

# Symmetric Cross Entropy for Robust Learning with Noisy Labels

Yisen Wang<sup>1\*</sup>† Xingjun Ma<sup>2\*</sup>† Zaiyi Chen<sup>3</sup> Yuan Luo<sup>1</sup> Jinfeng Yi<sup>4</sup> James Bailey<sup>2</sup>  
<sup>1</sup>Shanghai Jiao Tong University <sup>2</sup>The University of Melbourne <sup>3</sup>Cainiao AI <sup>4</sup>JD AI

## Abstract

Training accurate deep neural networks (DNNs) in the presence of noisy labels is an important and challenging task. Though a number of approaches have been proposed for learning with noisy labels, many open issues remain. In this paper, we show that DNN learning with Cross Entropy (CE) exhibits overfitting to noisy labels on some classes (“easy” classes), but more surprisingly, it also suffers from significant under learning on some other classes (“hard” classes). Intuitively, CE requires an extra term to facilitate learning of hard classes, and more importantly, this term should be noise tolerant, so as to avoid overfitting to noisy labels. Inspired by the symmetric KL-divergence, we propose the approach of **Symmetric cross entropy Learning (SL)**, boosting CE symmetrically with a noise robust counterpart Reverse Cross Entropy (RCE). Our proposed SL approach simultaneously addresses both the under learning and overfitting problem of CE in the presence of noisy labels. We provide a theoretical analysis of SL and also empirically show, on a range of benchmark and real-world datasets, that SL outperforms state-of-the-art methods. We also show that SL can be easily incorporated into existing methods in order to further enhance their performance.

## 1. Introduction

Modern deep neural networks (DNNs) are often highly complex models that have hundreds of layers and millions of trainable parameters, requiring large-scale datasets with clean label annotations such as ImageNet [2] for proper training. However, labeling large-scale datasets is a costly and error-prone process, and even high-quality datasets are likely to contain noisy (incorrect) labels. Therefore, training accurate DNNs in the presence of noisy labels has become a task of great practical importance in deep learning.

Recently, several works have studied the dynamics of DNN learning with noisy labels. Zhang *et.al* [28] argued that DNNs exhibit memorization effects whereby they first memorize the training data for clean labels and then subsequently memorize data for the noisy labels. Similar findings

\*Equal contribution. Part of work was done at JD AI Research.

†Correspondence to: Yisen Wang (eewangyisen@gmail.com) and Xingjun Ma (xingjun.ma@unimelb.edu.au).

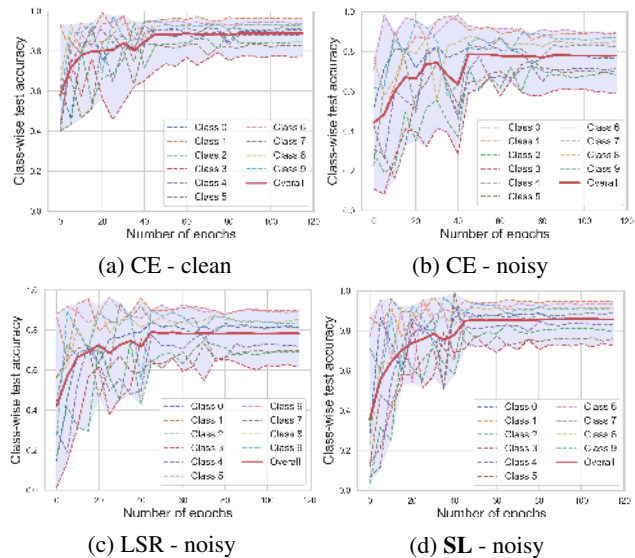


Figure 1: The class-wise test accuracy of an 8-layer CNN on CIFAR-10 trained by (a) CE on clean labels with class-biased phenomenon, (b) CE on 40% symmetric/uniform noisy labels with amplified class-biased phenomenon and under learning on hard classes (e.g., class 3), (c) LSR under the same setting to (b) with under learning on hard classes still existing, (d) our proposed SL under the same setting to (b) exhibiting overall improved learning on all classes.

are also reported in [1] that DNNs first learn clean and easy patterns and eventually memorize the wrongly assigned labels. Further evidence is provided in [13] that DNNs first learn simple representations via subspace dimensionality compression and then overfit to noisy labels via subspace dimensionality expansion. Different findings are reported in [19], where DNNs with a specific activation function (*i.e.*, tanh) undergo an initial label fitting phase then a subsequent representation compression phase where the overfitting starts. Despite these important findings, a complete understanding of DNN learning behavior, particularly their learning process for noisy labels, remains an open question.

In this paper, we provide further insights into the learning procedure of DNNs by investigating the learning dynamics across classes. While Cross Entropy (CE) loss is the most commonly used loss for training DNNs, we have found that DNN learning with CE can be **class-biased**:

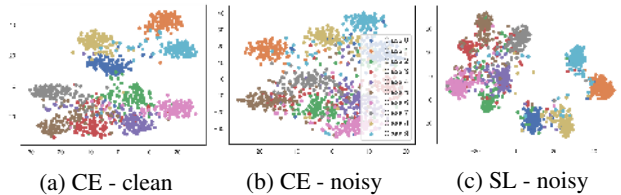


Figure 2: Visualization of learned representations on CIFAR-10 using t-SNE 2D embeddings of deep features at the last second dense layer with (a) CE on clean labels, (b) CE on 40% symmetric noisy labels, (c) the proposed SL on the same setting to (b).

some classes (“easy” classes) are easy to learn and converge faster than other classes (“hard” classes). As shown in Figure 1a, even when labels are clean, the class-wise test accuracy spans a wide range during the entire training process. As further shown in Figure 1b, this phenomenon is amplified when training labels are noisy: whilst easy classes (e.g., class 6) already overfit to noisy labels, hard classes (e.g., class 3) still suffer from significant *under learning* (class accuracy significantly lower than clean label setting). Specifically, class 3 (bottom curve) only has an accuracy of  $\sim 60\%$  at the end, considerably less than the  $>90\%$  accuracy of class 6 (top curve). Label Smoothing Regularization (LSR) [21, 17] is a widely known technique to ease overfitting issues, as shown in Figure 1c, which still exhibits significant under learning on hard classes. Comparing the overall test accuracy (solid red curve) in Figure 1, a low test accuracy (under learning) on hard classes is a barrier to high overall accuracy. This is a different finding from previous belief that poor performance is simply caused by overfitting to noisy labels. We also visualize the learned representations for the noisy label case in Figure 2b: some clusters are learned comparably well to those learned with clean labels (Figure 2a), while some other clusters do not have clear separated boundaries.

Intuitively, CE requires an extra term to improve its learning on hard classes, and more importantly, this term needs to be tolerant to label noise. Inspired by the symmetric KL-divergence, we propose such a noise tolerant term, namely Reverse Cross Entropy (RCE), which combined with CE forms the basis of the approach Symmetric cross entropy Learning (SL). SL not only promotes sufficient learning (class accuracy close to clean label setting) of hard classes, but also improves the robustness of DNNs to noisy labels. As a preview of this, we can inspect the improved learning curves of class-wise test accuracy and representations in Figure 1d and 2c. Under the same 40% noise setting, the variation of class-wise test accuracy has been narrowed by SL to 20% with 95% the highest and 75% the lowest (Figure 1d), and the learned representations are of better quality with more separated clusters (Figure 2c), both of which are very close to the clean settings.

Compared to existing approaches that often involve architectural or non-trivial algorithmic modifications, SL is

extremely simple to use. It requires minimal intervention to the training process and thus can be straightforwardly incorporated into existing models to further enhance their performance. In summary, our main contributions are:

- We provide insights into the class-biased learning procedure of DNNs with CE loss and find that the under learning problem of hard classes is a key bottleneck for learning with noisy labels.
- We propose a Symmetric Learning (SL) approach, to simultaneously address the hard class under learning problem and the noisy label overfitting problem of CE. We provide both theoretical analysis and empirical understanding of SL.
- We empirically demonstrate that SL can achieve better robustness than state-of-the-art methods, and can be also easily incorporated into existing methods to significantly improve their performance.

## 2. Related Work

Different approaches have been proposed to train accurate DNNs with noisy labels, and they can be roughly divided into three categories: 1) label correction methods, 2) loss correction methods, and 3) refined training strategies.

The idea of label correction is to improve the quality of the raw labels. One common approach is to correct noisy labels to their true labels via a clean label inference step using complex noise models characterized by directed graphical models [26], conditional random fields [23], neural networks [11, 24] or knowledge graphs [12]. These methods require the support from extra clean data or an expensive detection process to estimate the noise model.

Loss correction methods modify the loss function for robustness to noisy labels. One approach is to model the noise transition matrix that defines the probability of one class changed to another class [5]. Backward [16] and Forward [16] are two such correction methods that use the noise transition matrix to modify the loss function. However, the ground-truth noise transition matrix is not always available in practice, and it is also difficult to obtain accurate estimation [5]. Work in [4, 20] augments the correction architecture by adding a linear layer on top of the neural network. Bootstrap [18] uses a combination of raw labels and their predicted labels. There is also research that defines noise robust loss functions, such as Mean Absolute Error (MAE) [3], but a challenge is that training a network with MAE is slow due to gradient saturation. Generalized Cross Entropy (GCE) loss [29] applies a Box-Cox transformation to probabilities (power law function of probability with exponent  $q$ ) and can behave like a weighted MAE. Label Smoothing Regularization (LSR) [21, 17] is another technique using soft labels in place of one-hot labels to alleviate overfitting to noisy labels.

Refined training strategies design new learning paradigms for noisy labels. MentorNet [8, 27] supervises the training of a StudentNet by a learned sample weighting scheme in favor of probably correct labels. Decoupling training strategy [15] trains two networks simultaneously, and parameters are updated when their predictions disagree. Co-teaching [6] maintains two networks simultaneously during training, with one network learning from the other network’s most confident samples. These studies all require training of an auxiliary network for sample weighting or learning supervision. D2L [13] uses subspace dimensionality adapted labels for learning, paired with a training process monitor. The iterative learning framework [25] iteratively detects and isolates noisy samples during the learning process. The joint optimization framework [22] updates DNN parameters and labels alternately. These methods either rely on complex interventions into the learning process, which may be challenging to adapt and tune, or are sensitive to hyperparameters like the number of training epochs and learning rate.

### 3. Weakness of Cross Entropy

We begin by analyzing the Cross Entropy (CE) and its limitations for learning with noisy labels.

#### 3.1. Preliminaries

Given a  $K$ -class dataset  $\mathcal{D} = \{(\mathbf{x}, y)^{(i)}\}_{i=1}^n$ , with  $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d$  denoting a sample in the  $d$ -dimensional input space and  $y \in \mathcal{Y} = \{1, \dots, K\}$  its associated label. For each sample  $\mathbf{x}$ , a classifier  $f(\mathbf{x})$  computes its probability of each label  $k \in \{1, \dots, K\}$ :  $p(k|\mathbf{x}) = \frac{e^{z_k}}{\sum_{j=1}^K e^{z_j}}$ , where  $z_j$  are the logits. We denote the ground-truth distribution over labels for sample  $\mathbf{x}$  by  $q(k|\mathbf{x})$ , and  $\sum_{k=1}^K q(k|\mathbf{x}) = 1$ . Consider the case of a single ground-truth label  $y$ , then  $q(y|\mathbf{x}) = 1$  and  $q(k|\mathbf{x}) = 0$  for all  $k \neq y$ . The cross entropy loss for sample  $\mathbf{x}$  is:

$$\ell_{ce} = - \sum_{k=1}^K q(k|\mathbf{x}) \log p(k|\mathbf{x}). \quad (1)$$

#### 3.2. Weakness of CE under Noisy Labels

We now highlight some weaknesses of CE for DNN learning with noisy labels, based on empirical evidence on CIFAR-10 dataset [9] (10 classes of natural images). To generate noisy labels, we randomly flip a correct label to one of the other 9 incorrect labels uniformly (e.g., symmetric label noise), and refer to the portion of incorrect labels as the noise rate. The network used here is an 8-layer convolutional neural network (CNN). Detailed experimental settings can be found in Section 5.1.

We first explore in more detail the class-biased phenomenon shown in Figure 1a and 1b, focusing on three distinct learning stages: early (the 10-th epoch), middle (the

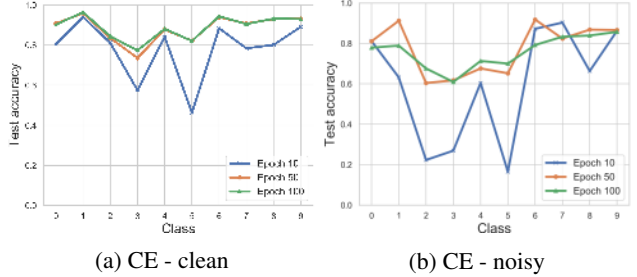


Figure 3: The class-wise test accuracy at epoch 10, 50 and 100 (120 epochs in total) trained by CE loss on CIFAR-10 with (a) clean labels or (b) 40% symmetric noisy labels.

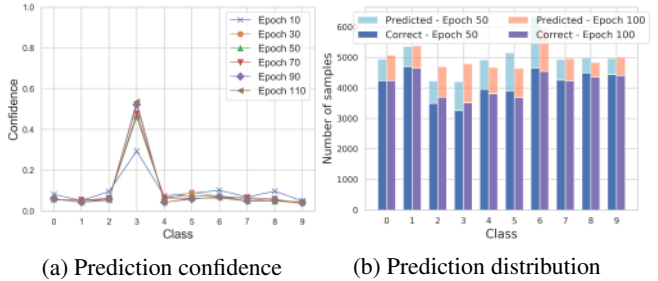


Figure 4: Intermediate results of CE loss on CIFAR-10 with 40% symmetric noisy labels. (a) Average confidence of the **clean portion** of class 3 samples. (b) The true positive samples (*correct*) out of predictions (*predicted*) for each class.

50-th epoch) and later (the 100-th epoch) stages, with respect to total 120 epochs of training. As illustrated in Figure 3, CE learning starts in a highly class-biased manner (the blue curves) for both clean labels and 40% noisy labels. This is because patterns inside of samples are intrinsically different. For clean labels, the network eventually manages to learn all classes uniformly well, reflected by the relatively flat accuracy curve across classes (the green curve in Figure 3a). However, for noisy labels, the class-wise test accuracy varies significantly across different classes, even at the later stage (the green curve in Figure 3b). In particular, the network struggles to learn hard classes (e.g., class 2/3) with up to a 20% gap to the clean setting, whereas some easy classes (e.g., class 1/6) are better learned and already start overfitting to noisy labels (accuracy drops from epoch 50 to 100). It appears that the under learning of hard classes is a major cause for the overall performance degradation, due to the fact that the accuracy drop caused by overfitting is relatively small.

We further investigate the reason behind the under learning of CE on hard classes from the perspective of representations. Due to their high similarities in representations to some other classes (see the red cluster for class 3 in Figure 2a), the predictions for hard class examples are likely to assign a relatively large probability to those similar classes. Under the noisy label scenario, class 3 has become even more scattered into other classes (red cluster in Figure 2b). As a consequence, no visible cluster was learned by CE,

even though there are still 60% correct labels in this scenario. Further delving into these 60% **clean portion** of class 3 samples, we show, in Figure 4a, their prediction confidence output of the neural network. Although the confidence at class 3 is the highest, it is only around 0.5, while for the other classes, the confidence is around 0.05 or 0.1 which is actually a relatively high value and an indication of insufficient learning of class 3 even on the clean labeled part. Another evidence of under learning can be obtained from Figure 4b, where hard classes (e.g., class 2/3) have fewer true positive samples throughout intermediate stages of learning.

Clearly, CE by itself is not sufficient for learning of hard classes, especially under the noisy label scenario. We note that this finding sheds new insights into DNN learning behavior under label noise, and differs from previous belief that DNNs overfit to all classes in general [1, 28]. In the next section, we propose a symmetric learning approach that can address both the hard class under learning and noisy label overfitting problems of CE.

## 4. Symmetric Cross Entropy Learning

In this section, we propose Symmetric cross entropy Learning (SL), an approach that strikes a balance between sufficient learning and robustness to noisy labels. We also provide theoretical analysis about the formulation and behavior of SL.

### 4.1. Definition

Given two distributions  $q$  and  $p$ , the relation between the cross entropy (denoted as  $H(q, p)$ ) and the KL-divergence (denoted as  $KL(q||p)$ ) is:

$$KL(q||p) = H(q, p) - H(q), \quad (2)$$

where  $H(q)$  is the entropy of  $q$ . In the context of classification,  $q = q(k|\mathbf{x})$  is the ground truth class distribution conditioned on sample  $\mathbf{x}$ , whilst  $p = p(k|\mathbf{x})$  is the predicted distribution over labels by the classifier  $f$ . From the perspective of KL-divergence, classification is to learn a prediction distribution  $p(k|\mathbf{x})$  that is close to the ground truth distribution  $q(k|\mathbf{x})$ , which is to minimize the KL-divergence  $KL(q||p)$  between the two distributions\*.

In information theory, given a true distribution  $q$  and its approximation  $p$ ,  $KL(q||p)$  measures the penalty on encoding samples from  $q$  using code optimized for  $p$  (penalty in the number of extra bits required). In the context of noisy labels, we know that  $q(k|\mathbf{x})$  does not represent the true class distribution, instead  $p(k|\mathbf{x})$  can reflect the true distribution to a certain extent. Thus, in addition to taking  $q(k|\mathbf{x})$  as the ground truth, we also need to consider the other direction of KL-divergence, that is  $KL(p||q)$ , to punish coding samples

\*In practice, the  $H(q(k|\mathbf{x}))$  term is a constant for a given class distribution and therefore omitted from Eq. (2) giving the CE loss in Eq. (1).

that come from  $p(k|\mathbf{x})$  when using a code for  $q(k|\mathbf{x})$ . The symmetric KL-divergence is:

$$SKL = KL(q||p) + KL(p||q). \quad (3)$$

Transferring this symmetric idea from KL-divergence to cross entropy gives us the *Symmetric Cross Entropy* (SCE):

$$SCE = CE + RCE = H(q, p) + H(p, q), \quad (4)$$

where  $RCE = H(p, q)$  is the reverse version of  $H(q, p)$ , namely, *Reverse Cross Entropy*. The RCE loss for a sample  $\mathbf{x}$  is:

$$\ell_{rce} = - \sum_{k=1}^K p(k|\mathbf{x}) \log q(k|\mathbf{x}). \quad (5)$$

The sample-wise SCE loss can then be defined as:

$$\ell_{sce} = \ell_{ce} + \ell_{rce}. \quad (6)$$

While the RCE term is noise tolerant as will be proved in Section 4.2, the CE term is not robust to label noise [3]. However, CE is useful for achieving good convergence [29], which will be verified empirically in Section 5. Towards more effective and robust learning, we propose a flexible symmetric learning framework with the use of two decoupled hyperparameters (e.g.,  $\alpha$  and  $\beta$ ), with  $\alpha$  on the overfitting issue of CE while  $\beta$  for flexible exploration on the robustness of RCE. Formally, the SL loss is:

$$\ell_{sl} = \alpha \ell_{ce} + \beta \ell_{rce}. \quad (7)$$

As the ground truth distribution  $q(k|\mathbf{x})$  is now inside of the logarithm in  $\ell_{rce}$ , this could cause computational problem when labels are one-hot: zero values inside the logarithm. To solve this issue, we define  $\log 0 = A$  (where  $A < 0$  is some constant), which shortly will be proved useful for the robustness of  $\ell_{rce}$  in Theorem 1. This technique is similar to the clipping operation implemented by most deep learning frameworks. Compared with another option label smoothing technique, our approach introduces less bias into the model (negligible bias (in the view of training) at finite number of points like  $q(k|x) = 0$  but no bias at  $q(k|x) = 1$ ). Note that, the effect of  $\beta$  on RCE can be reflected by different settings of  $A$  (refer to Eq. (4.3)).

### 4.2. Theoretical Analysis

**Robustness analysis:** In the following, we will prove that the RCE loss  $\ell_{rce}$  is robust to label noise following [3]. We denote the noisy label of  $\mathbf{x}$  as  $\hat{y}$ , in contrast to its true label  $y$ . Given any classifier  $f$  and loss function  $\ell_{rce}$ , we define the *risk* of  $f$  under clean labels as  $R(f) = \mathbb{E}_{\mathbf{x}, y} \ell_{rce}$ , and the *risk* under label noise rate  $\eta$  as  $R^\eta(f) = \mathbb{E}_{\mathbf{x}, \hat{y}} \ell_{rce}$ . Let  $f^*$  and  $f_\eta^*$  be the global minimizers of  $R(f)$  and  $R^\eta(f)$  respectively. Risk minimization under a given loss function is noise robust if  $f_\eta^*$  has the same probability of misclassification as that of  $f^*$  on noise free data. When the above is satisfied we also say that the loss function is noise-tolerant.

**Theorem 1.** In a multi-class classification problem,  $\ell_{rce}$  is noise tolerant under symmetric or uniform label noise if noise rate  $\eta < 1 - \frac{1}{K}$ . And, if  $R(f^*) = 0$ ,  $\ell_{rce}$  is also noise tolerant under asymmetric or class-dependent label noise when noise rate  $\eta_{yk} < 1 - \eta_y$  with  $\sum_{k \neq y} \eta_{yk} = \eta_y$ .

*Proof.* For symmetric noise:

$$\begin{aligned} R^\eta(f) &= \mathbb{E}_{\mathbf{x}, \hat{y}} \ell_{rce} = \mathbb{E}_{\mathbf{x}} \mathbb{E}_{y|\mathbf{x}} \mathbb{E}_{\hat{y}|\mathbf{x}, y} \ell_{rce} \\ &= \mathbb{E}_{\mathbf{x}} \mathbb{E}_{y|\mathbf{x}} \left[ (1 - \eta) \ell_{rce} + \frac{\eta}{K-1} \sum_{k \neq y}^K \ell_{rce} \right] \\ &= (1 - \eta) R(f) + \frac{\eta}{K-1} \left( \sum_{k=1}^K \ell_{rce} - R(f) \right) \\ &= R(f) \left( 1 - \frac{\eta K}{K-1} \right) - A\eta, \end{aligned}$$

where the last equality holds due to  $\sum_{k=1}^K \ell_{rce} = -(K-1)A$  following Eq. (5) and definition of  $\log 0 = A$ . Thus,

$$R^\eta(f^*) - R^\eta(f) = \left( 1 - \frac{\eta K}{K-1} \right) (R(f^*) - R(f)) \leq 0$$

because  $\eta < 1 - \frac{1}{K}$  and  $f^*$  is a global minimizer of  $R(f)$ . This proves  $f^*$  is also the global minimizer of risk  $R^\eta(f)$ , that is,  $\ell_{rce}$  is noise tolerate.

Similarly, we can prove the case for asymmetric noise, please refer Appendix A for details. ■

**Gradient analysis:** We next derive the gradients of a simplified SL with  $\alpha, \beta = 1$  to give a rough idea of how its learning process differs from that of CE<sup>†</sup>. For brevity, we denote  $p_k, q_k$  as abbreviations for  $p(k|\mathbf{x})$  and  $q(k|\mathbf{x})$ . Consider the case of a single true label, the gradient of the sample-wise RCE loss with respect to the logits  $z_j$  can be derived as:

$$\frac{\partial \ell_{rce}}{\partial z_j} = - \sum_{k=1}^K \frac{\partial p_k}{\partial z_j} \log q_k, \quad (8)$$

where  $\frac{\partial p_k}{\partial z_j}$  can be further derived based on whether  $k = j$ :

$$\frac{\partial p_k}{\partial z_j} = \begin{cases} p_k(1 - p_k), & k = j \\ -p_j p_k, & k \neq j. \end{cases} \quad (9)$$

According to Eq. (9) and the ground-truth distribution for the case of one single label (e.g.,  $q_y = 1$ , and  $q_k = 0$  for  $k \neq y$ ), the gradients of SL can be derived as:

$$\frac{\partial \ell_{sl}}{\partial z_j} = \begin{cases} \frac{\partial \ell_{ce}}{\partial z_j} - (Ap_j^2 - Ap_j), & q_j = q_y = 1 \\ \frac{\partial \ell_{ce}}{\partial z_j} + (-Ap_j p_y), & q_j = 0, \end{cases} \quad (10)$$

<sup>†</sup>Complete derivations can be found in the Appendix B.

where  $A$  is the smoothed/clipped replacement of  $\log 0$ . Note that the gradient of sample-wise CE loss  $\ell_{ce}$  is:

$$\frac{\partial \ell_{ce}}{\partial z_j} = \begin{cases} p_j - 1 \leq 0, & q_j = q_y = 1 \\ p_j \geq 0, & q_j = 0. \end{cases} \quad (11)$$

In the case of  $q_j = q_y = 1$  ( $\frac{\partial \ell_{ce}}{\partial z_j} \leq 0$ ), the second term  $Ap_j^2 - Ap_j$  is an adaptive acceleration term based on  $p_j$ . Specifically,  $Ap_j^2 - Ap_j$  is a convex parabolic function in the first quadrant for  $p_j \in [0, 1]$ , and has the maximum value at  $p_j = 0.5$ . Considering the learning progresses towards  $p_j \rightarrow 1$ , RCE increases DNN prediction on label  $y$  with larger acceleration for  $p_j < 0.5$  and smaller acceleration for  $p_j > 0.5$ . In the case of  $q_j = 0$  ( $\frac{\partial \ell_{ce}}{\partial z_j} \geq 0$ ), the second term  $-Ap_j p_y$  is an adaptive acceleration on the minimization of the probability at unlabeled class ( $p_j$ ), based on the confidence at the labeled class ( $p_y$ ). Larger  $p_y$  leads to larger acceleration, that is, if the network is more confident about its prediction at the labeled class, then the residual probabilities at other unlabeled classes should be reduced faster. When  $p_y = 0$ , there is no acceleration, which means if the network is not confident on the labeled class at all, then the label is probably wrong, no acceleration needed.

### 4.3. Discussion

An easy addition to improve CE would be to upscale its gradients with a larger coefficient (e.g., ‘2CE’, ‘5CE’). However, this will cause more overfitting (see the ‘5CE’ curve in the following Section 5 Figure 9a). There are also other options to consider, such as MAE. Although motivated from completely different perspectives, that is, CE and RCE are measures of (information theoretic) uncertainty, while MAE is a measure of distance, we can surprisingly show that MAE is a special case of RCE at  $A = -2$ , when there is a single true label for  $\mathbf{x}$  (e.g.  $q(y|\mathbf{x}) = 1$  and  $q(k \neq y|\mathbf{x}) = 0$ ). For MAE, we have,

$$\begin{aligned} \ell_{mae} &= \sum_{k=1}^K |p(k|\mathbf{x}) - q(k|\mathbf{x})| = (1 - p(y|\mathbf{x})) + \sum_{k \neq y} p(k|\mathbf{x}) \\ &= 2(1 - p(y|\mathbf{x})), \end{aligned}$$

while, for RCE, we have,

$$\begin{aligned} \ell_{rce} &= - \sum_{k=1}^K p(k|\mathbf{x}) \log q(k|\mathbf{x}) \\ &= -p(y|\mathbf{x}) \log 1 - \sum_{k \neq y} p(k|\mathbf{x}) A = -A \sum_{k \neq y} p(k|\mathbf{x}) \\ &= -A(1 - p(y|\mathbf{x})). \end{aligned}$$

That is, when  $A = -2$ , RCE is reduced to exactly MAE. Meanwhile, different from the GCE loss (i.e., a weighted MAE) [29], SL is a combination of two symmetrical learning terms.

## 5. Experiments

We first provide some empirical understanding of our proposed SL approach, then evaluate its robustness against noisy labels on MNIST, CIFAR-10, CIFAR-100, and a large-scale real-world noisy dataset Clothing1M.

**Noise setting:** We test two types of label noise: symmetric (uniform) noise and asymmetric (class-dependent) noise. Symmetric noisy labels are generated by flipping the labels of a given proportion of training samples to one of the other class labels uniformly. Whilst for asymmetric noisy labels, flipping labels only occurs within a specific set of classes [16, 29], for example, for MNIST, flipping  $2 \rightarrow 7, 3 \rightarrow 8, 5 \leftrightarrow 6$  and  $7 \rightarrow 1$ ; for CIFAR-10, flipping TRUCK  $\rightarrow$  AUTOMOBILE, BIRD  $\rightarrow$  AIRPLANE, DEER  $\rightarrow$  HORSE, CAT  $\leftrightarrow$  DOG; for CIFAR-100, the 100 classes are grouped into 20 super-classes with each has 5 sub-classes, then flipping between two randomly selected sub-classes within each super-class.

### 5.1. Empirical Understanding of SL

We conduct experiments on CIFAR-10 dataset with symmetric noise towards a deeper understanding of SL.

**Experimental setup:** We use an 8-layer CNN with 6 convolutional layers followed by 2 fully connected layers. All networks are trained using SGD with momentum 0.9, weight decay  $10^{-4}$  and an initial learning rate of 0.01 which is divided by 10 after 40 and 80 epochs (120 epochs in total). The parameter  $\alpha, \beta$  and  $A$  in SL are set to 0.1, 1 and -6 respectively.

**Class-wise learning:** The class-wise test accuracy of SL on 40% noisy labels has already been presented in Figure 1d. Here we provide further results for 60% noisy labels in Figure 5. Under both settings, each class is more sufficiently learned by SL than CE, accompanied by accuracy increases. Particularly for the hard classes (e.g., classes 2/3/4/5), SL significantly improves their learning performance. This is because SL facilitates an adaptive pace to encourage learning from hard classes. During learning, samples from easy classes can be quickly learned to have a high probability  $p_k > 0.5$ , while samples from hard classes still have a low probability  $p_k < 0.5$ . SL will balance this discrepancy by increasing the learning speed for samples with  $p_k < 0.5$  while decreasing the learning speed for those with  $p_k > 0.5$ .

**Prediction confidence and distribution:** In comparison to the low confidence of CE on the clean samples in Figure 4a, we train the same network using SL under the same setting. As shown in Figure 6a, on the **clean portion** of class 3 samples, SL successfully pulls up the confidence to 0.95, while at the same time, pushes down the residual confidence at other classes to almost 0. As further shown in Figure 6b, the prediction distributions demonstrate that each class contains more than 4000 true positive samples, including the hard classes (e.g., class 2/3/4/5). Some classes (e.g., class 1/6/7/8/9) even obtain close to 5000 true positive samples (the ideal case). Compared to the earlier results in Figure

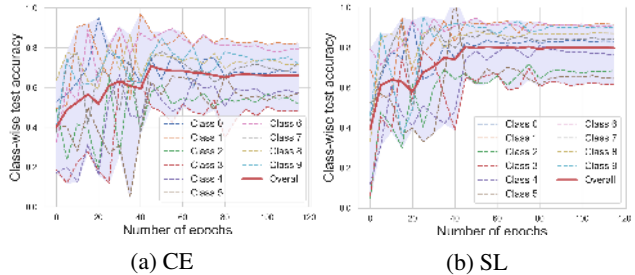


Figure 5: Class-wise test accuracy of CE and SL on CIFAR-10 dataset with 60% symmetric noisy labels. The red solid lines are the overall test accuracies.

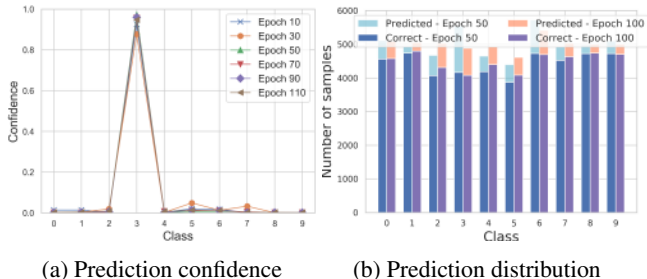


Figure 6: Effect of the proposed SL on prediction confidence/distribution on CIFAR-10 with 40% noisy labels. (a) Average confidence of the **clean portion** of class 3 samples. (b) The true positive samples (*correct*) out of predictions (*predicted*) for each class.

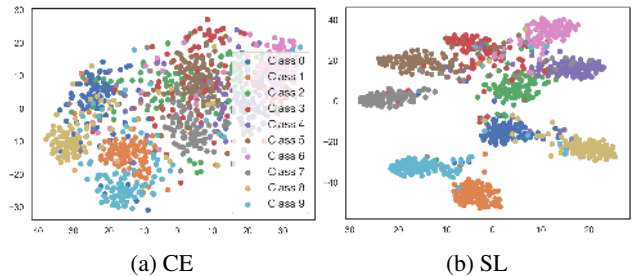


Figure 7: Representations learned by CE and SL on CIFAR-10 dataset with 60% symmetric noisy labels.

4b, SL achieves considerable improvement on each class. **Representations:** We further investigate the representations learned by SL compared to that learned by CE. We extract the high-dimensional representation at the second last dense layer, then project to a 2D embedding using t-SNE [14]. The projected representations are illustrated in Figures 2 and 7 for 40% and 60% noisy labels respectively. Under both settings, the representations learned by SL are of significantly better quality than that of CE with more separated and clearly bounded clusters.

**Parameter analysis:** We tune the parameters of SL:  $\alpha, \beta$  and  $A$ . As  $\beta$  can be reflected by  $A$ , here we only show results of  $\alpha$  and  $A$ . We tested  $A$  in  $[-8, -2]$  with step 2 and  $\alpha \in [10^{-2}, 1]$  on CIFAR-10 under 60% noisy labels. Figure 8a shows that large  $\alpha$  (e.g., 1.0/0.5) tends to cause more overfitting, while small  $\alpha$  (e.g., 0.1/0.01) can help ease

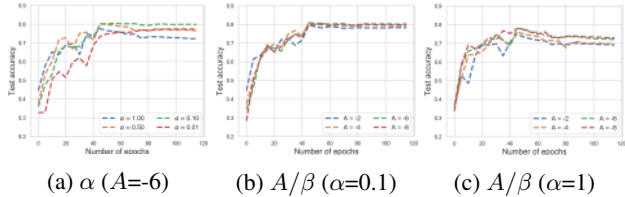


Figure 8: Parameter analysis for SL with an 8-layer CNN on CIFAR-10 dataset under 60% symmetric label noise: (a) Tuning  $\alpha$  (fix  $A = -6$ ); (b) Tuning  $A/\beta$  (fix  $\alpha = 0.1$ ); and (c) Tuning  $A/\beta$  (fix  $\alpha = 1$ ).

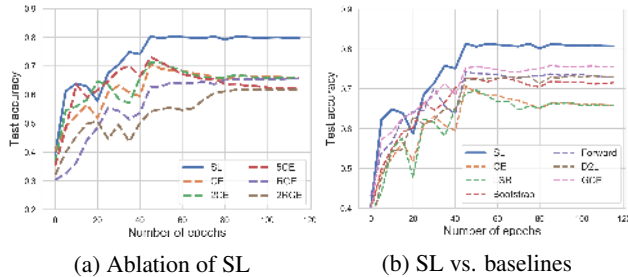


Figure 9: Accuracy of different models on CIFAR-10 with 60% symmetric label noise. (a) Ablation study of SL; (b) Comparison between SL and other baselines.

the overfitting of CE. Nevertheless, the convergence can become slow when  $\alpha$  is too small (e.g., 0.01), a behaviour like the single RCE. For this reason, a relatively large  $\alpha$  can help convergence on difficult datasets such as CIFAR-100. As for parameter  $A$ , if the overfitting of CE is well controlled by  $\alpha = 0.1$ , SL is not sensitive to  $A$  (Figure 8b). However, if CE overfitting is not properly addressed, SL becomes mildly sensitive to  $A$  (Figure 8c).

**Ablation study:** For a comprehensive understanding of each term in SL, we further conduct a series of ablation experiments on CIFAR-10 under 60% noisy labels. Figure 9a presents the following experiments: 1) removing the RCE term; 2) removing the CE term; 3) upscaling the CE term; and 4) upscaling the RCE term. We can observe that simply upscaling CE does not help learning, or even leads to more overfitting. The RCE term itself does not exhibit overfitting even when upscaled, but it converges slowly. But when CE and RCE are combined into the SL framework, the performance is drastically improved.

## 5.2. Robustness to Noisy Labels

**Baselines:** We compare SL with 5 recently proposed noisy label learning methods as well as the standard CE loss: (1) Forward [16]: Training with label correction by multiplying the network prediction with the ground truth noise matrix; (2) Bootstrap [18]: Training with new labels generated by a convex combination of the raw labels and the predicted labels; (3) GCE [29]: Training with a noise robust loss encompassing both MAE and CE; (4) D2L [13]: Training with subspace dimensionality adapted labels; (5) Label Smoothing Regularization (LSR) [17]: Training with CE on soft

labels, rather than the one-hot labels; and (6) CE: Training with standard cross entropy loss.

**Experimental setup:** Experiments are conducted on MNIST [10], CIFAR-10 [9] and CIFAR-100 [9]. We use a 4-layer CNN for MNIST, the same network as Section 5.1 for CIFAR-10 and a ResNet-44 [7] for CIFAR-100. Parameters for the baselines are configured according to their original papers. For our SL, we set  $A = -4$  for all datasets, and  $\alpha = 0.01, \beta = 1.0$  for MNIST,  $\alpha = 0.1, \beta = 1.0$  for CIFAR-10,  $\alpha = 6.0, \beta = 0.1$  for CIFAR-100 (a dataset known for hard convergence)<sup>‡</sup>. All networks are trained using SGD with momentum 0.9, weight decay  $5 \times 10^{-3}$  and an initial learning rate of 0.1. The learning rate is divided by 10 after 10 and 30 epochs for MNIST (50 epochs in total), after 40 and 80 epochs for CIFAR-10 (120 epochs in total), and after 80 and 120 epochs for CIFAR-100 (150 epochs in total). Simple data augmentation techniques (width/height shift and horizontal flip) are applied on CIFAR-10 and CIFAR-100. For symmetric noise, we test varying noise rates  $\eta \in [0\%, 80\%]$ , while for asymmetric noise, we test noise rates  $\eta \in [0\%, 40\%]$ .

**Robustness performance:** The classification accuracies are reported in Table 1. As can be seen, SL improves on the baselines via a large margin for almost all noise rates and all datasets. Note that Forward sometimes also delivers a relatively good performance, as we directly provide it with the ground truth noise matrix. We also find that SL can be more effective than GCE, particularly for high noise rates. The complete learning procedures of SL and baselines on CIFAR-10 are illustrated in Figure 9b. SL shows a clear advantage over other methods, especially in the later stages of learning with noisy labels. This is likely because that, in the later stages of DNN learning, other methods all suffer to some extent from under learning on hard classes, while SL ensures sufficient learning on them.

**Enhancing existing methods with SL:** We introduce some general principles to incorporate SL into existing methods to further enhance their performance. For methods that use robust loss functions or label corrections, the RCE term of SL can be directly added to the loss function, while for methods that still use the standard CE loss without label corrections, SL can be used with small  $\alpha$  and large  $\beta$  to replace the existing loss function. This is to avoid overfitting while promote sufficient learning. As a proof-of-concept, we conduct experiments to enhance Forward and LSR with SL. For “Forward+SL”, we add the RCE term to the Forward loss with  $\beta = 1.0/0.1$  for symmetric/asymmetric noise respectively, while for “LSR+SL”, we use the SL loss with the same setting in Table 1. Results on CIFAR-10 are presented in Table 2. Both the enhanced methods demonstrate a clear performance improvement over their original versions (Forward or LSR) both on symmetric and asymmetric noise. However, in some scenarios, the enhanced methods

<sup>‡</sup>For 40% asymmetric noise,  $\beta$  is set to 5.0 for CIFAR-10 and  $\alpha$  is set to 2.0 for CIFAR-100. Other parameters are unchanged.

Table 1: Test accuracy (%) of different models on benchmark datasets with various rates of symmetric and asymmetric noisy labels. The average accuracy and standard deviation of 5 random runs are reported and the best results are in **bold**.

Datasets	Methods	Symmetric Noise					Asymmetric Noise		
		Noise Rate $\eta$					Noise Rate $\eta$		
		0.0	0.2	0.4	0.6	0.8	0.2	0.3	0.4
MNIST	CE	99.02 ± 0.01	88.71 ± 0.05	69.56 ± 0.19	46.54 ± 0.28	21.77 ± 0.07	93.14 ± 0.04	87.91 ± 0.05	81.10 ± 0.07
	LSR	99.28 ± 0.01	89.56 ± 0.06	68.11 ± 0.24	45.01 ± 0.15	21.28 ± 0.27	94.18 ± 0.08	88.39 ± 0.20	81.09 ± 0.35
	Bootstrap	99.08 ± 0.01	88.72 ± 0.14	69.97 ± 0.36	47.06 ± 0.26	22.60 ± 0.27	93.31 ± 0.03	87.87 ± 0.09	80.46 ± 0.15
	Forward	99.03 ± 0.01	94.85 ± 0.07	86.02 ± 0.13	69.77 ± 0.41	49.72 ± 0.30	97.31 ± 0.05	96.25 ± 0.10	95.72 ± 0.09
	D2L	99.27 ± 0.01	98.80 ± 0.01	98.49 ± 0.01	93.61 ± 0.01	48.57 ± 0.04	98.71 ± 0.02	97.77 ± 0.04	93.32 ± 0.15
	GCE	99.04 ± 0.01	98.66 ± 0.01	97.17 ± 0.01	79.65 ± 0.14	31.55 ± 0.18	96.73 ± 0.08	88.46 ± 0.18	81.26 ± 0.11
	<b>SL</b>	<b>99.32 ± 0.01</b>	<b>99.02 ± 0.01</b>	<b>98.97 ± 0.01</b>	<b>97.40 ± 0.02</b>	<b>65.02 ± 0.19</b>	<b>99.18 ± 0.01</b>	<b>98.85 ± 0.01</b>	<b>98.80 ± 0.02</b>
CIFAR-10	CE	89.26 ± 0.03	82.96 ± 0.05	78.70 ± 0.07	66.62 ± 0.15	34.80 ± 0.25	85.98 ± 0.03	83.53 ± 0.08	78.51 ± 0.05
	LSR	88.57 ± 0.04	83.49 ± 0.05	78.41 ± 0.03	67.38 ± 0.15	36.30 ± 0.16	85.38 ± 0.05	82.89 ± 0.12	77.88 ± 0.20
	Bootstrap	88.77 ± 0.06	83.95 ± 0.10	79.97 ± 0.07	71.65 ± 0.05	41.44 ± 0.49	86.57 ± 0.08	84.86 ± 0.05	79.76 ± 0.07
	Forward	89.39 ± 0.04	85.83 ± 0.05	81.37 ± 0.03	73.59 ± 0.08	47.10 ± 0.14	87.68 ± 0.01	<b>86.86 ± 0.06</b>	<b>85.73 ± 0.04</b>
	D2L	86.66 ± 0.05	81.13 ± 0.06	76.80 ± 0.12	60.67 ± 0.12	19.83 ± 0.05	82.72 ± 0.06	80.41 ± 0.05	73.33 ± 0.12
	GCE	86.76 ± 0.03	84.86 ± 0.06	82.42 ± 0.10	75.20 ± 0.09	40.81 ± 0.24	84.61 ± 0.09	82.11 ± 0.13	75.32 ± 0.10
	<b>SL</b>	<b>89.28 ± 0.04</b>	<b>87.63 ± 0.06</b>	<b>85.34 ± 0.07</b>	<b>80.07 ± 0.02</b>	<b>53.81 ± 0.27</b>	<b>88.24 ± 0.05</b>	85.36 ± 0.14	80.64 ± 0.10
CIFAR-100	CE	64.34 ± 0.37	59.26 ± 0.39	50.82 ± 0.19	25.39 ± 0.09	5.27 ± 0.06	62.97 ± 0.19	63.12 ± 0.16	61.85 ± 0.35
	LSR	63.68 ± 0.54	58.83 ± 0.40	50.05 ± 0.31	24.68 ± 0.43	5.22 ± 0.07	63.03 ± 0.48	62.32 ± 0.48	61.59 ± 0.41
	Bootstrap	63.26 ± 0.39	57.91 ± 0.42	48.17 ± 0.18	12.27 ± 0.11	1.00 ± 0.01	63.44 ± 0.35	63.18 ± 0.35	62.08 ± 0.22
	Forward	63.99 ± 0.52	59.75 ± 0.34	53.13 ± 0.28	24.70 ± 0.26	2.65 ± 0.03	64.09 ± 0.61	64.00 ± 0.32	60.91 ± 0.36
	D2L	64.60 ± 0.31	59.20 ± 0.43	52.01 ± 0.37	35.27 ± 0.28	5.33 ± 0.54	62.43 ± 0.28	63.20 ± 0.27	61.35 ± 0.66
	GCE	64.43 ± 0.20	59.06 ± 0.27	53.25 ± 0.65	36.16 ± 0.74	8.43 ± 0.80	63.03 ± 0.22	63.17 ± 0.26	61.69 ± 1.15
	<b>SL</b>	<b>66.75 ± 0.04</b>	<b>60.01 ± 0.19</b>	<b>53.69 ± 0.07</b>	<b>41.47 ± 0.04</b>	<b>15.00 ± 0.04</b>	<b>65.58 ± 0.06</b>	<b>65.14 ± 0.05</b>	<b>63.10 ± 0.13</b>

Table 2: Accuracy (%) of SL-boosted Forward and LSR methods on CIFAR-10 under various label noise.

Method	Symmetric noise		Asymmetric noise
	0.4	0.6	0.4
Forward+SL	84.54 ± 0.03	79.64 ± 0.04	86.22 ± 0.18
LSR+SL	85.20 ± 0.01	79.28 ± 0.05	80.99 ± 0.30

are still not as good as SL. This often occurs when there is a large performance gap between the original methods and SL. We believe that with more adaptive incorporation and careful parameter tuning, SL can be combined with existing approaches to achieve even better performance.

### 5.3. Experiments on Real-world Noisy Dataset

In the above experiments, we have seen that SL achieves excellent performance on datasets with manually corrupted noisy labels. Next, we assess its applicability for a real-world large-scale noisy dataset: Clothing1M [26].

The Clothing1M dataset contains 1 million images of clothing obtained from online shopping websites with 14 classes: T-shirt, Shirt, Knitwear, Chiffon, Sweater, Hoodie, Windbreaker, Jacket, Down Coat, Suit, Shawl, Dress, Vest, and Underwear. The labels are generated by the surrounding text of images and are thus extremely noisy. The overall accuracy of the labels is  $\sim 61.54\%$ , with some pairs of classes frequently confused with each other (e.g., Knitwear and Sweater), which may contain both symmetric and asymmetric label noise. The dataset also provides 50k, 14k, 10k manually refined clean data for training, validation and testing respectively, but we did not use the 50k clean data. The classification accuracy on the 10k clean testing data is used as the evaluation metric.

**Experimental setup:** We use ResNet-50 with ImageNet pretrained weights similar to [16, 26]. For preprocessing, images are resized to  $256 \times 256$ , with mean value subtracted

Table 3: Accuracy (%) of different models on real-world noisy dataset Clothing1M. The best results are in **bold**.

Methods	CE	Bootstrap	Forward	D2L	GCE	<b>SL</b>
Acc	68.80	68.94	69.84	69.47	69.75	<b>71.02</b>

and cropped at the center of  $224 \times 224$ . We train the models with batch size 64 and initial learning rate  $10^{-3}$ , which is reduced by 1/10 after 5 epochs (10 epochs in total). SGD with a momentum 0.9 and weight decay  $10^{-3}$  are adopted as the optimizer. Other settings are the same as Section 5.2. **Results:** As shown in Table 3, SL obtains the highest performance compared to the baselines. We also find that Forward achieves a relatively good result, though it requires the use of the part of data that both have noisy and clean labels to obtain the noise transition matrix, which is not often available in real-world settings. SL only requires the noisy data and does not require extra auxiliary information.

## 6. Conclusions

In this paper, we identified a deficiency of cross entropy (CE) used in DNN learning for noisy labels, in relation to under learning of hard classes. To address this issue, we proposed the Symmetric cross entropy Learning (SL), boosting CE symmetrically with the noise robust Reverse Cross Entropy (RCE), to simultaneously addresses its under learning and overfitting problems. We provided both theoretical and empirical understanding on SL, and demonstrated its effectiveness against various types and rates of label noise on both benchmark and real-world datasets. Overall, due to its simplicity and ease of implementation, we believe SL is a promising loss function for training robust DNNs against noisy labels, and an attractive framework to be used along with other techniques for datasets containing noisy labels.



## References

- [1] Devansh Arpit, Stanisaw Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, and Simon Lacoste-Julien. A closer look at memorization in deep networks. In *ICML*, 2017. 1, 4
- [2] Jia Deng, Wei Dong, Richard Socher, Jia Li, Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 1
- [3] Aritra Ghosh, Himanshu Kumar, and PS Sastry. Robust loss functions under label noise for deep neural networks. In *AAAI*, 2017. 2, 4
- [4] Jacob Goldberger and Ehud Ben-Reuven. Training deep neural-networks using a noise adaptation layer. In *ICLR*, 2017. 2
- [5] Bo Han, Jiangchao Yao, Gang Niu, Mingyuan Zhou, Ivor Tsang, Ya Zhang, and Masashi Sugiyama. Masking: A new perspective of noisy supervision. In *NeurIPS*, 2018. 2
- [6] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: robust training deep neural networks with extremely noisy labels. In *NeurIPS*, 2018. 3
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 7
- [8] 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 *ICML*, 2018. 3
- [9] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Technical Report, University of Toronto*, 2009. 3, 7
- [10] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 7
- [11] Kuang-Huei Lee, Xiaodong He, Lei Zhang, and Linjun Yang. Cleannet: Transfer learning for scalable image classifier training with label noise. *arXiv preprint arXiv:1711.07131*, 2017. 2
- [12] Yuncheng Li, Jianchao Yang, Yale Song, Liangliang Cao, Jiebo Luo, and Jia Li. Learning from noisy labels with distillation. In *ICCV*, 2017. 2
- [13] Xingjun Ma, Yisen Wang, Michael E Houle, Shuo Zhou, Sarah M Erfani, Shu-Tao Xia, Sudanthi Wijewickrema, and James Bailey. Dimensionality-driven learning with noisy labels. In *ICML*, 2018. 1, 3, 7
- [14] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008. 6
- [15] Eran Malach and Shai Shalev-Shwartz. Decoupling” when to update” from” how to update”. In *NeurIPS*, 2017. 3
- [16] Giorgio Patrini, Alessandro Rozza, Aditya Menon, Richard Nock, and Lizhen Qu. Making neural networks robust to label noise: a loss correction approach. In *CVPR*, 2017. 2, 6, 7, 8
- [17] Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548*, 2017. 2, 7
- [18] 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. 2, 7
- [19] Ravid Shwartz-Ziv and Naftali Tishby. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*, 2017. 1
- [20] Sainbayar Sukhbaatar, Joan Bruna, Manohar Paluri, Lubomir Bourdev, and Rob Fergus. Training convolutional networks with noisy labels. *arXiv preprint arXiv:1406.2080*, 2014. 2
- [21] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016. 2
- [22] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *CVPR*, 2018. 3
- [23] Arash Vahdat. Toward robustness against label noise in training deep discriminative neural networks. In *NeurIPS*, 2017. 2
- [24] Andreas Veit, Neil Alldrin, Gal Chechik, Ivan Krasin, Abhinav Gupta, and Serge Belongie. Learning from noisy large-scale datasets with minimal supervision. In *CVPR*, 2017. 2
- [25] Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, and Shu-Tao Xia. Iterative learning with open-set noisy labels. In *CVPR*, 2018. 3
- [26] Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *CVPR*, 2015. 2, 8
- [27] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *ICML*, 2019. 3
- [28] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR*, 2017. 1, 4
- [29] Zhilu Zhang and Mert R Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. In *NeurIPS*, 2018. 2, 4, 5, 6, 7