


Alleviating Noisy-label Effects in Image Classification via Probability Transition Matrix

Ziqi Zhang¹
zq-zhang18@mails.tsinghua.edu.cn

Yuxiang Li ²
vicyxli@tencent.com

Hongxin Wei³
hongxin001@e.ntu.edu.sg

Kai Ma²
kylekma@tencent.com

Tao Xu ¹
taoxu@tsinghua.edu.cn

Yefeng Zheng²
yefengzheng@tencent.com

¹ Tsinghua-Berkeley Shenzhen Institute,
Tsinghua University
Shenzhen, China



² Tencent Jarvis Lab
Shenzhen, China

³ School of Computer Science and
Engineering, Nanyang Technological
University
Singapore

Abstract

Deep-learning-based image classification frameworks often suffer from the noisy label problem caused by the inter-observer variation. Recent studies employed learning-to-learn paradigms (*e.g.*, Co-teaching and JoCoR) to filter the samples with noisy labels from the training set. However, most of them use a simple cross-entropy loss as the criterion for noisy label identification. The hard samples, which are beneficial for classifier learning, are often mistakenly treated as noises in such a setting, since both the hard samples and the ones with noisy labels lead to a relatively larger loss value than the easy cases. In this paper, we propose a plugin module, namely noise ignoring block (NIB), consisting of a probability transition matrix and an inter-class correlation (IC) loss, to separate the hard samples from the mislabeled ones, and further boost the accuracy of image classification network trained with noisy labels. Concretely, our IC loss is calculated as Kullback-Leibler divergence between the network prediction and the accumulative soft label generated by the probability transition matrix. Such that, with lower value of IC loss, the hard cases can be easily distinguished from mislabeled cases. Extensive experiments are conducted on natural and medical image datasets (CIFAR-10 and ISIC 2019). The experimental results show that our NIB module consistently improves the performances of the state-of-the-art robust training methods.

1 Introduction

Witnessing the success of deep neural networks (DNNs) for computer vision tasks  , an increasing number of researchers began to implement deep-learning-based approaches

for image classification tasks. However, due to the inter-observer variation, the training set often contains noisy labels, which may significantly degrade the model performance, *i.e.*, the memorization effects of DNNs [28]. Recent studies have proposed algorithms to train deep learning models with noisy labels, which can be mainly grouped into two categories, *i.e.*, noise estimation [10, 13, 17] and instance selection [3, 6, 11, 24]. The former approaches detect the samples with noisy labels and accordingly revise their labels via the noise transition matrix or noise rate derived from the prior-knowledge, which limits their application only to the data with known noise characteristics. The latter category addresses the problem by proposing learning-to-learn selection paradigms (*e.g.*, MentorNet [5], Co-teaching [9], and JoCoR [24]) for the robust learning with unknown noisy labels. Generally, they treat the samples with the smaller loss as ‘clean’, and exploit them for network training. The existing learning-to-learn approaches share a common drawback—most of them use a simple cross-entropy loss as the criterion for noisy label identification. The hard samples, which are beneficial for classifier learning, are often mistakenly treated as noises in such a setting, since both the hard samples and the ones with noisy labels lead to a relatively larger loss value than the easy cases.

In this paper, we propose a simple-yet-effective plugin module, namely noise ignoring block (NIB), which can be easily integrated to the existing learning-to-learn selection paradigms, to distinguish the hard samples from the mislabeled ones and further boost the performance of robust learning with noisy labels. Our NIB module consists of a probability transition matrix and an inter-class correlation (IC) loss. In particular, we calculate the Kullback-Leibler (KL) divergence between the network prediction and the accumulative soft-label generated from the probability transition matrix, which represents the inter-class correlation, as an auxiliary loss (*i.e.*, IC loss). Such that, with lower value of IC loss, the hard cases can be easily distinguished from the ones with noisy labels. The proposed NIB module is evaluated on two publicly available natural and medical image datasets (CIFAR-10 and ISIC 2019). Experimental results demonstrate that the proposed NIB significantly improves the classification accuracy of state-of-the-art robust training frameworks, *e.g.*, Co-teaching and JoCoR.

2 Method

In this section, we first uncover the difference between hard and mislabeled samples from the aspect of inter-class correlation, and then present our noise ignoring block (NIB) in details.

Problem Formulation. We first illustrate the difference between hard samples and the ones with noisy labels. Generally, the hard samples are the ones, which are mistakenly identified to a wrong but semantic-related class by the model, due to the similar visual features. For example, the dog contains more similar visual characteristics to cat than the other classes such as truck and airplane in CIFAR-10. Hence, the ‘hard’ dog sample mistakenly classified to cat, *i.e.*, falling around the decision boundary of dog/cat in the latent space, can aid the model to refine its decision boundary and should be paid more attention, which have been verified by existing studies [8, 18]. In contrast, the noisy labels are not necessary to be semantic-related, *e.g.*, a dog sample with a noisy label of ship. Here, we reveal a common drawback shared by the conventional robust learning approaches—they calculate the cross-entropy loss between network predictions (\mathbf{p}_{hard} and \mathbf{p}_{noisy}) and one-hot label \mathbf{y} , result in

Algorithm 1 Training Procedure with Noise Ignoring Block (NIB)

Input: Network f with Θ , learning rate η , epoch number N_e , iteration number N_i , batch size n_b , training set $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ consists of N samples from K classes;

- 1: **for** $n = 1, 2, \dots, N_e$ **do**
- 2: **Shuffle** training set \mathcal{D} ;
- 3: **for** $j = 1, \dots, N_i$ **do**
- 4: **Fetch** mini-batch \mathcal{D}_j with size n_b from \mathcal{D} ;
- 5: $\mathbf{p}_x = f(\mathbf{x}, \Theta), \forall \mathbf{x} \in \mathcal{D}_j$;
- 6: **Obtain** the potential set of ‘clean’ samples $\tilde{\mathcal{D}}_j$ with size d from \mathcal{D}_j based on the loss defined in Eq. (5) using \mathbf{p}_x, \mathbf{y} and $\hat{\mathbf{y}}$ (obtained from $\mathbb{T}^{(n-1)}$);
- 7: **for** $k = 1, \dots, K$ **do**
- 8: **Update** $\mathbf{T}_k^{(n,j)}$ by (1);
- 9: **end for**
- 10: **Obtain** $\ell_{overall}$ by Eq. (5) on $\tilde{\mathcal{D}}_j$;
- 11: **Update** $\Theta = \Theta - \eta \nabla \ell_{overall}$;
- 12: **end for**
- 13: **Update** $\mathbb{T}^{(n)}$ by Eqs. (2) and (3)
- 14: **end for**

Output: Θ

large loss values for both hard samples and the ones with noisy labels. Therefore, the valuable hard samples are often wrongly identified as mislabeled ones and excluded for network training by current robust learning frameworks, which degrades the performance of trained models. We argue that this issue is caused by the insufficient information provided by the one-hot label and propose that the soft label $\hat{\mathbf{y}}$, describing the inter-class relationship, is a potential solution for the problem.

2.1 Noise Ignoring Block

The proposed noise ignoring block (NIB) aims to mitigate negative effects of noisy labels by minimizing the distribution distance between soft labels and instance predictions. The pipeline of our NIB module can be divided into two processes, *i.e.*, probability transition matrix estimation and inter-class correlation (IC) loss calculation, which are alternatively progressed, as presented in Alg. 1. The detailed information of each process is provided in the following:

Probability Transition Matrix. The probability transition matrix is formed by the accumulative soft-labels calculated using ‘clean’ data, which represents the inter-class correlation. Denote $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ as the training set containing N samples from K classes, where $\mathbf{x}_i \in \mathcal{R}^{H \times W \times C}$ is the image (H, W and C are the height, width and number of channels of the image, respectively), and $y_i \in \{1, \dots, K\}$ is the one-hot label. Let \mathbb{T} indicate a $K \times K$ probability transition matrix,¹ which can be written as $\mathbb{T} = [\mathbf{T}_1, \mathbf{T}_2, \dots, \mathbf{T}_K]$, where \mathbf{T}_k ($k \in \{1, \dots, K\}$) is the k -th row of \mathbb{T} . After selecting ‘clean’ samples $\tilde{\mathcal{D}}_j$ from a batch

¹The initial probability transition matrix is denoted as $\mathbb{T}^0 = \mathbf{0}_{K \times K}$ (*i.e.*, a zero matrix).

\mathcal{D}_j ,² the calculation of $\mathbf{T}_k^{(n,j)}$ for class k at current epoch n can be formulated as:

$$\mathbf{T}_k^{(n,j)} = \frac{1}{d_k} \sum_{(\mathbf{x},y) \in \tilde{\mathcal{D}}_{j,k}} \mathbf{p}_x, \quad (1)$$

where $\tilde{\mathcal{D}}_{j,k} = \{(\mathbf{x},y) \mid y = k, \mathbf{x} \in \tilde{\mathcal{D}}_j\}$; \mathbf{p}_x is the prediction (class-wise probability) yielded by the classification network f with parameter Θ for \mathbf{x} ($\mathbf{p}_x = f(\mathbf{x}, \Theta)$); and d_k is the number of samples from class k in the selected ‘clean’ set $\tilde{\mathcal{D}}_j$. After that, the probability transition matrix $\mathbb{T}^{(n,j)}$ for epoch n and batch \mathcal{D}_j is obtained via concatenation:

$$\mathbb{T}^{(n,j)} = [\mathbf{T}_1^{(n,j)}, \mathbf{T}_2^{(n,j)}, \dots, \mathbf{T}_K^{(n,j)}]. \quad (2)$$

The final probability transition matrix of epoch n is obtained by averaging, as defined:

$$\mathbb{T}^{(n)} = \frac{1}{n \times N_i} \sum_{m=1}^n \sum_{j=1}^{N_i} \mathbb{T}^{(m,j)}, \quad (3)$$

where $N_i = \frac{N}{n_b}$ is the number of iterations in one epoch (n_b is the batch size). The probability transition matrix $\mathbb{T}^{(n)}$ is used to calculate the inter-class correlation loss in the next epoch $n+1$ for the ‘clean’ data selection. Concretely, \mathbf{T}_k from \mathbb{T} is used as the accumulative soft label $\hat{\mathbf{y}}$ for class k , *i.e.*, $\hat{\mathbf{y}}_i = \mathbf{T}_k$ for $y_i = k$.

Inter-class Correlation Loss. To distinguish the ‘hard’ samples from the mislabeled ones, we implement an inter-class correlation (IC) loss calculated with the soft labels from the estimated probability transition matrix \mathbb{T} . As previously mentioned, the accumulative soft label $\hat{\mathbf{y}}$ for class k is derived from \mathbf{T}_k of \mathbb{T} . Then, for an input image \mathbf{x}_i , the IC loss ℓ_{IC} can be formulated as:

$$\ell_{IC}(\mathbf{x}_i) = D_{\text{KL}}(\hat{\mathbf{y}}_i \parallel \mathbf{p}_{\mathbf{x}_i}) = \sum_{k=1}^K \hat{y}_i^k \log \frac{\hat{y}_i^k}{p^k(\mathbf{x}_i)}, \quad (4)$$

where $p^k(\mathbf{x}_i)$ is the k -th element of $\mathbf{p}_{\mathbf{x}_i}$; \hat{y}_i^k is the k -th element of $\hat{\mathbf{y}}$; and D_{KL} is the KL divergence measuring the divergence between distributions.

Objective Function. The proposed NIB is a plugin module, which can be easily integrated to the existing robust training framework. Hence, the overall criterion ℓ_{overall} for ‘clean’ sample selection can be written as:

$$\ell_{\text{overall}}(\mathbf{x}_i) = \lambda \cdot \ell_{cls}(\mathbf{x}_i, y_i) + (1 - \lambda) \cdot \ell_{IC}(\mathbf{x}_i, \hat{\mathbf{y}}_i), \quad (5)$$

where ℓ_{cls} is the classification loss adopted by existing robust training approaches, *e.g.*, cross-entropy loss in Co-teaching, and λ is a factor balancing the ℓ_{cls} and ℓ_{IC} (λ is empirically set to 0.6 in our experiments). Consistent to [8, 24], the samples with lower ℓ_{overall} are selected as ‘clean’ data by instance ranking [8, 24]. The ℓ_{overall} is not only used for ‘clean’ data selection, but also for the optimization of classification network, as presented in Alg. 1.

²At the beginning of robust learning, we select ‘clean’ samples based on the cross-entropy loss calculated with the one-hot labels to generate \mathbb{T}^1 . Then, the joint criterion defined in Eq. 5 is adopted for sample selection and matrix update. The number of clean data is consistent to the setting of [8, 24]

Relationship between Label Smoothing and Soft Label. Label smoothing is a technique widely used for the training of deep learning models, wherein one-hot training labels are mixed with uniform label vectors [20]. Empirically, label smoothing has been proven to improve both predictive performance and model calibration [20, 29]. However, the label noises seem to be amplified by the label smoothing, since it is equivalent to injecting symmetric noise to the labels [26]. In contrast, the soft label used in this study is a data-driven label, which is derived from the probability transition matrix estimated from the ‘clean’ samples. In other words, the soft label is a representative of probability distribution corresponding to the class (*vs.* the uniform distribution adopted by label smoothing); hence, no extra label noise is introduced to the dataset. This is also the underlying reason that the proposed inter-class correlation (IC) loss, calculated using the soft label and model prediction, can further benefit the robust learning with noisy labels.

It is worthwhile to mention that soft label has been widely used for different tasks, *e.g.*, semi-supervised learning [2] and knowledge distillation [9]. Concretely, the former one [2] accumulated the network predictions from different training epochs as the supervision signal (soft label) for unlabeled data. The later one [9] used the class-wise soft label for knowledge distillation. To this end, the proposed NIB module can be seen as the combination of the two approaches, which accumulates the network predictions from different training epochs to form a class-wise transition matrix for noisy label rejection. Furthermore, we notice that there are some existing approaches [19] using the transition matrix for robust network training. The related noisy label defencing framework [19] was a typical meta-learning-based approach, which required a meta net to iteratively optimize the parameter of classifier. In contrast, the proposed NIB is a plug-in module without extra network parameters, which is flexible and easy-to-implement.

3 Experiments

In this section, we validate the proposed plugin module (noise ignoring block, NIB) on two publicly available natural and medical image datasets, and present the experimental results.

3.1 Datasets

CIFAR-10. CIFAR-10³ [6] is popularly used for evaluation of noisy labels in the literature [0, 2, 15, 16]. The dataset contains 60,000 images, which can be categorized to ten classes, with a uniform size of 32×32 pixels. The training set and test set consist of 50,000 images and 10,000 images, respectively.

ISIC 2019. Deep-learning-based medical image classification frameworks often suffer from the noisy label problem, due to the different experience levels of annotators (*i.e.*, doctors and radiologists). In this paper, the widely-used ISIC 2019 dataset⁴ is adopted to verify the effectiveness of our NIB module for medical image classification. The ISIC 2019 dataset [22] is from the challenge of prediction of eight skin disease categories with dermoscopic images, including melanoma (MEL), melanocytic nevus (NV), basal cell carcinoma (BCC), actinic keratosis (AK), benign keratosis (BKL), dermatofibroma (DF), vascular lesion (VASC), and

³<https://www.cs.toronto.edu/~kriz/cifar.html>

⁴<https://challenge2019.isic-archive.com/>

squamous cell carcinoma (SCC). The original ISIC dataset is highly imbalanced between classes [9]. To alleviate the potential effect of imbalance in continual experiment, we randomly sample 628 images from each class (Note that we take all images from two classes with fewer than 628 images), consistent to [9]. A total of 4,260 images are randomly divided into a training and a test set according to the ratio of 80:20.

3.2 Experimental Settings

Label Shuffling. Following [14, 16], we shuffle the labels of the training set by a noise transition matrix Q , where $Q_{ij} = Pr[\tilde{y} = j | y = i]$ denotes the probability of flipping class i to j . The widely-used structures of Q , *i.e.*, symmetry flipping [15, 23] and pair flipping [9], are adopted in our study. Note that, consistent to [9, 15, 23], we validate our NIB module with different noise ratios, denoting as ‘symmetry-10%’, ‘symmetry-20%’, ‘symmetry-40%’ and ‘pair-10%’. For example, the ‘symmetry-10%’ represents that 10% of the labels have been symmetrically flipped to be noisy labels.

Implementation & Evaluation Criterion. The proposed NIB is implemented using the PyTorch toolbox. All the frameworks use the same backbone architectures, *i.e.*, a 9-layer CNN network architecture is adopted for CIFAR-10, while we utilize the ResNet-18 and DenseNet-169 as the backbone for ISIC 2019, due to the larger image size and classification complexity of the ISIC dataset. The Adam optimizer (momentum=0.9) is used for network optimization with an initial learning rate of 0.001. The batch size is set to 128 and 64 for CIFAR-10 and ISIC dataset, respectively. The images from ISIC 2019 dataset are resized to 224×224 pixels for network processing. The average classification accuracy (ACC) on the test set is adopted as the metric to evaluate the performance of robust learning with noisy labels. We run 200 epochs in total and calculate ACC over the last 10 epochs.

3.3 Performance Evaluation

In this section, we present the experimental results on the two publicly available datasets, *i.e.*, CIFAR-10 and ISIC 2019. The state-of-the-art robust training approaches, *i.e.*, Co-teaching [9] and JoCoR [24], are involved as baselines.

CIFAR-10. The average classification accuracy on the CIFAR-10 test set yielded by the classification networks trained with different strategies is listed in Table 1. It can be observed that the proposed NIB module consistently improves the classification accuracy of existing robust learning approaches under different ratios of noises. Concretely, even when 40% of the labels are symmetrically flipped (*i.e.*, symmetry-40%), our NIB can still significantly increase the test ACC for Co-teaching and JoCoR by margins of +4.73% and +4.55%, respectively, which demonstrates the outstanding robustness of our NIB to label noises.

As shown in Table 1, we also conduct an ablation study on CIFAR-10 to evaluate the performance of models only using ℓ_{IC} for clean data selection. As the sample-wise class relationship information may differ from the class-wise one, which results in a larger value of ℓ_{IC} , some clean data is wrongly rejected while only using ℓ_{IC} for clean data selection. Hence, the classification accuracy of models only using ℓ_{IC} significantly degrades on CIFAR-10, compared to the original learning-to-learn paradigms using ℓ_{cls} . The experimental results also reveal the mechanism underlying our framework—the clean data is easily filtered by ℓ_{cls} ,

Table 1: Average classification accuracy (%) on the CIFAR-10 test set.

Method↓,Noise Type→	Clean	symmetry-10%	symmetry-20%	symmetry-40%	pair-10%
Co-teaching [9]	89.13	85.06	82.24	78.15	85.29
Co-teaching+NIB	90.87	88.44	86.69	82.88	88.33
Co-teaching+ ℓ_{IC} -only	90.13	87.80	86.07	76.92	86.34
Co-teaching+ ℓ_{DY} [10]	89.04	85.23	83.32	79.93	83.80
JoCoR [24]	89.30	84.04	81.61	77.40	84.40
JoCoR+NIB	90.70	87.81	85.88	81.95	87.61
JoCoR+ ℓ_{IC} -only	90.69	79.39	68.88	64.84	70.73
JoCoR+ ℓ_{DY} [10]	90.16	82.36	82.44	79.67	84.20

Table 2: Average test ACC (%) on the ISIC 2019 dataset.

Backbone	Method↓,Noise Type→	Clean	Symmetry-10%	Symmetry-20%	Symmetry-40%	Pair-10%
ResNet-18	Co-teaching [9]	62.00	58.47	55.43	47.81	58.40
	Co-teaching+NIB	64.40	60.08	59.37	49.33	60.08
	JoCoR [24]	62.11	57.78	56.71	47.62	58.67
	JoCoR+NIB	63.43	59.77	59.17	49.65	60.43
DenseNet-169	Co-teaching	70.11	65.67	63.29	48.87	65.16
	Co-teaching+NIB	71.32	68.86	63.74	56.12	68.11
	JoCoR	68.46	62.93	56.08	49.92	67.24
	JoCoR+NIB	70.99	65.79	60.07	53.42	68.22

Table 3: Classification accuracy (%) on the Clothing1M test set.

Method	Co-teaching [9]	Co-teaching+NIB	JoCoR [24]	JoCoR+NIB
Accuracy	69.74	71.02	70.25	71.28

while ℓ_{IC} is proposed to further separate the hard samples from the ones with noisy labels. In addition, we integrate the state-of-the-art approach (dynamic bootstrapping, ℓ_{DY}) [10] to existing learning-to-learn paradigms for comparison. As shown in Table 1, the frameworks using our NIB module surpass the ones using ℓ_{DY} by a large margin.

ISIC 2019. To validate the effectiveness of our NIB module on medical images, we conduct experiments on the ISIC 2019 dataset. The test ACC on the ISIC 2019 dataset are presented in Table 2. A similar trend to CIFAR-10 is observed—using our NIB module, the ACCs of existing approaches on the ISIC 2019 test set are consistently improved. The Co-teaching + NIB achieves the best test ACC under most noise ratios. Furthermore, we notice that the proposed NIB achieves significant improvements for not only ResNet-18, but also the ultra-deep DenseNet-169, even under the large noise ratio (40%).

Application on Real-world Dataset. An experiment is conducted on the Clothing1M⁵ [24], a real-world noisy dataset, to further validate the effectiveness of our NIB module. The dataset consists of one million images captured from online shopping websites. The label (14 classes) for each image is generated by extracting tags from the surrounding texts and keywords, which are naturally noisy. The Clothing1M dataset is separated to training and test sets according to the protocol [24]. In the experiment, frameworks are trained with noisy training set, and then evaluated on the test set using clean labels. The evaluation results are shown in Table 3. It can be observed that the proposed NIB module consistently boosts the classification accuracy of existing learning-to-learn paradigms, *i.e.*, improvements of +1.28% and +1.03% are yielded by our NIB module to Co-teaching and JoCoR, respec-

⁵https://github.com/Cysu/noisy_label

Table 4: Average label precision (%) on the CIFAR-10 over the last 10 epochs.

Method↓, Noise Type→	Symmetry-10%	Symmetry-20%	Symmetry-40%	Pair-10%
Co-teaching [3]	95.32	92.74	87.46	94.94
Co-teaching+NIB	96.01	93.43	88.28	95.73
JoCoR [24]	95.76	93.43	89.06	95.29
JoCoR+NIB	96.23	93.90	89.57	96.01

Table 5: Label precision (%) of the selected samples on the ISIC 2019 dataset.

Backbone	Method↓, Noise Type→	Symmetry-10%	Symmetry-20%	Symmetry-40%	Pair-10%
ResNet-18	Co-teaching	94.33	90.58	80.76	94.16
	Co-teaching+NIB	94.17	90.79	82.03	94.31
	JoCoR	94.69	90.77	81.64	95.26
	JoCoR+NIB	94.55	90.74	81.04	94.65
DenseNet-169	Co-teaching	94.76	90.41	78.42	94.17
	Co-teaching+NIB	95.30	91.18	82.06	94.71
	JoCoR	94.98	90.65	79.84	94.63
	JoCoR+NIB	95.62	91.35	81.73	95.27

tively. The experimental results demonstrate the effectiveness of the proposed NIB module for the realistic application.

Label Precision. Apart from ACC, the label precision is also measured and reported on CIFAR-10 and ISIC 2019 datasets to assess the pure ratio of selected samples:

$$\text{label precision} = \frac{\text{the number of selected samples with correct labels}}{\text{the total number of selected samples}}. \quad (6)$$

Intuitively, the higher label precision means the fewer noisy instances mistakenly identified as ‘clean’ by the approach, *i.e.*, the better robustness to noisy labels.

The label precision of different robust learning approaches on CIFAR-10 and ISIC 2019 are listed in Table 4 and 5, respectively. The label precision of the frameworks using the proposed NIB module is observed to consistently surpass the original ones on the CIFAR-10 dataset. For the ISIC 2019 dataset, as stated in recent studies [28], the ultra-deep networks (*e.g.*, DenseNet-169) more easily overfit to the noises and severely suffer from the noisy label problem, compared to the shallow ones. Table 2 reveals that the proposed NIB module can significantly alleviate such an overfitting problem, *i.e.*, further improving the test accuracy of DenseNet-169 even under the large noise ratio (40%). The underlying reason is the higher label precision achieved by our NIB module, as shown in Table 5. The samples selected by the frameworks with our NIB is ‘cleaner’ than the conventional ones, *i.e.*, a higher label precision is achieved by Co-teaching + NIB and JoCoR + NIB than that of Co-teaching and JoCoR, which accordingly benefits the training of ultra-deep networks on the classification task. On the other hand, we notice that although the label precision of JoCoR + NIB is slightly lower than JoCoR in Table 5 for ResNet-18, the test accuracy is still improved by using our NIB. This is because the JoCoR + NIB is joint optimized by ℓ_{cls} and ℓ_{IC} . The combination of these two losses refines the direction of optimization, compared to the ℓ_{cls} -only JoCoR. The curves of test accuracy and label precision *vs.* epochs during the whole training process of ResNet-18 and DenseNet-169 on the ISIC dataset are presented in Fig. 1. It can be observed that the frameworks using our NIB consistently outperform the original ones on both test accuracy and label precision, demonstrating the generalization of the proposed NIB module.

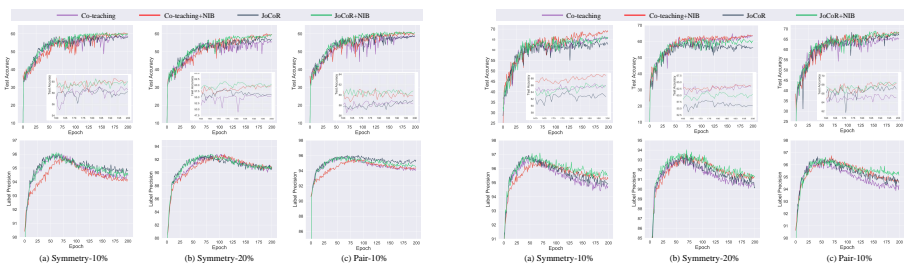


Figure 1: Training curves of **ResNet-18** (left) and **DenseNet-169** (right) on the ISIC dataset, respectively. Top: test ACC (%) vs. epochs; bottom: label precision (%) vs. epochs. The figures on the top row reveal that the test accuracy of Co-teaching/JoCoR + NIB is consistently higher than the original ones in the late stage of training (*i.e.*, after 160 epochs of training). For the clear observation, please refer to the arXiv version.

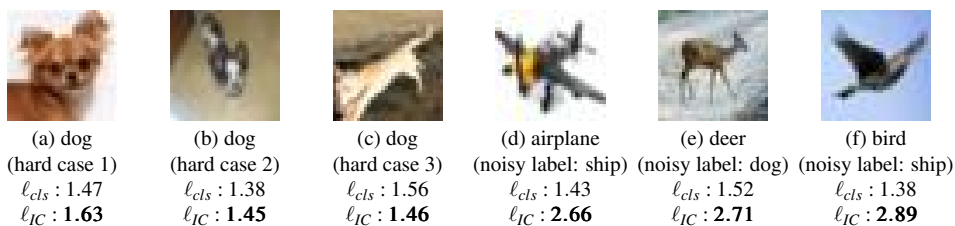


Figure 2: Examples of hard and mislabeled cases from CIFAR-10 identified by our NIB.

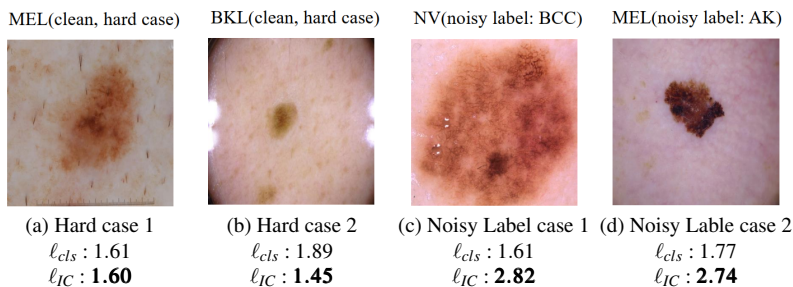


Figure 3: Examples of hard and mislabeled cases from ISIC 2019 identified by our NIB.

3.4 Analysis: Hard Samples vs. Samples with Noisy Labels

In this section, to validate the effectiveness of inter-class correlation loss for hard sample identification, we show the detected hard samples and samples with noisy labels from CIFAR-10 and ISIC 2019 in Fig. 2 and 3, respectively. The ℓ_{cls} and ℓ_{IC} calculated with one-hot label \mathbf{y} and accumulative soft label $\hat{\mathbf{y}}$, respectively, are also listed under each sample. We observe that the ℓ_{cls} of hard and mislabeled cases is close, *i.e.*, around 1.5 and 1.7 on CIFAR-10 and ISIC 2019, respectively. Therefore, the existing ℓ_{cls} -only robust learning frameworks cannot distinguish them and will assign them to the same category (either ‘clean’ or ‘noisy’). Thus, more samples with noise labels are involved for classifier training if treating them as ‘clean’ or the hard samples are over filtered, both leading to the degradation of classification performance. In contrast, our inter-class correlation loss ℓ_{IC} excellently distinguishes the

hard samples and the ones with noisy labels (around 1.5 vs. around 2.8) on both CIFAR-10 and ISIC datasets. Hence, the framework using our ℓ_{IC} can simultaneously benefit the classifier from hard samples and maintain the robustness to noise labels, which results in the improvements presented in the previous section.

Effectiveness of Hard Samples for Network Training. As aforementioned, the hard samples falling around the decision boundary in the latent space may provide rich information for decision boundary refinement and should be paid more attention during network training. Such a claim has been verified by the existing studies [8, 18]. In our experiments, the experimental results explicitly demonstrate the effectiveness of hard samples for network training. The original learning-to-learn paradigms (Co-teaching and JoCoR) exclude the hard samples illustrated in Fig. 2 and 3 from network training, due to the large values of ℓ_{cls} . Note that ℓ_{cls} is used as the overall loss $\ell_{overall}$ in the original paradigms for sample selection and the average ℓ_{cls} and ℓ_{IC} for clean samples are 1.13 and 1.37, respectively. In contrast, our ℓ_{IC} decreases the $\ell_{overall}$ (i.e., $\ell_{cls} + \ell_{IC}$) of hard samples; therefore, those samples can be separated from noisy ones and included for network training. As shown in Table 1 and 2, the classification accuracy is significantly improved by adding our ℓ_{IC} loss to $\ell_{overall}$, which validates the effectiveness of including hard samples for network training.

4 Conclusion

In this paper, we proposed a simple-yet-effective plugin module, namely noise ignoring block (NIB), which can be easily integrated to the existing learning-to-learn instance selection paradigms, to distinguish the hard samples from the mislabeled ones and further boost the performance of robust learning with noisy labels. Extensive experiments were conducted on natural and medical image datasets (CIFAR-10 and ISIC 2019). The experimental results showed that our NIB module consistently improved the performances of the state-of-the-art robust training methods, e.g., Co-teaching and JoCoR.

Limitation and Future Work. Currently, the proposed NIB module is only integrated to the framework designed for image classification. However, existing studies began to focus on the multi-label classification task, e.g., ReLabel [27]. In this regard, we plan to adjust the proposed NIB module to more image processing tasks, such as multi-label classification and semantic segmentation, in the future work.

Acknowledgements

This work was funded by the Key-Area Research and Development Program of Guangdong Province, China (No.2020B090923003 and No.2018B010111001), National Natural Science Foundation of China (Grant No. 52075285), National Key R&D Program of China (2018YFC-2000702) and the Scientific and Technical Innovation 2030-“New Generation Artificial Intelligence” Project (No. 2020AAA0104100).

References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel O'Connor, and Kevin Mcguinness. Unsupervised label noise modeling and loss correction. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.
- [2] Jacob Goldberger and Ehud Ben-Reuven. Training deep neural-networks using a noise adaptation layer. In *Proceedings of the 5th International Conference on Learning Representation*, 2016.
- [3] 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. *arXiv preprint arXiv:1804.06872*, 2018.
- [4] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2014.
- [5] 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, 2018.
- [6] Alex Krizhevsky. *Learning multiple layers of features from tiny images*, 2009.
- [7] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *International Conference on Learning Representations (ICLR)*, 2017.
- [8] Yuexiang Li, Dong Wei, Jiawei Chen, Shilei Cao, Hongyu Zhou, Yanchun Zhu, Jianrong Wu, Lan Lan, Wenbo Sun, Tianyi Qian, Kai Ma, Haibo Xu, and Yefeng Zheng. Efficient and effective training of covid-19 classification networks with self-supervised dual-track learning to rank. *IEEE Journal of Biomedical and Health Informatics*, 24(10):2787–2797, 2020.
- [9] Zhuoyun Li, Changhong Zhong, Ruixuan Wang, and Wei-Shi Zheng. Continual learning of new diseases with dual distillation and ensemble strategy. In *Medical Image Computing and Computer Assisted Intervention*, pages 169–178, 2020. ISBN 978-3-030-59710-8.
- [10] 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.
- [11] Eran Malach and Shai Shalev-Shwartz. Decoupling “when to update” from “how to update”. *arXiv preprint arXiv:1706.02613*, 2017.
- [12] Iacopo Masi, Yue Wu, Tal Hassner, and Prem Natarajan. Deep face recognition: A survey. In *2018 31st SIBGRAPI Conference on Graphics, Patterns and Images*, pages 471–478. IEEE, 2018.
- [13] Aditya Menon, Brendan Van Rooyen, Cheng Soon Ong, and Bob Williamson. Learning from corrupted binary labels via class-probability estimation. In *International Conference on Machine Learning*, pages 125–134, 2015.

- [14] 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.
- [15] Li Quan, Yan Li, Xiaoyi Chen, and Ni Zhang. An effective data refinement approach for upper gastrointestinal anatomy recognition. In *International Conference on Medical Image Computing and Computer Assisted Intervention*, pages 43–52, 2020.
- [16] Scott Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. Training deep neural networks on noisy labels with bootstrapping. *arXiv preprint arXiv:1412.6596*, 2014.
- [17] Tomoya Sakai, Marthinus Christoffel Plessis, Gang Niu, and Masashi Sugiyama. Semi-supervised classification based on classification from positive and unlabeled data. In *International Conference on Machine Learning*, pages 2998–3006, 2017.
- [18] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 761–769, 2016.
- [19] Jun Shu, Qian Zhao, Zongben Xu, and Deyu Meng. Meta transition adaptation for robust deep learning with noisy labels. *arXiv preprint arXiv:2006.05697*, 2020.
- [20] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016.
- [21] Mingxing Tan, Ruoming Pang, and Quoc V Le. EfficientDet: Scalable and efficient object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10781–10790, 2020.
- [22] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific Data*, 5(1):1–9, 2018.
- [23] Brendan Van Rooyen, Aditya Krishna Menon, and Robert C Williamson. Learning with symmetric label noise: The importance of being unhinged. *arXiv preprint arXiv:1505.07634*, 2015.
- [24] 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*, pages 13726–13735, 2020.
- [25] 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 (CVPR)*, 2015.
- [26] Lingxi Xie, Jingdong Wang, Zhen Wei, Meng Wang, and Qi Tian. Disturblabel: Regularizing CNN on the Loss Layer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4753–4762, 2016.

- [27] Sangdoon Yun, Seong Joon Oh, Byeongho Heo, Dongyoon Han, Junsuk Choe, and Sanghyuk Chun. Re-labeling imagenet: From single to multi-labels, from global to localized labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2340–2350, June 2021.
- [28] C Zhang, S Bengio, M Hardt, B Recht, and O Vinyals. Understanding deep learning requires rethinking generalization. In *Proceedings of the International Conference on Learning Representation*, 2017.
- [29] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8697–8710, 2018.