

Label Noise-Robust Learning using a Confidence-Based Sieving Strategy

Anonymous authors

Paper under double-blind review

Abstract

In learning tasks with label noise, boosting model robustness against overfitting is a pivotal challenge because the model eventually memorizes labels including the noisy ones. Identifying the samples with corrupted labels and preventing the model from learning them is a promising approach to address this challenge. Per-sample training loss is a previously studied metric that considers samples with small loss as *clean* samples on which the model should be trained. In this work, we first demonstrate the ineffectiveness of this small-loss trick. Then, we propose a novel discriminator metric called *confidence error* and a *sieving* strategy called CONFES to effectively differentiate between the clean and noisy samples. We experimentally illustrate the superior performance of our proposed approach compared to recent studies on various settings such as synthetic and real-world label noise. Moreover, we show CONFES can be combined with other approaches such as Co-teaching and DivideMix to further improve the model performance.

1 Introduction

The superior performance of deep neural networks (DNNs) in numerous application domains, ranging from medical diagnosis De Fauw et al. (2018); Liu et al. (2019) to autonomous driving Grigorescu et al. (2020) mainly relies on the availability of large-scale and high quality data Sabour et al. (2017); Marcus (2018). Supervised machine learning in particular requires correctly annotated datasets to train highly accurate DNNs. However, such datasets are rarely available in practice due to labeling errors (leading to *label noise*) stemming from high uncertainty Beyer et al. (2020) or lack of expertise Peterson et al. (2019). In medical applications for instance, there might be a disagreement between the labels assigned by radiology experts and those from the corresponding medical reports Majkowska et al. (2020); Bernhardt et al. (2022), yielding datasets with noisy labels.

In real-world scenarios, a corrupted label assigned to a sample depends on the feature values and the true label of the sample, which is known as *instance-dependent noise* Liu (2021); Zhang et al. (2021b). There are also other types of label noise in the literature including *symmetric noise*, where a sample is allocated a random label or *pairflip noise*, in which the label of a sample is flipped into the adjacent label Patrini et al. (2017); Xia et al. (2020); Bai et al. (2021). Instance-dependent noise is more challenging than the other label noise types Yao et al. (2020); Berthon et al. (2021). Training DNNs in the presence of label noise can lead to memorization of noisy labels, and consequently, reduction in model generalizability Zhang et al. (2021a); Chen et al. (2021b). Hence, it is indispensable to design and develop robust learning algorithms that are able to alleviate the adverse impact of noisy labels during training. Throughout this paper, we will refer to these methods as *label noise learning* methods.

Some of the existing studies Patrini et al. (2017); Berthon et al. (2021); Yao et al. (2021b); Xia et al. (2019); Yao et al. (2020) model the noise distribution as a transition matrix, encapsulating the probability of clean labels being flipped into noisy ones, and leverage loss correction to attenuate the effect of the noisy samples. Other studies Cheng et al. (2020); Xia et al. (2021); Wei et al. (2020) learn the clean label distribution and capitalize on regularization or selection of reliable samples to cope with the noisy labels. Although the former line of work theoretically guarantees the consistency between the classifiers learned with and without




Figure 1: **Distributions of loss values** for clean and noisy labels are similar, and consequently, loss value is an inefficient metric to distinguish between the clean labels and noisy ones. Experiments are performed using PreAct-ResNet18 trained on CIFAR-100 with instance-dependent label noise of level 60%.




Figure 2: **Distributions of confidence error values** for clean and noisy samples progressively diverge from each other as the training process continues. This indicates that confidence error is an effective metric to differentiate the clean samples from noisy ones. Experiments are conducted using PreAct-ResNet18 and CIFAR-100 with noise level of 60%.

noisy samples, its efficiency strongly depends on the accurate estimation of the transition matrix, which is a non-trivial task Xia et al. (2019); Bernhardt et al. (2022); Cheng et al. (2020).

A main challenge in the latter line of work, also known as sample selection (sieving), is to find a reliable criterion (or metric), which can efficiently differentiate between clean and noisy samples. The majority of the previous studies Jiang et al. (2018); Han et al. (2018); Yu et al. (2019) employ the *loss value* to this end, where the samples with small loss values are considered to likely be clean ones (*small-loss trick*). However, our observations (Fig 1) illustrate the distributions of the loss values for clean and noisy samples overlap widely, implying that a lot of noisy samples have small loss values and vice versa. Consequently, loss value is not an effective metric to distinguish between clean and noisy samples.

Contributions. We propose a novel metric called *confidence error* to more efficiently discriminate between the clean and noisy labels. The confidence error metric is defined as the difference between the softmax outputs/logits of the predicted and original label of a sample. Our observations (Fig 1) indicate that there exist a clear correlation between the confidence error value and the probability of being clean. That is, a sample with lower confidence error has much higher probability to be a clean sample than a noisy one.

Next, we integrate the confidence error criterion into a novel learning algorithm called *CONFidence Error Sieving* (CONFES) to robustly train DNNs in the instance-dependent, symmetric, and pairflip label noise settings. The CONFES algorithm computes the confidence error associated with training samples at the beginning of each epoch and only incorporates a subset of training samples with lower confidence error values

during training (i.e. likely clean samples). We draw a performance comparison between CONFES and its competitors using typical benchmark datasets for label noise learning including CIFAR-10/100 Krizhevsky et al. (2009) and Clothing1M Xiao et al. (2015).

In summary, we make the following contributions:

- We demonstrate that loss value is not an effective metric to discriminate between the clean and noisy samples.
- To address this shortcoming, we introduce the confidence error as a novel alternative metric and illustrate that it can efficiently differentiate clean samples from noisy ones.
- We propose the CONFES learning algorithm, which leverages the confidence error as a core building block to effectively *sieve* the training samples in an *online fashion* during training.
- Through extensive experiments, we show that CONFES outperforms the state-of-the-art learning algorithms of comparison, in terms of accuracy in symmetric, pairflip, and instance-dependent label noise scenarios.
- We combine CONFES with other learning algorithms including CoTeaching, JoCor, and DivideMix and illustrate the combination provides further improved accuracy.

2 Related Work

Overcoming the memorization of noisy labels plays a crucial role in label noise learning and improves model generalization by making the training process more robust to label noise Song et al. (2022); Zhang et al. (2021a); Natarajan et al. (2013); Arpit et al. (2017). The research community mainly tackled the memorization problem by adjusting the loss function Yao et al. (2020); Xia et al. (2019); Yao et al. (2021b), using implicit/explicit regularization techniques Xia et al. (2021); Liu et al. (2020); Zhang et al. (2018); Chen et al. (2021a); Song et al. (2019a), or refining the training data and performing sample sieving (also called sample selection) Han et al. (2018); Jiang et al. (2018); Yu et al. (2019); Cheng et al. (2020); Li et al. (2020); Yao et al. (2021a); Malach & Shalev-Shwartz (2017); Kim et al. (2021); Pleiss et al. (2020); Hu et al. (2021).

Adjusting the loss function according to the noise transition probabilities is an effective method for decreasing the adverse impact of noisy samples during the training but comes at the cost of accurate estimation of the transition matrix Patrini et al. (2017). Previous studies Yao et al. (2020); Xia et al. (2019); Yao et al. (2021b) have paved the way for this non-trivial estimation in different ways. For instance, T-Revision Xia et al. (2019) estimates the transition matrix without requiring anchor points (the data points whose associated class is known almost surely), which play an important role in learning the transition matrix effectively. Dual-T Yao et al. (2020) first divides the transition matrix into two matrices that are easier to estimate, and then aggregates their outputs for more accurate estimation of the original transition matrix.

A different line of work Liu et al. (2020); Xia et al. (2021); Zhang et al. (2018); Chen et al. (2021a) improves model generalization by introducing regularization effects suitable for learning with noisy labels. The regularization effect may be injected implicitly using methods such as data augmentation Zhang et al. (2018) and inducing stochasticity Chen et al. (2021a). For example, Mixup Zhang et al. (2018) augments the training data using a convex combination of a pair of examples and their labels which encourages the model to learn a simple interpolation between the samples. SLN Chen et al. (2021a) introduces stochastic label noise to help the optimizer to skip sharp minima in the optimization landscape.

Although conventional (implicit) regularization techniques such as dropout Srivastava et al. (2014) and data augmentation have been proven effective in alleviating overfitting and improving generalization, they are insufficient to tackle the label noise challenge Song et al. (2022). ELR (Early-Learning Regularization) Liu et al. (2020) is an explicit regularization approach based on the observation that at the beginning of training, there is an early-learning phase in which the model learns the clean samples without overfitting the noisy ones. Given that, ELR adds a regularization term to the Cross-Entropy (CE) loss, leading the model output towards its own (correct) predictions at the early-learning phase. Similarly, CDR Xia et al. (2021) groups

the model parameters into critical and non-critical in terms of their importance for generalization, and then penalizes the non-critical parameters.

Differentiating the clean samples from the noisy ones (known as sample selection/sieving) and employing them in the training process, is a promising direction to enhance robustness in label noise learning Han et al. (2018). MentorNet Jiang et al. (2018) uses an extra pretrained model (mentor) to help the main model (student) by providing it with small-loss samples. The decoupling algorithm Malach & Shalev-Shwartz (2017) trains two networks simultaneously using the samples on which the models disagree about the predicted label. Co-teaching Han et al. (2018), cross-trains two models such that each of them leverages the samples with small-loss values according to the other model. Co-teaching+ Yu et al. (2019) improves Co-teaching by considering the clean samples as the ones that not only have small loss but also those on which the models disagree. JoCoR Wei et al. (2020) computes a joint-loss to make the outputs of the two models become closer, then it considers the samples with small loss as clean samples.

Algorithm 1: Confidence error based sieving (CONFES)

Input: Noisy training dataset $\tilde{D} = \{(x_i, \tilde{y}_i)\}_{i=1}^n$, model \mathcal{F}_θ , number of training epochs T , initial sieving threshold α , number of warm-up epochs T_w , batch size M

Output: Trained model \mathcal{F}_θ

```

1 for  $i = 0, \dots, T$  do
2    $\alpha_i = \text{Max}(\alpha - i \cdot \frac{\alpha}{T_w}, 0)$  /* Set sieving threshold */
3    $O^i \leftarrow \mathcal{F}_\theta^i(\tilde{D})$  /* Feed the dataset to model */
4    $C^i \leftarrow \mathcal{P}(O^i)$  /* Calculate confidence on labels */
5    $E_C(S) = C^i(\tilde{y}_i) - C^i(\tilde{y}_i)$  /* Compute confidence error */
6    $D_c = \{S \mid E_C(S) \leq \alpha_i\}$  /* Sieve clean samples */
7    $D' = D_c \oplus (\{(\text{augment}(x_i), \tilde{y}_i)\} \mid (x_i, y_i) \in D_c)$  /* Build new dataset(clean $\oplus$ augmented) */
8   for mini-batch  $\beta = \{(x_i, \tilde{y}_i)\}_{i=1}^M \in D'$  do
9     /* Train the model on new dataset */
9     Update model  $\mathcal{F}_\theta$  on mini batch  $\beta$  using Equation equation 1
10  return Trained model  $\mathcal{F}_\theta$ 

```

3 CONFES: CONFidence Error based Sieving

We assume a classification task on a training dataset $D = \{(x_i, y_i) \mid x_i \in X, y_i \in Y\}_{i=1}^n$, where n is the number of samples and X and Y are the feature and label (class) space, respectively. The neural network model $\mathcal{F}(X_\beta; \theta) \in \mathbb{R}^{m \times k}$ is a k -class classifier with trainable parameters θ that takes mini-batches X_β of size m as input. In real life, a sample might be assigned a wrong label (e.g. due to human error). Consequently, *clean* (that is, label noise-free) training datasets are not always available in practice. Given that, assume $\tilde{Y} = \{\tilde{y}_i\}_{i=1}^n$ and $\tilde{D} = \{(x_i, \tilde{y}_i)\}_{i=1}^n$ indicate the noisy labels and noisy dataset, respectively. The training process is conducted by minimizing the empirical loss (e.g. cross-entropy) using mini-batches of samples from the noisy dataset:

$$\min_{\theta} \mathcal{L}(\mathcal{F}(X_\beta; \theta); \tilde{Y}_\beta) = \min_{\theta} \frac{1}{m} \sum_{i=1}^m \mathcal{L}(\mathcal{F}(x_i, \theta), \tilde{y}_i), \quad (1)$$

where \mathcal{L} is the loss function and $(X_\beta, \tilde{Y}_\beta)$ is a mini-batch of samples with size m from the noisy dataset \tilde{D} .

In the presence of label-noise, the efficiency of the training process mainly depends on the capability of the model to distinguish between clean and noisy labels and to diminish the impact of noisy ones on the training process. The loss value is a commonly used metric to that end, where a sample with lower loss value is considered to be more likely a clean sample than a noisy one Han et al. (2018); Jiang et al. (2018); Yu et al. (2019); Li et al. (2020). We design an experiment to investigate the effectiveness of this consideration: we employ the SGD optimizer and cross-entropy loss function to train PreActResNet18 on CIFAR-100, where 60% of the labels are made noisy using symmetric, asymmetric or instance-dependent noise. At the beginning

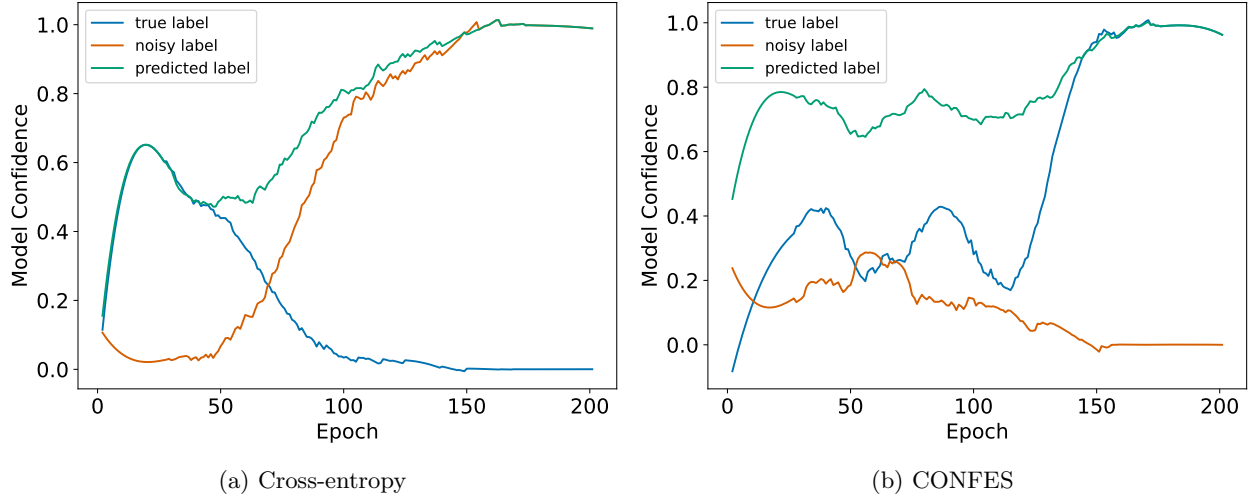


Figure 3: **Why confidence error works effectively in label noise training?** The model is confident about the true label of the sample during the initial epochs, although a noisy label is given to it as the original label. However, when using the cross-entropy training, the model is misled by the noisy sample as training proceeds, wrongly predicting the label as the noisy label at the end of training. CONFES, however, detects the noisy sample based on its confidence error value at the early epochs and prevents it from misleading the model. Therefore, the model still correctly remains confident about its prediction until the end of training.

of each epoch, the model computes the loss value for all training samples, and sorts them in ascending order by loss value. The model considers the first 40% of the samples with lower loss values as clean, and only incorporate them during training. This procedure is repeated for 200 epochs.

As seen in Figure 1, the distribution of the loss values for the clean and noisy labels are alike. Moreover, the corresponding distributions still remain similar as the training process continues. This implies that there is a weak correlation between the loss value and the likelihood of being a clean sample, and consequently, loss value is not an effective metric to discriminate between the clean and noisy samples. This observation motivates us to propose an alternative metric called *confidence error* for more efficient sieving of the samples during training.

Confidence Error Consider a sample $S = (x_i, \tilde{y}_i)$ from the noisy dataset \tilde{D} . The k -class/label classifier $\mathcal{F}(x_i; \theta)$ takes x_i as input and computes the weight value associated with each class as output. Moreover, assume $\mathcal{P}(\cdot)$ is the softmax activation function such that $\mathcal{P}(\mathcal{F}(x_i; \theta)) \in [0, 1]^k$ takes classifier’s output and computes the predicted probability for each class. We define the *model confidence* for a given label $l \in \{1, \dots, k\}$ associated with sample S as the prediction probability assigned to the label:

$$C^{(l)} = \mathcal{P}(\mathcal{F}(x_i; \theta))^{(l)} \quad (2)$$

The class with the maximum probability is considered as the predicted class, i.e. \hat{y}_i , for sample S :

$$\hat{y}_i = \arg \max_{j \in \{1, \dots, k\}} \mathbb{P}^{(j)}(\mathcal{F}(x_i; \theta)), \quad (3)$$

where \mathbb{P} indicates probability. The *confidence error* $E_C(S)$ for sample S is defined as the difference between the probability assigned to the predicted label \hat{y}_i and the probability associated with the original label \tilde{y}_i :

$$E_C(S) = C^{(\hat{y}_i)} - C^{(\tilde{y}_i)}, \quad (4)$$

where $E_C(S) \in [0, 1]$. In other words, the confidence error states the possibility that the predicted class is the same as the original class. The confidence error of zero implies that the original and predicted classes are the same.




Figure 4: Test accuracy on CIFAR-100 using PreAct-ResNet18: CONFES is robust against overfitting, whereas some of its competitors including SLN and ELR start overfitting after some epochs of training.

We repeat the previous experiment, but with confidence error as the metric to differentiate the clean samples from noisy ones. As shown in Figure 7, the distribution of confidence error values for the clean and noisy samples becomes more and more dissimilar as the training process proceeds. For instance at epoch 50, a sample with high confidence error (e.g. near 1.0) is much more likely to be a noisy sample than clean one. Likewise, a sample with very low confidence error is most probably a clean sample. Given that, we conclude that confidence error is a more efficient metric than loss value to distinguish between the clean and noisy samples.

Why confidence error works effectively in label noise training? We design a simple experiment to answer this question. We consider a sample with true label of *Cat*, and assign the noisy label of *Horse* to it. According to Figure 3a, in initial epochs the model is confident that the label of the sample is *Cat*, although the label of *Horse* is given to the model as the original label. However, the model is misled by the noisy sample as training proceeds with the naive cross-entropy approach, wrongly predicting the label as *Horse* at the end of training. CONFES, on the other hand, identifies that the label of the sample is corrupted based on the confidence error value in the early epochs and excludes it from training. According to Figure 3b, the model remains confident about its prediction until the end of training.

CONFES algorithm Previous studies Bai et al. (2021); Liu et al. (2020) show that deep neural networks tend to memorize noisy samples, which can have detrimental effect on the model utility. Therefore, it is crucial to detect the noisy samples and alleviate their adverse impact, especially in the early steps of training. The CONFES algorithm takes this into consideration by sieving the training samples using the confidence error metric and *completely excluding* the identified noisy samples during training. CONFES (Algorithm 1) consists of three main steps at each epoch: Sieving samples, building the refined training set, and training the model.

In the sieving step, the confidence error for each training sample is computed using Equation 4; then, the samples whose confidence error is less than or equal to α_i (sieving threshold at epoch i) are considered as clean, whereas the remaining samples are assumed to be noisy and excluded from training. CONFES has two hyper-parameters: initial sieving threshold α , and the number of warm-up epochs T_w , which are used to compute per-epoch sieving threshold α_i at each epoch i . In the second step, a new training dataset is created by concatenating (\oplus) only the identified clean samples and their augmentations such that this dataset becomes as large as the initial training set. Finally, the model is trained on the augmented dataset, which consists of the most probably clean samples according to their confidence error.

4 Experiments

In this section, we draw a performance comparison between CONFES and recent baseline approaches Han et al. (2018); Liu et al. (2020); Chen et al. (2021a); Cheng et al. (2020); Bai et al. (2021) on three label noise

settings: symmetric, pairflip, and instance-dependent. We first describe our experimental setting, and then provide and discuss the empirical results.

Datasets We employ the CIFAR-10/100 datasets Krizhevsky et al. (2009) and Clothing1M datasets Xiao et al. (2015), which is naturally noisy and utilised as a standard benchmark for label noise training tasks. CIFAR-10/100 contain 50000 training samples and 10000 testing samples of shape 32×32 from 10/100 classes. For the CIFAR datasets, we perturb the training labels using symmetric, pairflip, and instance-dependent label noise introduced in Xia et al. (2020), but keep the test set clean. For data augmentation/preprocessing, the training samples are horizontally flipped with probability 0.5, randomly cropped with size 32×32 and padding 4×4 , and normalized using the mean and standard deviation of the dataset. Clothing1M is a real-world dataset of 1 million images of size 224×224 with noisy labels (whose estimated noise level is approximately 38% Wei et al. (2022); Song et al. (2019b)) and 10k clean test images in 14 classes. The data augmentation methods performed on this dataset include 256×256 resizing, 224×224 random crops and random horizontal flips. In the clothing1M training dataset, the number of samples for each class is imbalanced. Thus, we follow Li et al. (2020) and sample a class-balanced subset of the training dataset at each epoch.

State-of-the-art methods On all of the examined datasets, we compare CONFES with the most recent related studies including: (1) standard cross-entropy loss (CE), (2) co-teaching Han et al. (2018) that cross-trains two models and uses the small-loss trick for selecting clean samples and exchanges them between the two models, (3) ELR Liu et al. (2020), an early-learning regularization method that leverages the model’s output during the early-learning phase, (4) CORES² Cheng et al. (2020), a sample sieving approach that uses confidence regularization which leads the model towards having more confident predictions, (5) PES Bai et al. (2021), a progressive early-stopping strategy and (6) SLN Chen et al. (2021a) that improves regularization by introducing stochastic label noise. Co-teaching and CORES² are based on sample selection, ELR and SLN are regularization-based methods. For all of these methods, the specific hyper-parameters are set according to the corresponding manuscript or the published source code, if available.

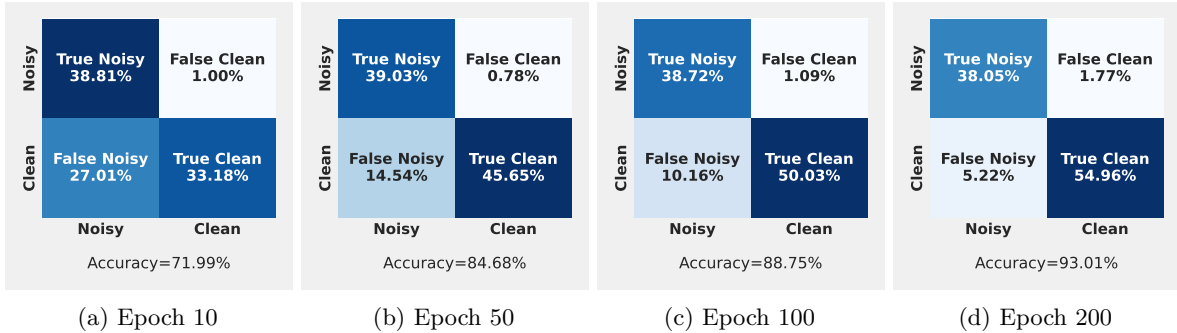


Figure 5: Confusion matrix for the CONFES algorithm. In early epochs, CONFES correctly identifies the majority of noisy labels (around 38% out of 40%), but wrongly identifies many clean labels as noisy ones (about 27%). As training proceeds, the algorithm not only still remains effective in identifying the noisy labels (around 38% out of 40%), but also correctly recognizes the clean labels (about 55% out of 60%). In the experiment, the dataset is CIFAR-100 with instance-dependant label noise of rate 40% trained using PreActResNet18.

Neural Networks and optimization We conduct the experiments on a single GPU system equipped with an NVIDIA RTX A6000 graphic processor and 48GB of GPU memory. Our method is implemented in PyTorch v1.9. For all methods, we evaluate the average test accuracy on the last 5 epochs and for co-teaching, we report the average of this metric for the two networks. Following previous works Li et al. (2020); Bai et al. (2021) on CIFAR-10 and CIFAR-100, we train the PreActResNet-18 He et al. (2016) model using the SGD optimizer with a momentum of 0.9, weight decay of $5e-4$ and batch size of 128. The initial learning rate is set to 0.02 and is decreased by 0.01 in 300 epochs using a cosine annealing scheduler Loshchilov & Hutter (2017). On Cloting1M dataset, we adopt the setting from Li et al. (2020) and train the ResNet-50 model

pre-trained on ImageNet Krizhevsky et al. (2012). The network is trained using SGD with momentum of 0.9 and weight decay of $1e-3$ for 80 epochs, starting with learning rate of 0.002 which decreased by 10 at epoch 40. At each epoch the network is trained on 1000 mini-batches of size 32.

4.1 Evaluation

CIFAR-10/100 datasets Tables 1 and 2 list test accuracy values for different noise types and noise rates, respectively. According to Table 1, CONFES outperforms the competitors for all considered symmetric, pairflip, and instance-dependent noise types. Similarly, CONFES delivers higher accuracy than the competitors for different noise rates as shown in Table 2. Moreover, as the noise level increases, the accuracy gap between CONFES and its competitors widens in favor of CONFES. Figure 4 illustrates the test accuracy versus epoch for the different learning algorithms. As shown in the figure, CONFES is robust against overfitting because the corresponding test accuracy continues to increase as training moves forward, and stays at the maximum after the model converges. Some of the other algorithms such as SLN and ELR on the other hand suffer from an overfitting problem, where their final accuracy values are lower than the maximum accuracy they achieve.

Table 1: Test accuracy for different noise types. The model is PreActResNet18, and noise level is 40%.

(a) **CIFAR-10**

Method	Symmetric	Pairflip	Instance
CE	66.61 \pm 0.4	59.25 \pm 0.1	66.04 \pm 0.2
Co-teaching Han et al. (2018)	87.42 \pm 0.2	84.57 \pm 0.2	86.90 \pm 0.1
ELR Liu et al. (2020)	85.74 \pm 0.2	86.15 \pm 0.1	85.37 \pm 0.3
CORES ² Cheng et al. (2020)	83.9 \pm 0.4	58.38 \pm 0.6	76.71 \pm 0.4
PES Bai et al. (2021)	90.55 \pm 0.1	85.56 \pm 0.1	85.63 \pm 0.5
SLN Chen et al. (2021a)	83.69 \pm 0.2	85.26 \pm 0.5	67.71 \pm 0.4
CONFES (ours)	90.62 \pm 0.2	86.18 \pm 0.3	90.28 \pm 0.2

(b) **CIFAR-100**

Method	Symmetric	Pairflip	Instance
CE	47.85 \pm 0.4	45.62 \pm 0.4	48.92 \pm 0.3
Co-teaching Han et al. (2018)	66.26 \pm 0.1	55.42 \pm 0.5	66.55 \pm 0.3
ELRLiu et al. (2020)	48.33 \pm 0.2	44.31 \pm 0.3	49.15 \pm 0.2
CORES ² Cheng et al. (2020)	52.26 \pm 0.2	45.34 \pm 0.2	50.98 \pm 0.2
PES Bai et al. (2021)	68.37 \pm 0.2	59.18 \pm 0.2	64.83 \pm 0.2
SLN Chen et al. (2021a)	46.98 \pm 0.2	45.72 \pm 0.2	46.08 \pm 0.3
CONFES (ours)	69.63 \pm 0.2	62.37 \pm 0.4	69.68 \pm 0.2

Table 3 shows the accuracy values of CoTeaching, JoCor, and DivideMix if confidence error is used as the discriminator metric instead of the training loss. As shown in the table, the accuracy from these algorithms is enhanced by 2-5%, indicating that confidence error is not only effective as the main building block of the proposed CONFES algorithm, but also combined with other state-of-the-art methods such as Co-teaching, JoCor and DivideMix, which is a complex method employing data augmentation and relabeling.

Clothing1M dataset Table 4 summarises the performance of methods on Clothing1M dataset. CORES² and PES provide slight or no accuracy gain compared to the baseline cross-entropy training, respectively. CONFES, on the other hand, outperforms the competitors including ELR and SLN.

Table 2: Test accuracy for different noise rates; dataset: CIFAR-100; model: PreActResNet18; noise: instance-dependent.

Method	Inst.-30%	Inst.-50%	Inst.-60%
CE	56.54 \pm 0.3	39.37 \pm 0.3	30.65 \pm 0.4
Co-teaching Han et al. (2018)	69.67 \pm 0.2	62.19 \pm 0.3	57.18 \pm 0.2
ELR Liu et al. (2020)	56.36 \pm 0.3	40.18 \pm 0.6	29.88 \pm 0.6
CORES ² Cheng et al. (2020)	58.23 \pm 0.3	42.42 \pm 0.3	33.93 \pm 0.5
PES Bai et al. (2021)	69.49 \pm 0.2	55.47 \pm 0.2	41.10 \pm 0.5
SLN Chen et al. (2021a)	53.98 \pm 0.4	38.06 \pm 0.5	29.77 \pm 0.4
CONFES (ours)	71.96\pm0.5	66.44\pm0.2	59.48\pm0.1

Table 3: Test accuracy for CONFES (CNF) combined with other approaches; dataset: CIFAR-100; model: PreActResNet18; noise rate: 40%.

Method	Symmetric	Pairflip	Instance
Co-teaching Han et al. (2018)	66.26 \pm 0.1	55.42 \pm 0.5	66.55 \pm 0.3
CNF-Co-teaching	69.94 \pm 0.1	57.90 \pm 0.2	69.51 \pm 0.1
Improvement	+3.68	+2.48	+2.96
DivideMix Li et al. (2020)	74.63 \pm 0.2	74.9 \pm 0.1	66.79 \pm 0.3
CNF-DivideMix	76.31 \pm 0.2	76.51 \pm 0.1	69.03 \pm 0.1
Improvement	+1.68	+1.61	+2.24
JoCoR Wei et al. (2020)	67.05 \pm 0.2	54.96 \pm 0.3	67.46 \pm 0.2
CNF-JoCoR	70.48 \pm 0.2	59.61 \pm 0.1	70.24 \pm 0.4
Improvement	+3.43	+4.65	+2.78

Table 4: Test accuracy on Clothing1M dataset

Method	CE	ELR	CORES ²	PES	SLN	CONFES (ours)
Test Accuracy	69.21%	71.39%	69.50%	69.18%	72.80%	73.24%

5 Discussion

According to the experimental results, CONFES outperforms all baseline methods in the considered symmetric, pairflip, and instant-dependent noise settings. As the noise rate increases, the efficiency of the CONFES algorithm becomes more apparent (e.g. noise rate of 50% in CIFAR-100). Moreover, CONFES is robust to overfitting unlike some of its competitors such as SLN and ELR. This indicates that the underlying confidence error metric can effectively differentiate the clean labels from the noisy ones, and eliminate their adverse impact by excluding them from the training process.

Figure 5 shows the confusion matrix for the CONFES algorithm. According to the figure, CONFES is effective in recognizing the noisy samples from the beginning to the end of training, where it correctly identifies around 38% out of 40% of noisy labels. On the other hand, the algorithm wrongly identifies many clean samples as noisy in the early epochs (around 27%). However, as training moves forward, CONFES becomes more and more efficient in distinguishing the clean samples, where it correctly identifies around 55% out of 60% of the clean samples.

The initial sieving threshold α and number of warm-up epochs T_w are the hyper-parameter of our proposed method. The per-epoch sieving threshold is computed using the aforementioned hyper-parameters. For

CIFAR-100, we set $\alpha=0.2$ and $T_w=30$ for all noise types and noise rates. For CIFAR-10, $\alpha=0.1$ and $T_w=25$ in symmetric and instance-dependent noise types. For Clothing1M, α and T_w are set to 0.05 and 3, respectively. In general, the value of these hyper-parameters should be set to higher values for more complex training processes, which depends on the model and dataset combination.

The CONFES strategy sieves the clean samples accurately and improves the robustness without adding too much overhead to the standard Cross-Entropy approach. In terms of computational overhead, CONFES has one additional forward pass for constructing the refined dataset, which only includes clean samples according to the confidence error metric. However, methods such as Co-teaching Han et al. (2018) employ two networks in the training process, which makes them substantially more computationally inefficient compared to our approach.

6 Conclusion

We present an effective label noise learning algorithm called CONFES for different noise settings such as symmetric, pairflip, and instance-dependent. CONFES is based on the proposed confidence error metric, which exploits the model’s predictive confidence for discriminating between clean and noisy samples. We illustrate the efficacy of the confidence error metric in differentiation of the clean samples from the noisy ones. CONFES refines the training samples by keeping only the identified clean samples and filtering out the noisy ones. Our empirical results verify the robustness of CONFES under different noise type scenarios, especially when noise levels are high. Moreover, we demonstrate that confidence error can be employed by other algorithms such as co-teaching and DivideMix to further improve the model accuracy.

References

- Devansh Arpit, Stanisław Jastrzębski, 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*, pp. 233–242. PMLR, 2017.
- Yingbin Bai, Erkun Yang, Bo Han, Yanhua Yang, Jiatong Li, Yinian Mao, Gang Niu, and Tongliang Liu. Understanding and improving early stopping for learning with noisy labels. *Advances in Neural Information Processing Systems*, 34, 2021.
- Mélanie Bernhardt, Daniel C Castro, Ryutaro Tanno, Anton Schwaighofer, Kerem C Tezcan, Miguel Monteiro, Shruthi Bannur, Matthew P Lungren, Aditya Nori, Ben Glocker, et al. Active label cleaning for improved dataset quality under resource constraints. *Nature communications*, 13(1):1–11, 2022.
- Antonin Berthon, Bo Han, Gang Niu, Tongliang Liu, and Masashi Sugiyama. Confidence scores make instance-dependent label-noise learning possible. In *International Conference on Machine Learning*, pp. 825–836. PMLR, 2021.
- Lucas Beyer, Olivier J Hénaff, Alexander Kolesnikov, Xiaohua Zhai, and Aäron van den Oord. Are we done with imagenet? *arXiv preprint arXiv:2006.07159*, 2020.
- Pengfei Chen, Guangyong Chen, Junjie Ye, Jingwei Zhao, and Pheng-Ann Heng. Noise against noise: stochastic label noise helps combat inherent label noise. In *International Conference on Learning Representations*, 2021a.
- Pengfei Chen, Junjie Ye, Guangyong Chen, Jingwei Zhao, and Pheng-Ann Heng. Robustness of accuracy metric and its inspirations in learning with noisy labels. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021b.
- Hao Cheng, Zhaowei Zhu, Xingyu Li, Yifei Gong, Xing Sun, and Yang Liu. Learning with instance-dependent label noise: A sample sieve approach. *International Conference on Learning Representations*, 2020.
- Jeffrey De Fauw, Joseph R Ledsam, Bernardino Romera-Paredes, Stanislav Nikolov, Nenad Tomasev, Sam Blackwell, Harry Askham, Xavier Glorot, Brendan O’Donoghue, Daniel Visentin, et al. Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature medicine*, 24(9):1342–1350, 2018.

- Sorin Grigorescu, Bogdan Trasnea, Tiberiu Cocias, and Gigel Macesanu. A survey of deep learning techniques for autonomous driving. *Journal of Field Robotics*, 37(3):362–386, 2020.
- 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. *Advances in neural information processing systems*, 31, 2018.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016.
- Wei Hu, QiHao Zhao, Yangyu Huang, and Fan Zhang. P-diff: Learning classifier with noisy labels based on probability difference distributions. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pp. 1882–1889. IEEE, 2021.
- 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*, pp. 2304–2313. PMLR, 2018.
- Taehyeon Kim, Jongwoo Ko, Sangwook Cho, JinHwan Choi, and Se-Young Yun. FINE samples for learning with noisy labels. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=QZpx42n0BWr>.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Citeseer*, 2009.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- Junnan Li, Richard Socher, and Steven C.H. Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *International Conference on Learning Representations*, 2020.
- 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.
- Xiaoxuan Liu, Livia Faes, Aditya U Kale, Siegfried K Wagner, Dun Jack Fu, Alice Bruynseels, Thushika Mahendiran, Gabriella Moraes, Mohith Shamdass, Christoph Kern, et al. A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. *The lancet digital health*, 1(6):e271–e297, 2019.
- Yang Liu. The importance of understanding instance-level noisy labels. *arXiv e-prints*, pp. arXiv–2102, 2021.
- Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with warm restarts. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- Anna Majkowska, Sid Mittal, David F Steiner, Joshua J Reicher, Scott Mayer McKinney, Gavin E Duggan, Krish Eswaran, Po-Hsuan Cameron Chen, Yun Liu, Sreenivasa Raju Kalidindi, et al. Chest radiograph interpretation with deep learning models: assessment with radiologist-adjudicated reference standards and population-adjusted evaluation. *Radiology*, 294(2):421–431, 2020.
- Eran Malach and Shai Shalev-Shwartz. "Decoupling" when to update" from "how to update". *Advances in Neural Information Processing Systems*, 30, 2017.
- Gary Marcus. Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*, 2018.
- Nagarajan Natarajan, Inderjit S Dhillon, Pradeep K Ravikumar, and Ambuj Tewari. Learning with noisy labels. *Advances in neural information processing systems*, 26, 2013.
- 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*, pp. 1944–1952, 2017.

- Joshua C Peterson, Ruairidh M Battleday, Thomas L Griffiths, and Olga Russakovsky. Human uncertainty makes classification more robust. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9617–9626, 2019.
- Geoff Pleiss, Tianyi Zhang, Ethan Elenberg, and Kilian Q Weinberger. Identifying mislabeled data using the area under the margin ranking. *Advances in Neural Information Processing Systems*, 33:17044–17056, 2020.
- Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. Dynamic routing between capsules. *Advances in neural information processing systems*, 30, 2017.
- Hwanjun Song, Minseok Kim, and Jae-Gil Lee. Selfie: Refurbishing unclean samples for robust deep learning. In *International Conference on Machine Learning*, pp. 5907–5915. PMLR, 2019a.
- Hwanjun Song, Minseok Kim, Dongmin Park, and Jae-Gil Lee. How does early stopping help generalization against label noise? In *ICML 2020 Workshop on Uncertainty and Robustness in Deep Learning*, 2019b.
- Hwanjun Song, Minseok Kim, Dongmin Park, Yooju Shin, and Jae-Gil Lee. Learning from noisy labels with deep neural networks: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1): 1929–1958, 2014.
- Hongxin Wei, Lei Feng, Xiangyu Chen, and Bo An. Combating noisy labels by agreement: A joint training method with co-regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13726–13735, 2020.
- Jiaheng Wei, Zhaowei Zhu, Hao Cheng, Tongliang Liu, Gang Niu, and Yang Liu. Learning with noisy labels revisited: A study using real-world human annotations. In *International Conference on Learning Representations*, 2022.
- Xiaobo Xia, Tongliang Liu, Nannan Wang, Bo Han, Chen Gong, Gang Niu, and Masashi Sugiyama. Are anchor points really indispensable in label-noise learning? *Advances in Neural Information Processing Systems*, 32, 2019.
- Xiaobo Xia, Tongliang Liu, Bo Han, Nannan Wang, Mingming Gong, Haifeng Liu, Gang Niu, Dacheng Tao, and Masashi Sugiyama. Part-dependent label noise: Towards instance-dependent label noise. *Advances in Neural Information Processing Systems*, 33:7597–7610, 2020.
- Xiaobo Xia, Tongliang Liu, Bo Han, Chen Gong, Nannan Wang, Zongyuan Ge, and Yi Chang. Robust early-learning: Hindering the memorization of noisy labels. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=Eq15b1_hTE4.
- 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*, pp. 2691–2699, 2015.
- 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*, pp. 5192–5201, 2021a.
- 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. *Advances in neural information processing systems*, 33:7260–7271, 2020.
- Yu Yao, Tongliang Liu, Mingming Gong, Bo Han, Gang Niu, and Kun Zhang. Instance-dependent label-noise learning under a structural causal model. *Advances in Neural Information Processing Systems*, 34, 2021b.

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*, pp. 7164–7173. PMLR, 2019.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021a.

Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=r1Ddp1-Rb>.

Yivan Zhang, Gang Niu, and Masashi Sugiyama. Learning noise transition matrix from only noisy labels via total variation regularization. In *International Conference on Machine Learning*, pp. 12501–12512. PMLR, 2021b.

A Appendix

More details on experiment setup For all experiments on the CIFAR-10, CIFAR-100, and Clothing1M datasets, there are some general hyper-parameters such as learning rate, batch size and weight decay, specified in the original manuscript and are summarized in Table 5. The source code for CONFES is included in the supplementary material and will be publicly available. The method-specific hyper-parameters used in the experiments are set based on the corresponding manuscript or the published source code: Co-teaching Han et al. (2018)¹, ELR Liu et al. (2020)², CORES² Cheng et al. (2020)³, PES Bai et al. (2021)⁴, SLN Chen et al. (2021a)⁵, DivideMix Li et al. (2020)⁶ and JoCoR Wei et al. (2020)⁷.

Table 5: General training hyperparameters (common for all methods of comparison)

	CIFAR-10	CIFAR-100	Clothing1M
model	PreAct-ResNet18	PreAct-ResNet18	Pre-trained ResNet50
batch size	128	128	32
learning rate (lr)	2e-2	2e-2	2e-3
lr decay	Cosine annealing	Cosine annealing	By 0.1 at 40th
weight decay	5e-4	5e-4	1e-3
epochs	300	300	80

¹<https://github.com/bhanML/Co-teaching>

²<https://github.com/shengliu66/ELR>

³<https://github.com/UCSC-REAL/cores>

⁴<https://github.com/tmllab/PES>

⁵<https://github.com/chenpf1025/SLN>

⁶<https://github.com/LiJunnan1992/DivideMix>

⁷<https://github.com/hongxin001/JoCoR>

More experiments on the effectiveness of confidence error : We extended the experiments corresponding to Fig.1 and Fig.2 of the main manuscript, to the symmetric and pairflip label noise. Experiments are conducted using PreAct-ResNet18 and CIFAR-100 with noise level of 40%.




Figure 6: Distributions of **confidence error** values for **pairflip** label noise




Figure 7: Distributions of **confidence error** values for **symmetric** label noise