) with state-of-the-art methods on CIFAR-10/100 datasets with symmetric noise.

Dataset	CIFAR-10				CIFAR-100			
	20%	50%	80%	90%	20%	50%	80%	90%
Method/Noise ratio								
Co-teaching+ [62]	89.5	85.7	67.4	47.9	65.6	51.8	27.9	13.7
Mixup [66]	95.6	87.1	71.6	52.2	67.8	57.3	30.8	14.6
PENCIL [61]	92.4	89.1	77.5	58.9	69.4	57.5	31.1	15.3
Meta-Learning [21]	92.9	89.3	77.4	58.7	68.5	59.2	42.4	19.5
M-correction [2]	94.0	92.0	86.8	69.1	73.9	66.1	48.2	24.3
AugDesc [32]	96.3	95.4	93.8	91.9	79.5	77.2	66.4	41.2
GCE [8]	90.0	89.3	73.9	36.5	68.1	53.3	22.1	8.9
Sel-CL+ [24]	95.5	93.9	89.2	81.9	76.5	72.4	59.6	48.8
MLC [70]	92.6	88.1	77.4	67.9	66.8	52.7	21.8	15.0
MSLC [52]	93.4	89.9	69.8	56.1	72.5	65.4	24.3	16.7
MOIT+ [34]	94.1	91.8	81.1	74.7	75.9	70.6	47.6	41.8
TCL [14]	95.0	93.9	92.5	89.4	78.0	73.3	65.0	54.5
DivideMix [22]	96.1	94.6	93.2	76.0	77.3	74.6	60.2	31.5
L2B-DivideMix	96.1	<b>95.4</b>	<b>94.0</b>	<b>91.3</b>	<b>77.9</b>	<b>75.9</b>	<b>62.2</b>	<b>35.8</b>
UniCon [16]	96.0	95.6	93.9	90.8	78.9	77.6	63.9	44.8
L2B-UniCon	<b>96.5</b>	<b>95.8</b>	<b>94.7</b>	<b>92.8</b>	78.8	77.3	<b>67.6</b>	<b>49.6</b>
C2D [69]	96.3	95.2	94.4	93.5	78.7	76.4	67.8	58.7
L2B-C2D	<b>96.7</b>	<b>95.6</b>	<b>94.8</b>	<b>94.4</b>	<b>80.1</b>	<b>78.1</b>	<b>69.6</b>	<b>60.7</b>

Table 3. Comparison with 40% asymmetric noise in test accuracy on the CIFAR-10 dataset.

Method	Acc
Cross-Entropy	85.0
F-correction [35]	87.2
M-correction [2]	87.4
Chen et al. [4]	88.6
P-correction [61]	88.5
REED [64]	92.3
Tanaka et al. [44]	88.9
NLNL [17]	89.9
JNPL [18]	90.7
DivideMix [22]	93.4
MLNT [21]	89.2
L2RW [38]	89.2
MW-Net [40]	89.7
MSLC [52]	91.6
Meta-Learning [21]	88.6
Distilling [68]	90.2
L2B-Naive (Ours)	91.8
L2B-C2D (Ours)	<b>94.0</b>

with 40% asymmetric noise and summarize the testing accuracy in Table 3. Among all compared methods, we reimplement L2RW under the same setting and report the performance of all other competitors from previous papers including [17, 18, 22]. Compared with previous meta-learning-based methods (e.g., [4], [64]), and other methods (e.g., [38], [52], [40]), our L2B achieves superior results.

Table 4. Comparison with state-of-the-art methods in test accuracy (%) on Clothing 1M.

Method	Acc (%)
Cross-Entropy	69.2
M-correction [2]	71.0
PENCIL [61]	73.5
DivideMix [22]	74.8
Nested [5]	74.9
AugDesc [32]	75.1
RRL [23]	74.9
GCE [8]	73.3
C2D [69]	74.3
MLNT [21]	73.5
MLC [70]	75.8
MSLC [52]	74.0
Meta-Cleaner [67]	72.5
Meta-Weight [40]	73.7
FaMUS [56]	74.4
MSLG [1]	76.0
DISC [26]	73.7
InstanceGM [7]	74.4
DivideMix+SNSCL [51]	75.3
<b>L2B-Naive (Ours)</b>	<b>77.5 ± 0.2</b>

**Generalization to real-world noisy labels.** We test L2B on Clothing 1M [55], a large-scale dataset with real-world noisy labels. The results of all competitors are reported from published papers. As shown in Table 4, our L2B-Naive attains an average performance of 77.5% accuracy from 3 independent runs with different random seeds, outperforming all competing methods.

Table 5. Segmentation performance comparison under noisy-supervision on PROMISE12.

Method	Dice (%) $\uparrow$	HD (voxel) $\downarrow$	ASD (voxel) $\downarrow$
UNet++ [72]	73.74	11.63	3.70
UNet++ meta	73.04	17.06	5.50
NL reweighting [30]	76.64	8.33	2.75
Mix-up [65]	69.18	13.25	4.56
L2B (Ours)	<b>80.83</b>	<b>6.68</b>	<b>2.10</b>

Table 6. L2B for segmentation under different noise levels.

Method	Dice (%) $\uparrow$
baseline - $L_1$	59.77
<b>L2B</b> - $L_1$	<b>77.70</b>
baseline - $L_2$	73.74
<b>L2B</b> - $L_2$	<b>80.83</b>
baseline - $L_3$	80.03
<b>L2B</b> - $L_3$	<b>82.01</b>

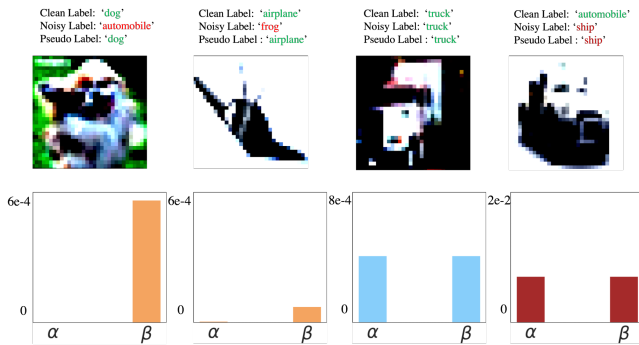


Figure 2. Examples of  $\alpha$  and  $\beta$  on CIFAR-10 with asymmetric noise fraction of 20%. When the estimated pseudo label is of high-quality, i.e., the pseudo label is different from the noisy label but equal to the clean label, our model will automatically assign a much higher weight to  $\beta$  than to  $\alpha$  for corrupted training samples. When the pseudo label is equal to the noisy label (i.e., the two loss terms are equal to each other),  $\alpha$  and  $\beta$  are almost identical.

**Generalization to image segmentation** L2B can be easily generalized for segmentation tasks. Specifically, the learnable weights  $\alpha$  and  $\beta$  are replaced with pixel-wise weight maps corresponding to noisy labels and pseudo labels (model predictions). L2B dynamically assigns these weight maps, adjusting for both noisy and pseudo labels to optimize the bootstrapping process via a meta-process. To assess L2B’s performance in segmentation, we employed the PROMISE12 dataset [27] which contains 50 3D transversal T2-weighted MR images. Specifically, 40/10 cases were used for training/evaluation. 3 out of the 40 training cases are chosen randomly as the meta set. Following [42, 48], we utilized 2D slices in the axial view for both training and testing. All images are resized to  $144 \times 144$  and splits are randomized. Noisy labels used in Table 5 were synthesized using random rotation, erosion, or dilation, achieving approximately a 60% corruption ratio and

an average Dice coefficient of 0.6206. And visualizations of the corrupted noisy labels (shown in yellow) as well as the ground-truth (shown in red) are illustrated in Figure ?? in Appendix. As presented in Table 5, we compare our method with 1) UNet++ [72], 2) UNet++ meta, which trains exclusively on the meta data, 3) NL reweighting [30], which only reweights the noisy labels, 4) Mix-up [65], a regularization based method. L2B outperforms others in all evaluation metrics of Dice, Hausdorff Distance (HD) and Average Surface Distance (ASD). Furthermore, we also investigate the robustness of our method by varying the noise level of the corrupted training set from  $\{L_1, L_2, L_3\}$ , where the average Dice coefficients are  $Dice_{L_1} = 0.4148$ ,  $Dice_{L_2} = 0.6206$ , and  $Dice_{L_3} = 0.8031$  (i.e., the corrupted ratios are around 60% ( $L_1$ ), 40% ( $L_2$ ), and 20% ( $L_3$ )). At each noise level, we compare the baseline UNet++ which is directly trained on the noisy training data with our generalized L2B. As shown in Table 6, we report the averaged dice coefficient over 5 repetitions for each series of experiments. The standard deviation for all experiments is within 0.5%. We could notice that while the noise level increases, performances of baseline drop from 80.03% to 59.77%, but performances of L2B only drop from 82.01% to 77.70% which indicates that our L2B is robust to different noisy levels and shows larger improvements under a much severer noisy situation.

**Qualitative Results** We also demonstrate a set of qualitative examples to illustrate how our proposed L2B benefits from the joint instance and label reweighting paradigm. In Figure 2, we can see that when the estimated pseudo label is of high-quality, i.e., the pseudo label is different from the noisy label but equal to the clean label, our model will automatically assign a much higher weight to  $\beta$  for corrupted training samples. On the contrary,  $\alpha$  can be near zero in this case. This indicates that our L2B algorithm will pay more attention to the pseudo label than the real noisy label when computing the losses. In addition, we also show several cases where the pseudo label is equal to the noisy label, where we can see that  $\alpha$  and  $\beta$  are almost identical under this circumstance since the two losses are of the same value. Note that the relatively small values of  $\alpha$  and  $\beta$  are due to that we use a large batch size (i.e., 512) for CIFAR-10 experiments. By normalizing the weights in each training batch (see equation 9), the value of  $\alpha$  and  $\beta$  can be on the scale of  $10^{-4}$ .

#### 4.4. Ablation Study

**On the importance of  $\alpha, \beta$ .** To understand why our proposed new learning objective can outperform previous meta-learning-based instance reweighting methods, we conduct the following analysis to understand the importance of hyper-parameter  $\alpha$  and  $\beta$  in our method. Specifically, we set  $\alpha = 0$  and  $\beta = 0$  respectively to investigate the importance of each loss term in equation 3. In addition, we also

Table 7. Ablation on size of validation data on CIFAR-10 and CIFAR-100 datasets.

	Validation Size	CIFAR-10				CIFAR-100			
		20%	50%	80%	90%	20%	50%	80%	90%
L2B-DivideMix	baseline	96.1	94.6	93.2	76.0	77.3	74.6	60.2	31.5
	0	96.3	95.3	93.5	82.6	77.6	75.3	60.8	31.0
	500	96.1	95.3	93.8	91.1	78.2	75.3	62.5	34.0
	1000	96.1	95.4	94.0	91.3	77.9	75.9	62.2	35.8
L2B-UniCon	baseline	96.0	95.6	93.9	90.8	78.9	77.6	63.9	44.8
	0	96.4	95.6	94.2	92.5	78.7	77.4	68.0	48.6
	500	96.3	95.6	94.5	92.7	78.5	77.5	67.8	51.1
	1000	96.5	95.8	94.7	92.8	78.8	77.3	67.6	49.6
L2B-C2D	baseline	96.4	95.3	94.4	93.5	78.7	76.4	67.8	58.7
	0	96.4	95.6	94.9	93.7	79.1	77.8	68.5	60.3
	500	96.6	95.5	94.9	94.0	79.5	77.9	69.0	60.8
	1000	96.7	95.6	94.8	94.4	80.1	78.1	69.6	60.7

Table 8. Ablation of  $\alpha, \beta$ . L2B ( $\alpha, \beta \geq 0$ ) consistently achieves superior results to L2B ( $\alpha + \beta = 1$ ) under different noise levels on CIFAR-100.

Method	20%	40%
baseline (CE)	59.6	49.2
$\alpha = 0$	55.7	47.1
$\beta = 0$	63.2	57.5
$\alpha + \beta = 1$	64.8	59.1
$\alpha, \beta \geq 0$	<b>71.8</b>	<b>67.3</b>

show how the restriction of  $\alpha_i + \beta_i = 1$  (equation 2) would deteriorate our model performance as follows.

- $\alpha = 0$ . As shown in Table 8, in this case, the performance even decreases compared with the baseline approach. This is due to that when only pseudo-labels are included in the loss computation, the error which occurs in the initial pseudo-label will be reinforced by the network during the following iterations.
- $\beta = 0$ . From equation 3, we can see that setting  $\beta$  as 0 is essentially equivalent to the baseline meta-learning-based instance reweighting method L2RW [38]. In this case, the performance is largely improved compared to the baseline, but still inferior to our method, which jointly optimizes  $\alpha$  and  $\beta$ .
- $\alpha + \beta = 1$ . We also investigate whether the restriction of  $\alpha + \beta = 1$  is required for obtaining optimal weights during the meta-update, as in [68]. As shown in Table 8, L2B ( $\alpha, \beta \geq 0$ ) consistently achieves superior results than L2B ( $\alpha + \beta = 1$ ) under different noise levels on CIFAR-100. The reason may be the latter is only reweighting different loss terms, whereas the former not only explores the optimal combination between the two loss terms but also jointly adjusts the contribution of different training samples.

**The number of clean validation samples** In Table 7, our L2B method is shown to require few to no validation samples for LNL problems, highlighting its practicality. In the

absence of a dedicated validation set, L2B adeptly generates an online meta set directly from the training data using a Gaussian mixture model [36], following [22]. L2B consistently boosts baseline methods such as DivideMix, UniCon, and C2D. Specifically, L2B-DivideMix has showcased its efficacy, particularly at high noise levels. Specifically, in a scenario with 90% noise on CIFAR-10, our approach outstripped the baseline by 8.7%, achieving an accuracy of 82.6% compared to 76.0%, and this was achieved without the need for clean validation samples. The advantage of L2B-DivideMix becomes even more pronounced when we incorporate a minimal amount of clean labels. With just 500 clean labels (equivalent to 2% of the training data), our performance lead over the baseline extends to a remarkable 15.1%. However, as we double the clean samples to 1000, the incremental benefit tapers off, yielding a mere 0.2% boost. This behavior underscores the efficiency of L2B-DivideMix, demonstrating that it can deliver impressive results with minimal or even no clean validation data, making it a highly adaptable and practical solution for real-world applications.

## 5. Conclusion

Our paper presents Learning to Bootstrap (L2B), a new technique using joint reweighting for model training. L2B dynamically balances weights between actual labels, pseudo-labels, and different samples, mitigating the challenges of erroneous pseudo-labels. Notably, L2B operates effectively without a clean validation set and can be well generalized to other tasks, highlighting its practicality in real-world settings. Extensive experiments on CIFAR-10, CIFAR-100, ISIC2019, and Clothing 1M datasets demonstrate the superiority and robustness compared to other existing methods under various settings.

**Acknowledgments** This work is partially supported by the AWS Public Sector Cloud Credit for Research Program.



## References

- [1] Görkem Algan and Ilkay Ulusoy. Meta soft label generation for noisy labels. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7142–7148. IEEE, 2021. 6
- [2] Eric Arazo, Diego Ortego, Paul Albert, Noel O’Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. In *International Conference on Machine Learning*, pages 312–321. PMLR, 2019. 1, 2, 3, 6
- [3] Devansh Arpit, Stanislaw Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *International Conference on Machine Learning*, pages 233–242. PMLR, 2017. 1, 2
- [4] Pengfei Chen, Benben Liao, Guangyong Chen, and Shengyu Zhang. Understanding and utilizing deep neural networks trained with noisy labels. In *ICML*, 2019. 6
- [5] Yingyi Chen, Xi Shen, Shell Xu Hu, and Johan AK Suykens. Boosting co-teaching with compression regularization for label noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2688–2692, 2021. 6
- [6] Tongtong Fang, Nan Lu, Gang Niu, and Masashi Sugiyama. Rethinking importance weighting for deep learning under distribution shift. In *Neurips*, 2020. 2
- [7] Arpit Garg, Cuong Nguyen, Rafael Felix, Thanh-Toan Do, and Gustavo Carneiro. Instance-dependent noisy label learning via graphical modelling. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2288–2298, 2023. 6
- [8] Aritra Ghosh and Andrew Lan. Contrastive learning improves model robustness under label noise. In *CVPR*, pages 2703–2708, 2021. 2, 6
- [9] Jacob Goldberger and Ehud Ben-Reuven. Training deep neural-networks using a noise adaptation layer. In *ICLR*, 2017. 1, 2
- [10] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *International Conference on Machine Learning*, 2018. 1, 2
- [11] Jiangfan Han, Ping Luo, and Xiaogang Wang. Deep self-learning from noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5138–5147, 2019. 2
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 5
- [13] Dan Hendrycks, Mantas Mazeika, Duncan Wilson, and Kevin Gimpel. Using trusted data to train deep networks on labels corrupted by severe noise. *Advances in neural information processing systems*, 31, 2018. 1, 2, 4
- [14] Zhizhong Huang, Junping Zhang, and Hongming Shan. Twin contrastive learning with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11661–11670, 2023. 6
- [15] 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 *International Conference on Machine Learning*, pages 2304–2313. PMLR, 2018. 1, 2, 4
- [16] Nazmul Karim, Mamshad Nayeem Rizve, Nazanin Rahnavard, Ajmal Mian, and Mubarak Shah. Unicon: Combating label noise through uniform selection and contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9676–9686, 2022. 2, 5, 6
- [17] Youngdong Kim, Junho Yim, Juseung Yun, and Junmo Kim. Nlnl: Negative learning for noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 101–110, 2019. 4, 6
- [18] Youngdong Kim, Juseung Yun, Hyounguk Shon, and Junmo Kim. Joint negative and positive learning for noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9442–9451, 2021. 6
- [19] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICML Workshop on challenges in representation learning*, 2013. 1, 3
- [20] Kuang-Huei Lee, Xiaodong He, Lei Zhang, and Linjun Yang. Cleannet: Transfer learning for scalable image classifier training with label noise. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5447–5456, 2018. 2
- [21] Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S Kankanhalli. Learning to learn from noisy labeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5051–5059, 2019. 2, 6
- [22] Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *ICLR*, 2020. 2, 4, 5, 6, 8
- [23] Junnan Li, Caiming Xiong, and Steven CH Hoi. Learning from noisy data with robust representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9485–9494, 2021. 2, 6
- [24] Shikun Li, Xiaobo Xia, Shiming Ge, and Tongliang Liu. Selective-supervised contrastive learning with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 316–325, 2022. 6
- [25] Yuncheng Li, Jianchao Yang, Yale Song, Liangliang Cao, Jiebo Luo, and Li-Jia Li. Learning from noisy labels with distillation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1910–1918, 2017. 3
- [26] Yifan Li, Hu Han, Shiguang Shan, and Xilin Chen. Disc: Learning from noisy labels via dynamic instance-specific selection and correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24070–24079, 2023. 6
- [27] Geert Litjens, Robert Toth, Wendy van de Ven, Caroline Hoeks, Sjoerd Kerkstra, Bram van Ginneken, Graham Vincent, Gwenael Guillard, Neil Birbeck, Jindang Zhang, et al.

- Evaluation of prostate segmentation algorithms for mri: the promise12 challenge. *Medical image analysis*, 18(2):359–373, 2014. 7
- [28] Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. *Advances in neural information processing systems*, 33:20331–20342, 2020. 1
- [29] Tongliang Liu and Dacheng Tao. Classification with noisy labels by importance reweighting. *IEEE Transactions on pattern analysis and machine intelligence*, 38(3):447–461, 2015. 1
- [30] Zahra Mirikharaji, Yiqi Yan, and Ghassan Hamarneh. Learning to segment skin lesions from noisy annotations. In *Domain Adaptation and Representation Transfer and Medical Image Learning with Less Labels and Imperfect Data*, pages 207–215. Springer, 2019. 2, 7
- [31] David F Nettleton, Albert Orriols-Puig, and Albert Fornells. A study of the effect of different types of noise on the precision of supervised learning techniques. *Artificial intelligence review*, 33(4):275–306, 2010. 1
- [32] Kento Nishi, Yi Ding, Alex Rich, and Tobias Hollerer. Augmentation strategies for learning with noisy labels. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8022–8031, 2021. 6
- [33] Diego Ortego, Eric Arazo, Paul Albert, Noel E. O’Connor, and Kevin McGuinness. Multi-objective interpolation training for robustness to label noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6606–6615, 2021. 2
- [34] 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, 2021. 6
- [35] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1944–1952, 2017. 1, 2, 6
- [36] Haim Permuter, Joseph Francos, and Ian Jermyn. A study of gaussian mixture models of color and texture features for image classification and segmentation. *Pattern recognition*, 39(4):695–706, 2006. 2, 8
- [37] Scott E. Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. Training deep neural networks on noisy labels with bootstrapping. In *ICLR*, 2015. 1, 2, 3, 4, 5
- [38] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning*, pages 4334–4343. PMLR, 2018. 1, 2, 3, 4, 5, 6, 8
- [39] Yanyao Shen and Sujay Sanghavi. Learning with bad training data via iterative trimmed loss minimization. In *International Conference on Machine Learning*, pages 5739–5748. PMLR, 2019. 2
- [40] 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 *Advances in Neural Information Processing Systems*, 2019. 1, 2, 4, 6
- [41] Jun Shu, Xiang Yuan, Deyu Meng, and Zongben Xu. Cmwnet: Learning a class-aware sample weighting mapping for robust deep learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2
- [42] Simon John Christoph Soerensen, Richard E Fan, Arun Seetharaman, Leo Chen, Wei Shao, Indrani Bhattacharya, Yong-hun Kim, Rewa Sood, Michael Borre, Benjamin I Chung, et al. Deep learning improves speed and accuracy of prostate gland segmentations on magnetic resonance imaging for targeted biopsy. *The Journal of Urology*, 206(3):604–612, 2021. 7
- [43] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. Selfie: Refurbishing unclean samples for robust deep learning. In *ICML*, 2019. 1
- [44] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5552–5560, 2018. 1, 2, 4, 6
- [45] Yuanpeng Tu, Boshen Zhang, Yuxi Li, Liang Liu, Jian Li, Yabiao Wang, Chengjie Wang, and Cai Rong Zhao. Learning from noisy labels with decoupled meta label purifier. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19934–19943, 2023. 2
- [46] Arash Vahdat. Toward robustness against label noise in training deep discriminative neural networks. In *NIPS*, 2017. 2
- [47] Andreas Veit, Neil Alldrin, Gal Chechik, Ivan Krasin, Abhinav Gupta, and Serge Belongie. Learning from noisy large-scale datasets with minimal supervision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 839–847, 2017. 2
- [48] Kaiping Wang, Bo Zhan, Chen Zu, Xi Wu, Jiliu Zhou, Luping Zhou, and Yan Wang. Triple-uncertainty guided mean teacher model for semi-supervised medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 450–460. Springer, 2021. 7
- [49] Xiaosong Wang, Ziyue Xu, Dong Yang, Leo Tam, Holger Roth, and Daguang Xu. Learning image labels on-the-fly for training robust classification models. *arXiv preprint arXiv:2009.10325*, 2020. 1
- [50] Hongxin Wei, Lei Feng, Xiangyu Chen, and Bo An. Combating noisy labels by agreement: A joint training method with co-regularization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [51] Qi Wei, Lei Feng, Haoliang Sun, Ren Wang, Chenhui Guo, and Yilong Yin. Fine-grained classification with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11651–11660, 2023. 6
- [52] Yichen Wu, Jun Shu, Qi Xie, Qian Zhao, and Deyu Meng. Learning to purify noisy labels via meta soft label corrector. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10388–10396, 2021. 2, 6
- [53] Zhi-Fan Wu, Tong Wei, Jianwen Jiang, Chaojie Mao, Mingqian Tang, and Yu-Feng Li. Ngc: A unified framework for learning with open-world noisy data. In *Proceedings*

- of the *IEEE/CVF International Conference on Computer Vision*, pages 62–71, 2021. 1
- [54] 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*, 2019. 1, 2
- [55] Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2691–2699, 2015. 2, 5, 6
- [56] Youjiang Xu, Linchao Zhu, Lu Jiang, and Yi Yang. Faster meta update strategy for noise-robust deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 144–153, 2021. 2, 6
- [57] Cheng Xue, Qi Dou, Xueying Shi, Hao Chen, and Pheng-Ann Heng. Robust learning at noisy labeled medical images: Applied to skin lesion classification. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 1280–1283. IEEE, 2019. 2, 4
- [58] Shuo Yang, Erkun Yang, Bo Han, Yang Liu, Min Xu, Gang Niu, and Tongliang Liu. Estimating instance-dependent label-noise transition matrix using dnns. 2021. 2
- [59] Yu Yao, Tongliang Liu, Bo Han, Mingming Gong, Jiankang Deng, Gang Niu, and Masashi Sugiyama. Dual t: Reducing estimation error for transition matrix in label-noise learning. In *Neurips*, 2020. 2
- [60] Yazhou Yao, Zeren Sun, Chuanyi Zhang, Fumin Shen, Qi Wu, Jian Zhang, and Zhenmin Tang. Jo-src: A contrastive approach for combating noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5192–5201, 2021. 1
- [61] Kun Yi and Jianxin Wu. Probabilistic end-to-end noise correction for learning with noisy labels. In *CVPR*, pages 7017–7025, 2019. 1, 2, 6
- [62] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *International Conference on Machine Learning*, pages 7164–7173. PMLR, 2019. 1, 2, 6
- [63] Boshen Zhang, Yuxi Li, Yuanpeng Tu, Jinlong Peng, Yabiao Wang, Cunlin Wu, Yang Xiao, and Cairong Zhao. Learning from noisy labels with coarse-to-fine sample credibility modeling. In *European Conference on Computer Vision*, pages 21–38. Springer, 2022. 2
- [64] Hui Zhang and Quanming Yao. Decoupling representation and classifier for noisy label learning. *arXiv preprint arXiv:2011.08145*, 2020. 6
- [65] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 7
- [66] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. 6
- [67] Weihe Zhang, Yali Wang, and Yu Qiao. Metacleaner: Learning to hallucinate clean representations for noisy-labeled visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7373–7382, 2019. 6
- [68] Zizhao Zhang, Han Zhang, Serkan O Arik, Honglak Lee, and Tomas Pfister. Distilling effective supervision from severe label noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9294–9303, 2020. 2, 6, 8
- [69] Evgenii Zheltonozhskii, Chaim Baskin, Avi Mendelson, Alex M Bronstein, and Or Litany. Contrast to divide: Self-supervised pre-training for learning with noisy labels. In *WACV*, pages 1657–1667, 2022. 2, 5, 6
- [70] Guoqing Zheng, Ahmed Hassan Awadallah, and Susan Dumais. Meta label correction for noisy label learning. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, 2021. 1, 2, 4, 6
- [71] Xiong Zhou, Xianming Liu, Chenyang Wang, Deming Zhai, Junjun Jiang, and Xiangyang Ji. Learning with noisy labels via sparse regularization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 72–81, 2021. 1
- [72] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation. In *Deep learning in medical image analysis and multimodal learning for clinical decision support*, pages 3–11. Springer, 2018. 7
- [73] Zhaowei Zhu, Tongliang Liu, and Yang Liu. A second-order approach to learning with instance-dependent label noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10113–10123, 2021. 1