

# Two Wrongs Don't Make a Right: Combating Confirmation Bias in Learning with Label Noise

Mingcai Chen, Hao Cheng, Yuntao Du, Ming Xu, Wenyu Jiang, Chongjun Wang\*

State Key Laboratory for Novel Software Technology at Nanjing University

Nanjing University, Nanjing 210023, China

{chenmc, chengh, duyuntao, lygjwy}@smail.nju.edu.cn, xuming0830@gmail.com, chjwang@nju.edu.cn

## Abstract

Noisy labels damage the performance of deep networks. For robust learning, a prominent two-stage pipeline alternates between eliminating possible incorrect labels and semi-supervised training. However, discarding part of noisy labels could result in a loss of information, especially when the corruption has a dependency on data, e.g., class-dependent or instance-dependent. Moreover, from the training dynamics of a representative two-stage method DivideMix, we identify the domination of confirmation bias: pseudo-labels fail to correct a considerable amount of noisy labels, and consequently, the errors accumulate. To sufficiently exploit information from noisy labels and mitigate wrong corrections, we propose Robust Label Refurbishment (Robust LR)—a new hybrid method that integrates pseudo-labeling and confidence estimation techniques to refurbish noisy labels. We show that our method successfully alleviates the damage of both label noise and confirmation bias. As a result, it achieves state-of-the-art performance across datasets and noise types, namely CIFAR under different levels of synthetic noise and Mini-WebVision and ANIMAL-10N with real-world noise.

## Introduction

Given certain capacity, deep networks have the capability of fitting arbitrary complex functions (Cybenko 1989). However, the randomization tests on common architectures (Edgington and Onghena 2007; Zhang et al. 2016; Arpit et al. 2017) show that they also easily fit training data with random labels. This phenomenon naturally raises the question of how deep learning continues to succeed in the presence of label noise.

Recently, the state-of-the-art two-stage methods have significantly improved noise robustness by incorporating Semi-Supervised Learning (SSL) (Ding et al. 2018; Nguyen et al. 2019; Li, Socher, and Hoi 2020; Zhou, Wang, and Bilmes 2021). The pipeline of a representative algorithm DivideMix (Li, Socher, and Hoi 2020) is shown in Fig. 1(a). In the first stage, problematic labels are identified and removed according to the per-example loss, i.e., the so-called “small-loss trick”. Therefore, the noisy dataset is divided into a labeled subset and an unlabeled subset. In the second stage, Di-

videMix calls an SSL algorithm named MixMatch (Berthelot et al. 2019), which minimizes the entropy of predictions on unlabeled examples through pseudo-labels. Such a pipeline leverages mislabeled data, improving the robustness to heavy and complex label noise.

However, we conclude that the two-stage pipeline suffers from two drawbacks. On the one hand, according to Vapnik’s principle (Vapnik 1998; Chapelle, Scholkopf, and Zien 2006),<sup>1</sup> discarding possible noisy labels to construct an SSL setting is inefficient. Specifically, some correct labels are wrongly filtered. What’s more, incorrect labels may also contain knowledge about the targets (Yu et al. 2018; Ishida et al. 2017; Kim et al. 2019; Berthon et al. 2021). For example, when an airplane image is mislabeled as a bird, the noisy label encodes the similarity information between the object of interest and the “bird” class. On the other hand, when introducing pseudo-labels during the SSL stage, confirmation bias (Tarvainen and Valpola 2017; Arazo et al. 2020) appears: Those confident but wrong predictions would be used to guide subsequent training, leading to a loop of self-reinforcing errors. *Label noise, together with confirmation bias, damage the performance.*

To observe the erroneous pseudo-labeling, we draw the training dynamics of a recent two-stage method DivideMix (Li, Socher, and Hoi 2020) on the corrupted training set of CIFAR-10 (Krizhevsky 2012) (under 90% synthetic symmetric noise). In every epoch, examples are grouped according to the relationship between their predicted labels, corrupted labels, and underlying ground-truth labels as in Fig. 1(b). The yellow color indicates the examples whose predicted labels agree with given noisy labels, i.e., III. predicted label = noisy label  $\neq$  ground-truth. The small yellow region at the bottom of Fig. 1(c) suggests that the model only agrees with a small fraction of noisy labels. It’s because DivideMix would filter possible wrong labels and avoid fitting them. On the other side, the red color indicates those predictions which fail to correct the noisy labels, i.e., IV. predicted label  $\neq$  noisy label and predicted label  $\neq$  ground-truth. From the red region at the top of Fig. 1(c), incorrect corrections comprise a large part throughout the training process. Considering the wrong pseudo-labels would be used for self-

\*Corresponding authors

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>When solving a problem of interest, do not solve a more general problem as an intermediate step.

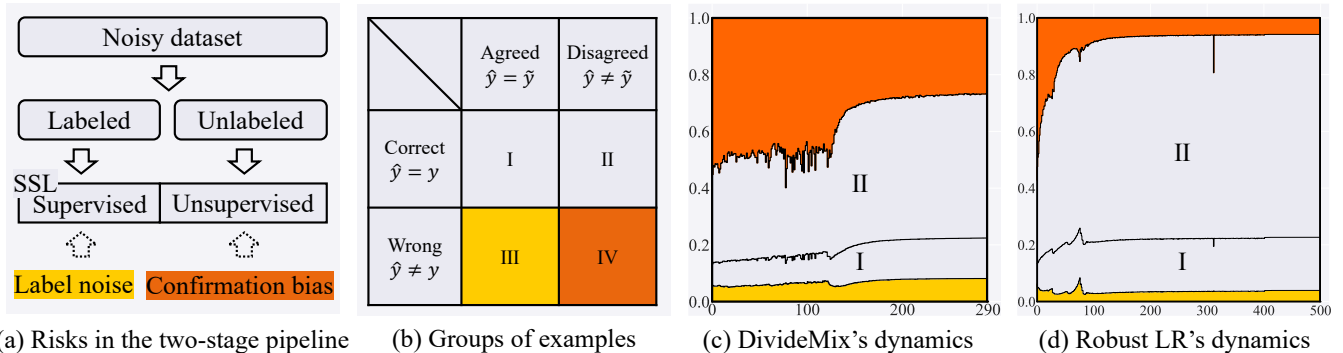


Figure 1: Two-stage pipeline fails to correct a large proportion of wrong labels, evidenced by the training dynamics. Underlying ground-truth label, noisy label, and predicted label are denoted as  $y$ ,  $\tilde{y}$ ,  $\hat{y}$  respectively. In every epoch, the examples are divided into four groups as shown in (b): I. The predicted labels agree with the clean labels. II. The predicted labels correct the noisy labels. III. The predicted labels agree with the noisy labels. IV. The predicted labels fail to correct the given labels. In (c) and (d), the x-axis denotes the epoch, and the y-axis denotes the proportion of different groups. Best viewed in color.

training, it causes the confirmation bias problem, affecting performance adversely.

Our work begins by suggesting that better robustness can be achieved by sufficiently exploiting the information in the noisy labels and mitigating the side-effect of SSL. We observe one of the recent two-stage methods as Fig. 1(c): The pseudo-labels dominate the given noisy labels during training. We propose a hybrid method named Robust LR to address the problem. It estimates the label confidence by modeling the per-example loss and then accordingly refurbishes noisy labels through a dynamic convex combination with pseudo-labels. Robust LR improves upon the two-stage pipeline by leveraging all noisy labels and constructing target labels in a more fine-grained manner. To further alleviate confirmation bias: 1). Two models are trained simultaneously, where each model interacts with its peer through pseudo-labeling and confidence estimation. 2). Different augmentation strategies are deployed for loss modeling and learning following recent findings (Chen et al. 2020b; Nishi et al. 2021). For comparison, we draw Fig. 1(d) using our method under the same setting. Compared with Fig. 1(c), the red region, which indicates wrong corrections, are much smaller. It shows that our approach alleviates the damage of wrong pseudo-labels while combating label noise. To sum up, we highlight the contributions of this paper as follows:

- We analyze the inefficiency of the two-stage pipeline and suggest that there is a loss of information when transforming the label noise problem into SSL. Moreover, the visualization of the training dynamics helps us identify the domination of confirmation bias (see Fig. 1(c)).
- To address this, we propose a hybrid method named Robust LR. By integrating pseudo-labeling and confidence estimation techniques into label refurbishment, it successfully leverages all noisy labels and alleviates the damage of both label noise and confirmation bias.
- We experimentally show that our method advances state-of-the-art results on CIFAR with synthetic label noise, as well as the real-world noisy dataset Mini-WebVision

and ANIMAL-10N. Besides, we systematically study the components of Robust LR to examine their impacts.

## Related Work

The label noise is ubiquitous in real-world data. When the noise rate is insignificant, it can be implicitly dealt with. For example, the noise labels in MNIST, CIFAR, and ImageNet (some of them are reported in <https://labelerrors.com/>), are usually neglected. Regularization techniques, including Dropout (Srivastava et al. 2014), weight decay (Krogh and Hertz 1992), and the inherent robustness in deep networks (Zhang et al. 2016) combat label noise.

The damage of noisy labels gradually appears as noise becomes non-negligible. Some methods assume a class-dependent (or instance-independent) label noise, i.e., the distribution of noisy labels only dependent on the ground-truth label:

$$p(\tilde{y} = j | y = i, X = x) = p(\tilde{y} = j | y = i) \quad (1)$$

The corruption process thus can be modeled by a label transition matrix  $T \in [0, 1]^{C \times C}$  where  $T_{ij} := p(\tilde{y} = j | y = i)$  and  $C$  is the number of classes. Webly learning (Chen and Gupta 2015) adds an extra noise adaptation layer on top of the base model to mimic the transition behavior. The base model is first trained on easy examples, and then the entire model is trained on the noisy dataset. Backward correction (Patrini et al. 2017) estimates the label transition through the outputs of a network trained on the noisy dataset. Then it trains another network with weighted loss, where the weights are from the estimated label transition matrix. Forward correction (Patrini et al. 2017) does the same to obtain the matrix. But it instead corrects the outputs during forward pass when trains a new network. To better estimate the transition matrix, Dual T (Yao et al. 2020) factorizes it into two easy-to-estimate matrices. The effectiveness of these approaches depends on whether the transition matrix is accurate. Besides, the noise type could be more complex in real-world, e.g., instance-dependent:

$$p(\tilde{y} = j | y = i, X = x) = T_{i,j}(x)p(y = j | y = i) \quad (2)$$

Method	Label Refurbishment	Two-Stage	Robust LR
Fully explore	inputs	☺	☺
	labels	☺	☺
Complex & Heavy Noise	☹	☹	☺

Table 1: Comparison of training schemes.

where  $T_{i,j}(x)$  is the instance-dependent noise model. The aforementioned methods have difficulty in modeling such complex noise.

A large part of the methods achieves robustness by relying on the internal noise tolerance of deep networks. They mainly differ in the example selection, loss weighting, or label refurbishment strategies (Frénay and Verleysen 2013; Algan and Ulusoy 2021; Song, Kim, and Lee 2019). Bootstrapping (Reed et al. 2014) uses the interpolation of labels and model predictions for training. Decouple (Malach and Shalev-Shwartz 2017) updates two predictors with only disagreed examples. Activate bias (Chang, Learned-Miller, and McCallum 2017) emphasizes high variance examples. MentorNet (Jiang et al. 2018) weights examples using a pre-trained teacher network. Co-teaching (Han et al. 2018) maintains two models where one selects examples with small losses to update another. Based on Co-teaching, Co-teaching+ (Yu et al. 2019) prevents two models from converging to a consensus by only considering disagreed examples. D2L (Ma et al. 2018) adopts a measure called local intrinsic dimensionality. Labels are refurbished to prevent the increase of intrinsic dimension. SELFIE (Song, Kim, and Lee 2019) only considers examples with consistent predictions for refurbishment. TopoFilter (Wu et al. 2020) adopts a different selection criteria by exploring the latent representational space. Self-adaptive training (Huang, Zhang, and Zhang 2020) uses the exponential moving average of predictions as pseudo-labels. SEAL (Chen et al. 2020a) retrain a model with the average predictions of a teacher model. However, these methods may suffer from big performance drops under heavy noise due to inaccurate correction, weighting, or refurbishment.

Recently, the two-stage pipeline has gained much attention. SELF (Nguyen et al. 2019) first uses the ensemble of predictions to filter problematic labels. In the second stage, it performs an SSL method named Mean Teacher (Tarvainen and Valpola 2017). DivideMix (Li, Socher, and Hoi 2020) uses the Gaussian Mixture Model (GMM) to separate examples with small and big losses, and they are treated as clean and noisy examples, respectively. Then the SSL method MixMatch (Berthelot et al. 2019) is used to leverage the feature information. RoCL (Zhou, Wang, and Bilmes 2021) selects clean examples according to the consistency of the loss and output, followed by a self-training method. This type of method utilizes SSL to leverage mislabeled examples. However, we suggest that they fail to exploit all noisy labels and suffer from wrong corrections. The proposed method Robust LR leverages possible noisy labels. It preserves label information in a soft manner by adopting successful ideas from the two-stage pipeline and SSL into the classic label

refurbishment process, as shown in Table 1. Furthermore, Robust LR is dedicated to alleviating confirmation bias. Different augmentation strategies and co-training are combined to form a hybrid method.

## Method

### Overview of Robust LR

Robust LR refurbishes the noisy labels before training. To reduce the marginalized effect of wrong labels, the refurbished label  $y^* \in \Delta^{C-1}$  (where  $\Delta^{C-1}$  is the probability simplex) comes from a dynamic convex combination of the noisy label  $\tilde{y}$  (one-hot label over  $C$  classes) and the soft pseudo-label  $\hat{y}$  (predicted probability distribution over  $C$  classes).

$$y^* = w\tilde{y} + (1 - w)\hat{y} \quad (3)$$

The pseudo-label  $\hat{y}$  is obtained from the models’ prediction. The weight  $w$ , i.e., the clean probability, is estimated using a two-component GMM fitted on the per-example loss. To further alleviate confirmation bias, two models are simultaneously trained, where one model contributes to another’s confidence estimation and pseudo-labeling process. They have the same structure but different parameters  $\theta^{(0)}, \theta^{(1)}$ . The overall pipeline of Robust LR is shown in Fig. 2 and Algorithm 1. In every training round, the confidence estimation and pseudo-labeling are performed first. Then the model is trained with the refurbished labels.

### Warm-Up

As shown in (Arpit et al. 2017), deep models tend to fit clean examples first. Therefore, Robust LR warms two models up by shortly training them on the noisy dataset. The commonly used mini-batch gradient descent algorithm is performed to update the parameters. For illustration, we denote this process as Train(dataset, parameters, number of iterations). Thus, the warm-up process is:

$$\text{Train}(\tilde{\mathcal{D}}, \theta^{(m)}, I_{warm}) \quad \text{for } m = 0, 1 \quad (4)$$

where  $I_{warm}$  is a small number of iterations so that the training ends before models fitting too many noisy labels.

### Main Training Round

**Confidence estimation** It has been shown that models are prone to present smaller losses on clean examples (Arpit et al. 2017; Chen et al. 2019; Han et al. 2018; Li, Socher, and Hoi 2020). Therefore, Robust LR estimates the label confidence based on the loss value. Specifically, the per-example cross-entropy loss  $H$  between the noisy label and the prediction is first calculated,

$$\ell_i = H(\tilde{y}_i, p(y | x_i; \theta^{(1-m)})) \quad (5)$$

Then a two-component one-dimensional GMM is used to model the distribution of per-example loss,

$$\mathcal{W} = \text{GMM}(\{(\ell_i)\}_{i=1}^N) \quad (6)$$

where  $\mathcal{W} = \{w_i\}_{i=1}^N$  is the label confidence which equals to the probability of each loss value belonging to the GMM component with a smaller mean. The parameters of GMM

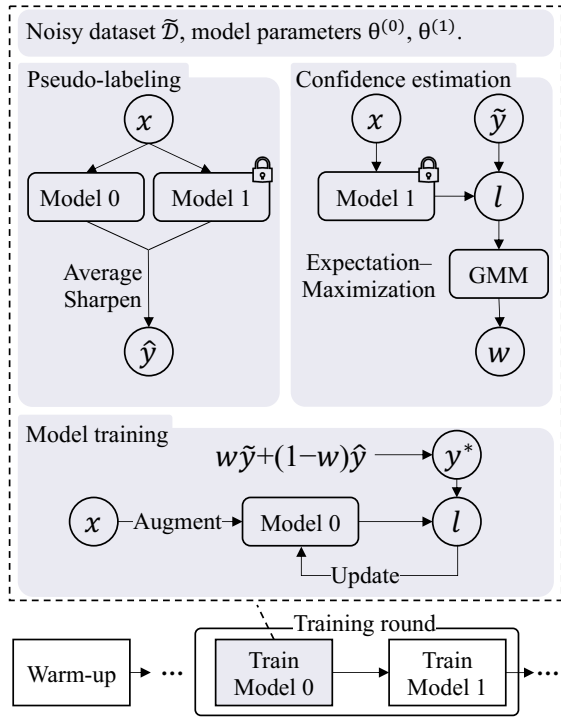


Figure 2: Pipeline of Robust LR. The visualization of fitted GMM can be found in the appendix.

are determined using the expectation-maximization algorithm. The procedure follows the standard practice, so we don't elaborate on the details here. Note that, to alleviate confirmation bias, the label confidence for the current model  $m$  comes from the predictions of another model  $1 - m$ .

**Pseudo-labeling** To correct the noisy labels with accurate pseudo-labels, two models' predictions are averaged and then sharpened,

$$\hat{y}_i = \text{Sharpen}\left(\frac{\text{p}(y | x_i; \theta^{(m)}) + \text{p}(y | x_i; \theta^{(1-m)})}{2}\right) \quad (7)$$

where the sharpening function scales the categorical distribution with a hyper-parameter  $T$ ,

$$\text{Sharpen}(p)_i = \frac{p_i^{\frac{1}{T}}}{\sum_{j=1}^C p_j^{\frac{1}{T}}} \quad (8)$$

where  $C$  is the number of classes.  $p_i$  is the probability of  $i$ -th class of input distribution  $p$ .

**Model training** After label refurbishment using the estimated confidence and pseudo-labels according to Equation 3, current model  $m$  is trained with the refurbished labels for  $I$  iterations,

$$\text{Train}(\{(\text{Aug}(x_i), y_i^*)\}_{i=1}^N, \theta^{(m)}, I) \quad (9)$$

where  $\text{Aug}(\cdot)$  is the data augmentation function introduced in the next section. The cross-entropy between the soft labels and predictions is used as loss function here. After the

### Algorithm 1: Robust LR

**Input:** Noisy dataset  $\tilde{\mathcal{D}} = \{(x_i, \tilde{y}_i)\}_{i=1}^N$ , # iterations for warm up  $I_{warm}$ , # iterations in main training round  $I$ , # training rounds  $R$ , training strategy  $\text{Train}(\text{dataset}, \text{parameters}, \# \text{ iterations})$ .

**Output:** model's parameters  $\theta^{(0)}, \theta^{(1)}$

```

1: Randomly initialize  $\theta^{(0)}, \theta^{(1)}$ 
2:  $\text{Train}(\tilde{\mathcal{D}}, \theta^{(m)}, I_{warm})$  for  $m = 0, 1$   $\triangleright$  warm up
3: for  $r = 1$  to  $R$  do
4:   for  $m = 0$  to  $1$  do  $\triangleright$  train two models separately
5:     for  $i = 0$  to  $N$  do
6:        $\ell_i = \text{H}(y_i, \text{p}(y | x_i; \theta^{(1-m)}))$ 
7:        $\triangleright$  obtain per-example loss
8:     end for
9:      $W = \text{GMM}(\{\{\ell_i\}_{i=1}^N\})$   $\triangleright$  fit GMM
10:    for  $i = 0$  to  $N$  do
11:       $\hat{y}_i = \text{Sharpen}\left(\frac{\text{p}(y|x_i; \theta^{(m)}) + \text{p}(y|x_i; \theta^{(1-m)})}{2}\right)$ 
12:       $\triangleright$  pseudo-label
13:       $y_i^* = w_i \tilde{y}_i + (1 - w_i) \hat{y}_i$   $\triangleright$  refurbish
14:    end for
15:     $\text{Train}(\{(\text{Aug}(x_i), y_i^*)\}_{i=1}^N, \theta^{(m)}, I)$ 
16:  end for
17: end for

```

training of model  $m$ , another model  $1 - m$  is trained similarly. This process proceeds until reaching a fixed number of training rounds.

During implementation, a regularization loss term is used as in (Tanaka et al. 2018; Arazo et al. 2019; Li, Socher, and Hoi 2020). It encourages the network to output uniform distribution across examples in the mini-batch.

$$L_{reg} = \sum_c \pi_c \log\left(\frac{\pi_c}{\bar{p}_c}\right) \quad (10)$$

$$\bar{p}_c = \frac{1}{B} \sum_{i=1}^B \text{p}(y = c | x_i; \theta)$$

where  $\pi$  is the uniform prior distribution, we set  $\pi_c = \frac{1}{C}$ .

For asymmetric noise, we add a negative entropy loss term during warm-up following (Pereyra et al. 2017; Li, Socher, and Hoi 2020).

$$\sum_c \text{p}(y | x; \theta) \log(\text{p}(y | x; \theta)) \quad (11)$$

### The Different Augmentation Strategies

Due to the lack of accurate supervised information, improving the generalization ability is the core task of learning with label noise. Data augmentation is a common technique that approaches such a problem via applying stochastic transformation on images.

In Robust LR, forward pass serves three purposes: loss modeling, pseudo-labeling, and learning. We use basic image augmentation for loss modeling and pseudo-labeling but stronger augmentations for learning. This design is based on two recent findings: 1). In learning with label noise, using

Dataset		CIFAR-10					CIFAR-100			
Noise type		Sym.			Asym.	Sym.				
Method/Noise ratio		20%	50%	80%	90%	40%	20%	50%	80%	90%
F-correction (Patrini et al. 2017)	Best	86.8	79.8	63.3	42.9	87.2	61.5	46.6	19.9	10.2
	Last	83.1	59.4	26.2	18.8	83.1	61.4	37.3	9.0	3.4
Co-teaching+ (Yu et al. 2019)	Best	89.5	85.7	67.4	47.9	-	65.6	51.8	27.9	13.7
	Last	88.2	84.1	45.5	30.1	-	64.1	45.3	15.5	8.8
P-correction (Yi and Wu 2019)	Best	92.4	89.1	77.5	58.9	88.5	69.4	57.5	31.1	15.3
	Last	92.0	88.7	76.5	58.2	88.1	68.1	56.4	20.7	8.8
Meta-Learning (Li et al. 2019)	Best	92.9	89.3	77.4	58.7	89.2	68.5	59.2	42.4	19.5
	Last	92.0	88.8	76.1	58.3	88.6	67.7	58.0	40.1	14.3
M-correction (Arazo et al. 2019)	Best	94.0	92.0	86.8	69.1	87.4	73.9	66.1	48.2	24.3
	Last	93.8	91.9	86.6	68.7	86.3	73.4	65.4	47.6	20.5
DivideMix (Li, Socher, and Hoi 2020)	Best	96.1	94.6	93.2	76.0	93.4	77.3	74.6	60.2	31.5
	Last	95.7	94.4	92.9	75.4	92.1	76.9	74.2	59.6	31.0
AugDesc* (Nishi et al. 2021)	Best	96.1	-	-	89.6	-	78.1	-	-	36.8
	Last	96.0	-	-	89.4	-	77.8	-	-	36.7
Ours	Best	<b>96.5</b>	<b>95.8</b>	<b>94.3</b>	<b>92.8</b>	<b>94.4</b>	<b>79.1</b>	<b>75.3</b>	<b>66.7</b>	<b>37.5</b>
	Last	<b>96.4</b>	<b>95.7</b>	<b>94.2</b>	<b>92.8</b>	<b>93.7</b>	<b>78.6</b>	<b>74.6</b>	<b>66.2</b>	<b>37.3</b>

Table 2: Comparison with state-of-the-art methods on CIFAR10 and CIFAR-100 with synthetic noise. Sym. and Asym. are symmetric and asymmetric for short, respectively. The results of other methods are from (Li, Socher, and Hoi 2020). The best results are indicated in bold. \*AugDesc uses the same augmentation technique (RandAugment) as our method.

different augmentations for loss modeling and learning is more effective (Nishi et al. 2021). 2). Unsupervised learning benefits from stronger data augmentation (Chen et al. 2020b), and we find the same preference can also be extended to this problem.

In particular, the basic image augmentation for loss modeling and pseudo-labeling consists of random crop and random horizontal flip. The strong transformation Aug( $\cdot$ ) consists of RandAugment (Cubuk et al. 2020) and Cutout (DeVries and Taylor 2017). RandAugment first randomly selects a given number of operations from a pre-defined set of transformations. The set consists of geometric and photometric transformations, such as affine transformation and color adjustment. In the next, these operations are applied with given magnitudes. Cutout randomly masks out square regions of images. These augmentations are sequentially applied to the input images. The settings of RandAugment are reported in the supplementary material.

## Experiment

### Comparison with State-of-the-Art Methods

We benchmark the proposed method on experimental settings using CIFAR-10, CIFAR-100 (Krizhevsky 2012) with different levels of synthetic noises, as well as the real-world noisy dataset Mini-WebVision (Li et al. 2017), ANIMAL-10N (Song, Kim, and Lee 2019).

**Synthetic label noise on CIFAR-10, CIFAR-100** Following previous methods (Kim et al. 2019; Li, Socher, and Hoi

2020), two types of synthetic noises are experimented: symmetric and asymmetric noise. Symmetric noise is generated by assigning examples to random classes with the same probability. The noise rate ranges from 20% to 90% (note that the noise labels are randomly distributed throughout  $C$  classes, and the true labels may be maintained after corruption). Asymmetric noise is generated by randomly corrupting labels according to a pre-defined transition matrix. Examples would only be corrupted to similar classes, such as deer to horse. 40% asymmetric noise is experimented (50% being indistinguishable).

We report the average performance of Robust LR over 3 trials with different random seeds for generating noise and parameters initialization. The backbone structure is PreAct Resnet (He et al. 2016). The training details are reported in the supplementary material. Following previous work, the best test accuracy across all epochs and the averaged test accuracy over the last 10 epochs are both reported. A validation set with 5,000 examples is drawn from the noisy training set for hyper-parameters tuning. We find that two main hyper-parameters in Robust LR, namely temperature value and the weight for regularization term (Tanaka et al. 2018; Arazo et al. 2019), don't need to be heavily tuned. Specifically, there are only two sets of hyper-parameters for light and heavy noise, respectively. For light noise, namely CIFAR-10 under 20% to 80% symmetric noise, 40% asymmetric noise, and CIFAR-100 under 20% symmetric noise, the temperature is 1, and the coefficient for the regularization term is 2. For heavy noise, namely CIFAR-10 under

Method	Mini-WebVision		ILSVRC12	
	top-1	top-5	top-1	top-5
F-correction	61.12	82.68	57.36	82.36
Decoupling	62.54	84.74	58.26	82.26
D2L	62.68	84.00	57.80	81.36
MentorNet	63.00	81.40	57.80	79.92
Co-teaching	63.58	85.20	61.48	84.70
Iterative-CV	65.24	85.34	61.60	84.98
DivideMix	77.32	91.64	75.20	90.84
Robust LR	<b>81.84</b>	<b>94.12</b>	<b>75.48</b>	<b>93.76</b>

Table 3: Comparison with other methods on Mini-WebVision. The results of other methods are from (Li, Socher, and Hoi 2020).

SELFIE	PLC	NCT	Robust LR
81.8	83.4	84.1	<b>88.5</b>

Table 4: Comparison with other methods on ANIMAL-10N. The results of other methods are from (Chen et al. 2021).

90% symmetric noise and CIFAR-100 under 50% to 90% symmetric noise, the temperature is  $1/3$ , and the coefficient for regularization term is 10.

As shown in Table 2, our method consistently outperforms previous best results on all the settings. The improvement is substantial, especially when the noise is heavy. For example, Robust LR obtains 92.8% accuracy on CIFAR-10 under 90% noise, surpassing the previous best by more than 3%. We remark that previous methods underperform under heavy noise because they fail to avoid confirmation bias. It’s worth noting that Robust LR outperforms AugDesc even with the same augmentation. It shows that our improvement also comes from other components.

The comparison between Robust LR and other methods is shown in Table 2. Our method outperforms the previous best method by over 1%. The distribution of asymmetric noise in the corrupted training set is shown in the appendix. Robust LR resists the mimicked class-dependent noise and correctly predicts most of them.

**Real-world label noise on Mini-WebVision and ANIMAL-10N** To verify the effectiveness of our approach on the real-world large-scale noisy dataset, we then conduct experiments on Mini-WebVision and ANIMAL-10N. WebVision is crawled from Flickr and Google using the same 1,000 classes as the ImageNet ILSVRC12 dataset for querying. The estimated noise rate is 20%. Following the setting of previous work (Chen et al. 2019; Li, Socher, and Hoi 2020): The first 50 classes of the ImageNet ILSVRC12 dataset are compared, and its validation set is used. In terms of the hyper-parameters, the temperature is 3, and the coefficient for the regularization term is 1. ANIMAL-10N consists of 50000 train animal images and 10000 test animal images in 10 classes, with an 8% estimated error rate. The temperature is 1, and the coefficient for the regularization term is 2.

For comparison, results of F-correction (Patrini et al. 2017), Decoupling (Malach and Shalev-Shwartz 2017), D2L (Ma et al. 2018), MentorNet (Jiang et al. 2018), Co-teaching (Han et al. 2018), Iterative-CV (Chen et al. 2019), DivideMix (Li, Socher, and Hoi 2020), SELFIE (Song, Kim, and Lee 2019), PLC (Zhang et al. 2021), NCT (Chen et al. 2021) are reported.

As shown in Table 3, Robust LR improves the performance by a considerable margin, namely, 4.5% top-1 accuracy against the previous best on the test set of Mini-WebVision and 4.4% on ANIMAL-10N. The results verify that our method can cope with complex real-world noise.

## Ablation Study

We further study the components of Robust LR. Specifically, we analyze the results of:

1. To study the effect of label refurbishment, we remove label refurbishment and directly use either given noisy labels or pseudo-labels. When the probability of being clean is larger than 0.5, the noisy label is used. Otherwise, the pseudo-label is used.
2. To study the effect of strong augmentation, we replace it with basic transformation.
3. To study the effect of GMM for dynamic confidence estimation, we replace it with 0.5 fixed confidence.
4. To study the effect of co-training, we only use one model.

The results on CIFAR-10 with four levels of symmetry noise are reported. From Table 5, other training schemes suffer from different degrees of performance drops. This verifies that the incorporation of the components in Robust LR is effective. In the next, we analyze each component.

**Label refurbishment** The label refurbishment alleviates the marginalized effect of wrong labels and thus, contributes to the final performance. Under light noise, i.e., when the noise is insignificant or can be corrected easily, the gain is limited. Under heavy noise, e.g., 90% noise rate, the model is much more sensitive to its absence.

To understand the large gap between the best and last performance (73.8% vs. 23.9%) under heavy noise, we further observe models’ behaviors. We find that the training is unstable under heavy noise, e.g., the GMM may not converge in some rounds and assigns more than 95% of examples with bigger clean probabilities. The bad confidence estimation would affect later training in return. The training can be stabilized after further tuning the hyper-parameters, such as the learning rate. For consistency, we only report the performance under the same hyper-parameters.

**Data augmentation** Replacing the strong augmentation makes the model fails to converge. We remark that it’s because Robust LR is a holistic method. Strong data augmentation not only serves the common purpose of regularization (Shorten and Khoshgoftaar 2019), but also is part of the different augmentation strategies (Nishi et al. 2021).

One may still argue that the augmentation is more important than other components. We show that other components all improve upon the Robust LR with strong augmentation



Figure 3: Some mislabeled or indistinguishable examples in the training set of CIFAR-10 found by Robust LR. The wrong annotations, the predicted classes (in the parentheses), and the IDs of images are shown. The airp. and auto. are airplane and automobile for short.

Method/Noise ratio		20%	50%	80%	90%
Robust LR	Best	96.5	95.8	94.5	92.8
	Last	96.4	95.7	94.2	92.8
1. w/o LR	Best	96.3	95.8	94.5	73.8
	Last	96.2	95.6	94.1	23.9
2. w/o strong aug.	Best	92.6	88.1	65.3	48.7
	Last	92.5	72.7	36.5	24.3
3. w/o GMM	Best	94.6	91.4	88.0	87.6
	Last	92.7	80.3	43.4	31.1
4. w/o co-training	Best	96.4	95.7	94.3	82.8
	Last	95.6	94.4	93.1	79.9

Table 5: Ablation study. Results on CIFAR-10 with different levels of symmetry noise are reported.

in Table 5. Besides, as shown in Table 2, our method outperforms AugDesc, a method with the same augmentation Robust LR uses.

**GMM** The GMM is also essential, and removing the dynamic confidence estimation damages the performance. We also notice that, for four levels of corruption, GMM assigns 18%, 44%, 70%, 78% examples bigger noisy probability ( $w < 0.5$ ) at the end of training, respectively. It is an accurate estimation of the real noise rate (for 20%, 50%, 80%, 90% noise rate, there is actually 18%, 45%, 72%, 81% noisy labels). For Mini-WebVision, the GMM assigns 18% examples bigger noisy probability in the end, which is also approximate to the reported noise rate 20% (Li et al. 2017). We envision this could be used to estimate the noise rate in real-world datasets.

**Co-training** Removing co-training leads to considerable drops in performance. It is also noteworthy that our single model’s performance already surpasses previous co-training methods, such as DivideMix or Co-teaching. We suggest that co-training alleviates confirmation bias, and the ensemble of two models also produces better self-training signals.

### Finding the Noisy Labels in CIFAR-10

Apart from combating label noise, Robust LR can be directly used to find the noisy labels in the training set. Standard empirical risk minimization would easily fit the training set with only a small amount of noisy labels. Instead, Robust LR could avoid the fitting on the possible noisy labels. We use CIFAR-10 to illustrate how we can use Robust LR to find noisy labels in a mostly correctly labeled dataset.

We first train Robust LR on the CIFAR-10 training set (without corruption) for 100 epochs without modifying the algorithm. In the next, examples with top-50 big losses are selected and hand-picked. We successfully find some mislabeled or indistinguishable examples in the training set as in Figure 3 (note that there is no ground-truth or high-resolution originals, we can only subjectively tell whether the noisy labels are right or wrong). Some of them are mislabeled probably because of the similarity between two classes, such as 25095 (bird vs. airplane), 38775 (dog vs. cat). Some classes don’t usually consider similar, but images in these classes can still be ambiguous, e.g., image 36782 (dog vs. bird) and 33079 (dog vs. horse). These verify that the noise we are facing in the real world could be complex.

### Conclusion

In this paper, we study the problem of learning with label noise. We analyze the drawbacks of the two-stage pipeline and identify its confirmation bias problem by visualizing the training dynamics. The observation motivates us to propose Robust LR, a new training algorithm that dynamically refurbishes labels using confidence estimation and pseudo-labeling techniques. We demonstrate that our approach combats both confirmation bias and label noise. As a result, it significantly advances the state-of-the-art. We then conduct ablation experiments to study the effects of the components. Finally, we attempt to find the mislabeled examples in CIFAR-10 with Robust LR. In future work, we are interested in further incorporating ideas from weakly supervised learning into hybrid methods and continuing to combat complex label noise.

## Acknowledgements

This paper is supported by the National Natural Science Foundation of China (Grant No. 62192783, U1811462), the Collaborative Innovation Center of Novel Software Technology and Industrialization at Nanjing University.

## References

- Algan, G.; and Ulusoy, I. 2021. Image classification with deep learning in the presence of noisy labels: A survey. *Knowledge-Based Systems*, 215: 106771.
- Arazo, E.; Ortego, D.; Albert, P.; O'Connor, N.; and McGuinness, K. 2019. Unsupervised label noise modeling and loss correction. In *International Conference on Machine Learning*, 312–321. PMLR.
- Arazo, E.; Ortego, D.; Albert, P.; O'Connor, N. E.; and McGuinness, K. 2020. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Arpit, D.; Jastrzebski, S.; Ballas, N.; Krueger, D.; Bengio, E.; Kanwal, M. S.; Maharaj, T.; Fischer, A.; Courville, A. C.; Bengio, Y.; and Lacoste-Julien, S. 2017. A Closer Look at Memorization in Deep Networks. In Precup, D.; and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, 233–242. PMLR.
- Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; and Raffel, C. 2019. Mixmatch: A holistic approach to semi-supervised learning. *arXiv preprint arXiv:1905.02249*.
- Berthon, A.; Han, B.; Niu, G.; Liu, T.; and Sugiyama, M. 2021. Confidence scores make instance-dependent label-noise learning possible. In *International Conference on Machine Learning*, 825–836. PMLR.
- Chang, H.-S.; Learned-Miller, E.; and McCallum, A. 2017. Active bias: Training more accurate neural networks by emphasizing high variance samples. *arXiv preprint arXiv:1704.07433*.
- Chapelle, O.; Scholkopf, B.; and Zien, A., eds. 2006. *Semi-Supervised Learning*. The MIT Press. ISBN 9780262033589.
- Chen, P.; Liao, B. B.; Chen, G.; and Zhang, S. 2019. Understanding and utilizing deep neural networks trained with noisy labels. In *International Conference on Machine Learning*, 1062–1070. PMLR.
- Chen, P.; Ye, J.; Chen, G.; Zhao, J.; and Heng, P.-A. 2020a. Beyond class-conditional assumption: A primary attempt to combat instance-dependent label noise. *arXiv preprint arXiv:2012.05458*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020b. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Chen, X.; and Gupta, A. 2015. Webly supervised learning of convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 1431–1439.
- Chen, Y.; Shen, X.; Hu, S. X.; and Suykens, J. A. K. 2021. Boosting Co-Teaching With Compression Regularization for Label Noise. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2021, virtual, June 19-25, 2021*, 2688–2692. Computer Vision Foundation / IEEE.
- Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. V. 2020. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 702–703.
- Cybenko, G. 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4): 303–314.
- DeVries, T.; and Taylor, G. W. 2017. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*.
- Ding, Y.; Wang, L.; Fan, D.; and Gong, B. 2018. A semi-supervised two-stage approach to learning from noisy labels. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 1215–1224. IEEE.
- Edgington, E.; and Onghena, P. 2007. *Randomization tests*. CRC press.
- Fréney, B.; and Verleysen, M. 2013. Classification in the presence of label noise: a survey. *IEEE transactions on neural networks and learning systems*, 25(5): 845–869.
- Han, B.; Yao, Q.; Yu, X.; Niu, G.; Xu, M.; Hu, W.; Tsang, I.; and Sugiyama, M. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *arXiv preprint arXiv:1804.06872*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Huang, L.; Zhang, C.; and Zhang, H. 2020. Self-adaptive training: beyond empirical risk minimization. *Advances in Neural Information Processing Systems*, 33.
- Ishida, T.; Niu, G.; Hu, W.; and Sugiyama, M. 2017. Learning from complementary labels. *arXiv preprint arXiv:1705.07541*.
- Jiang, L.; Zhou, Z.; Leung, T.; Li, L.-J.; and Fei-Fei, L. 2018. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International Conference on Machine Learning*, 2304–2313. PMLR.
- Kim, Y.; Yim, J.; Yun, J.; and Kim, J. 2019. Nlnl: Negative learning for noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 101–110.
- Krizhevsky, A. 2012. Learning Multiple Layers of Features from Tiny Images. *University of Toronto*.
- Krogh, A.; and Hertz, J. A. 1992. A simple weight decay can improve generalization. In *Advances in neural information processing systems*, 950–957.
- Li, J.; Socher, R.; and Hoi, S. C. 2020. Dividemix: Learning with noisy labels as semi-supervised learning. *arXiv preprint arXiv:2002.07394*.



- Li, J.; Wong, Y.; Zhao, Q.; and Kankanhalli, M. S. 2019. Learning to learn from noisy labeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5051–5059.
- Li, W.; Wang, L.; Li, W.; Agustsson, E.; and Van Gool, L. 2017. Webvision database: Visual learning and understanding from web data. *arXiv preprint arXiv:1708.02862*.
- Ma, X.; Wang, Y.; Houle, M. E.; Zhou, S.; Erfani, S.; Xia, S.; Wijewickrema, S.; and Bailey, J. 2018. Dimensionality-driven learning with noisy labels. In *International Conference on Machine Learning*, 3355–3364. PMLR.
- Malach, E.; and Shalev-Shwartz, S. 2017. “Decoupling” when to update” from” how to update”. *arXiv preprint arXiv:1706.02613*.
- Nguyen, D. T.; Mummadi, C. K.; Ngo, T. P. N.; Nguyen, T. H. P.; Beggel, L.; and Brox, T. 2019. Self: Learning to filter noisy labels with self-ensembling. *arXiv preprint arXiv:1910.01842*.
- Nishi, K.; Ding, Y.; Rich, A.; and Hollerer, T. 2021. Augmentation strategies for learning with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8022–8031.
- Patrini, G.; Rozza, A.; Krishna Menon, A.; Nock, R.; and Qu, L. 2017. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1944–1952.
- Pereyra, G.; Tucker, G.; Chorowski, J.; Kaiser, Ł.; and Hinton, G. 2017. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548*.
- Reed, S.; Lee, H.; Anguelov, D.; Szegedy, C.; Erhan, D.; and Rabinovich, A. 2014. Training deep neural networks on noisy labels with bootstrapping. *arXiv preprint arXiv:1412.6596*.
- Shorten, C.; and Khoshgoftaar, T. M. 2019. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1): 1–48.
- Song, H.; Kim, M.; and Lee, J. 2019. SELFIE: Refurbishing Unclean Samples for Robust Deep Learning. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, 5907–5915. PMLR.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1): 1929–1958.
- Tanaka, D.; Ikami, D.; Yamasaki, T.; and Aizawa, K. 2018. Joint optimization framework for learning with noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5552–5560.
- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv preprint arXiv:1703.01780*.
- Vapnik, V. N. 1998. *Statistical Learning Theory*. Wiley-Interscience.
- Wu, P.; Zheng, S.; Goswami, M.; Metaxas, D.; and Chen, C. 2020. A topological filter for learning with label noise. *arXiv preprint arXiv:2012.04835*.
- Yao, Y.; Liu, T.; Han, B.; Gong, M.; Deng, J.; Niu, G.; and Sugiyama, M. 2020. Dual T: Reducing estimation error for transition matrix in label-noise learning. *arXiv preprint arXiv:2006.07805*.
- Yi, K.; and Wu, J. 2019. Probabilistic end-to-end noise correction for learning with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7017–7025.
- Yu, X.; Han, B.; Yao, J.; Niu, G.; Tsang, I.; and Sugiyama, M. 2019. How does disagreement help generalization against label corruption? In *International Conference on Machine Learning*, 7164–7173. PMLR.
- Yu, X.; Liu, T.; Gong, M.; and Tao, D. 2018. Learning with biased complementary labels. In *Proceedings of the European conference on computer vision (ECCV)*, 68–83.
- Zhang, C.; Bengio, S.; Hardt, M.; Recht, B.; and Vinyals, O. 2016. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*.
- Zhang, Y.; Zheng, S.; Wu, P.; Goswami, M.; and Chen, C. 2021. Learning with Feature-Dependent Label Noise: A Progressive Approach. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Zhou, T.; Wang, S.; and Bilmes, J. 2021. Robust Curriculum Learning: From clean label detection to noisy label self-correction. In *Proceedings of the International Conference on Learning Representations, Lisbon, Portugal*, 28–29.