

A Novel Self-Supervised Re-labeling Approach for Training with Noisy Labels

Devraj Mandal *
Indian Institute of Science
devrajm@iisc.ac.in

Shrisha Bharadwaj *†
University of Tübingen
shrishabharadwaj5@gmail.com

Soma Biswas
Indian Institute of Science
somabiswas@iisc.ac.in

Abstract

The major driving force behind the immense success of deep learning models is the availability of large datasets along with their clean labels. This is very difficult to obtain and thus has motivated research on training deep neural networks in the presence of label noise. In this work, we build upon the seminal work in this area, Co-teaching and propose a simple, yet efficient approach termed mCT-S2R (modified co-teaching with self-supervision and re-labeling) for this task. Firstly, to deal with significant amount of noise in the labels, we propose to use self-supervision to generate robust features without using any labels. Furthermore, using a parallel network architecture, an estimate of the clean labeled portion of the data is obtained. Finally, using this data, a portion of the estimated noisy labeled portion is re-labeled, before resuming the network training with the augmented data. Extensive experiments on three standard datasets show the effectiveness of the proposed framework.

1. Introduction

The success of deep learning models like Alexnet [16], VGG [34], ResNet [12], etc. for image classification tasks can be attributed to the availability of large, annotated datasets like ImageNet [17]. But, obtaining clean annotations of large datasets is very expensive, and thus recent research has focused on learning from weakly supervised data, where the labels are often *noisy*, since they have been acquired from web searches [19] or crowd-sourcing [37]. The presence of noisy labels can severely degrade the performance of the learned classifiers, since deep neural networks can even over-fit on the noisy labels with sufficient training, due to their memorization capability [40] [3].

Recently, several approaches have been proposed to handle label noise in the training data [40] [11] [38] [13] [21] [10] [30]. One direction to address this problem is to estimate the noise transition matrix itself by utilizing an ad-

ditional softmax layer [10] for modeling the channel noise. An alternative two-step approach along the same lines is proposed in [30]. However, it is observed that estimation of the noise transition matrix is often hard and computationally expensive, especially when large number of classes are present in the data [11].

A more recent and elegant direction to handle label noise is to utilize the concept of small-loss instances [13] [32] [11]. Here, the algorithms estimate which samples are correctly labeled, and uses them to train the network subsequently. MentorNet [13] uses a pre-trained network (trained in a self-paced manner with a curriculum loss) to select clean labeled data to train the final classification model. Co-teaching [11] trains two networks in parallel and updates the weights of the networks using only the small loss samples. In addition, the gradients of the two networks are switched during the parameter update, which leads to better performance. It is observed [11] that when the training continues for a long time, the two networks generally converge to the same state and start performing similarly leading to the accumulation of errors.

Here, we propose a framework based on the Co-teaching approach [11], which also uses the concept of small-loss instances along with self-supervision and re-labeling, to significantly improve the training of deep networks with very noisy training data. The proposed approach has four main steps. First, to deal with significant amount of label noise, we utilize self-supervision as a pre-training step, so that the network can learn robust features without the use of any labels. Second, we use a parallel network architecture similar to the one in [11] to get an estimate of the small loss samples. In the third step, utilizing a portion of the small loss samples, the per class means for all the categories are computed, which are then used to re-label the large loss samples. Finally, the training of the network is resumed by taking all the small loss samples along with a portion of the re-labeled large loss samples based on a confidence measure. The proposed framework is termed as mCT-S2R (modified co-teaching with self-supervision and re-labeling). The main contributions of this work are as follows:

- We develop an approach by utilizing self-supervision

*Both authors contributed equally.

†work done during internship at Indian Institute of Science, Bangalore.

and re-labeling of the large loss data using the small loss samples for training deep networks under significant label noise.

- Our framework uses two parallel networks only to determine the small loss instances. Unlike [11] [21] which requires two networks for the entire training, the final training of the proposed framework after re-labeling proceeds using a single network. This makes our model computationally much lighter.
- We propose using a self-supervised training paradigm like [9] to learn about the data distribution without using the labels, which helps to avoid over-fitting on the noisy labels.
- Extensive experiments on three benchmark datasets show the effectiveness of the proposed framework.

2. Related Work

Training deep neural networks in the presence of label noise is an active research area [21][11][38][8]. Several directions have been proposed in the literature to address this problem. Some works [31] [7] have focused on designing an appropriate loss function, which acts as a regularization term while training. [31] proposes to increase the weight of the regularization term during the training paradigm so that the network focuses more on the clean examples with subsequent iterations. An unbiased estimator to handle label noise is proposed in [25], while a robust non-convex loss function is designed in [22] for this task.

Another research direction is to model the noise behavior itself [10][24]. The work in [24] models the conditional probability of observing a wrong label, and considers the correct label to be a latent variable. However, [24] assumes that the label-flip probability is known apriori, which requires extra knowledge about the training data. This has been mitigated in [10], where an additional softmax layer is used to model the noise transition matrix. Noise rate estimation is also addressed in the works of [23] [18] and [33].

Most of the current approaches in this area utilize the concept of *small loss* instances [13][21][11], where the network tries to estimate those samples which are probably correctly labeled, and ignore those samples which are probably noisy. Due to the ease of training these models and their generalization capability, they have been well accepted in the machine learning community. Our approach falls in this category, and we propose a framework mCT-S2R for this task.

There are some other related works [26][4][2][43], which assume the availability of a small clean set of data in addition to the large, noisy or unlabeled dataset for training the classification model. One recent work in [35] approaches this problem from the semi-supervised paradigm

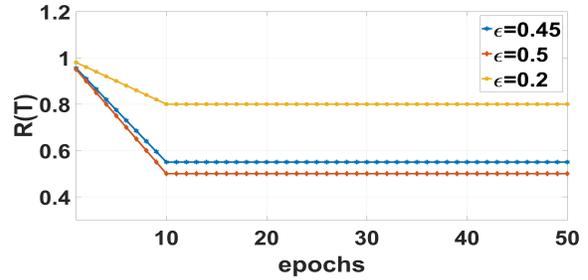


Figure 1. Plot of $R(T)$ for different values of label corruption $\epsilon = \{0.45, 0.5, 0.2\}$. $R(T)$ is used [11] to control how many samples are used to update the network weights per epoch. The proposed approach augments the training set by adding a portion of the large loss samples after re-labeling them using the small loss instances.

and uses the consistency loss along with the rotation self-supervised loss in an end-to-end manner to design noise robust classification models. The work [35] also needs a clean validation set (which is not available in our case) to tune the hyper-parameters of the model to get the best possible results. In addition, since the classification loss, consistency loss and the self-supervision objectives are used in an end-to-end manner [35], proper weighing of each of the factors for each individual dataset is required. These approaches are beyond the scope of our work, since we do not assume the presence of any examples with clean labels, which makes the tuning of the hyper-parameters of our model very difficult [11][38]. Thus, we try to design models which are very robust to the choice of the hyper-parameters.

3. Motivation

To develop the proposed framework, first, we analyze the recently proposed methods which utilize the small loss instances [11]. These approaches typically use two network architecture, and each network updates the peer network weights based on the small loss samples. The training procedure uses a factor $R(T)$ (at epoch T) to quantify the amount of small loss samples per mini-batch to be used to update the network weights. $R(T)$ is directly related to the amount of label corruption ϵ (ϵ is either assumed to be known or can be estimated by using the strategies outlined in [18] [39]), and is defined as $R(T) = 1 - \min\{\frac{T}{T_k}\epsilon, \epsilon\}$, where T_k is the epoch after which $R(T)$ becomes constant. Typical profiles of $R(T)$ for different ϵ is shown in Figure 1. We observe that the model initially considers all the input data and slowly zones into the cleanly labeled data (guided by the concept of small loss instances) to train the model. After completion of T_k epochs (here $T_k = 10$), the network weights are updated using $(1 - \epsilon)$ portion of the training data (marked as small loss instances) and ignores the large loss samples under the assumption that they are wrongly labeled. Even though the small and large loss samples for

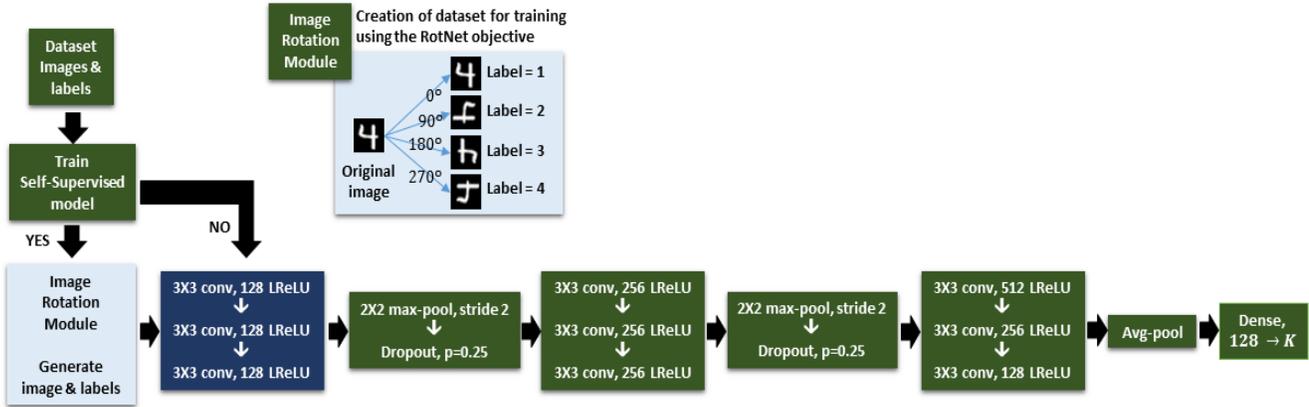


Figure 2. The architecture of the network used for the pre-training task using the self-supervised objective. For the pretext task, each image is fed into the image rotation module (marked with light blue shade) to generate different rotated images and their corresponding labels (i.e., the degrees of rotation), and the network is trained using the standard cross-entropy loss. For training the final noise-tolerant classification model, two networks of the same architectural design as shown are taken. The weights are randomly initialized and only the first 3 convolutional layer weights (marked with deep blue shade) are loaded with the weights from the pre-trained self-supervised model. The two parallel networks are trained using the images and their provided noisy labels. (Figure best viewed in color)

every mini-batch of an epoch are sampled randomly, the basic premise is to only focus on those samples whose labels are possibly correct. This automatically leads to ignoring a large portion of the data samples for training the model. In this work, we analyze whether the small loss samples can be used to correct the labels of some of the incorrectly labeled samples, which can then be utilized to better train the network.

4. Proposed Approach

The proposed framework has four main steps: (1) pre-training the neural network in a self-supervised fashion to learn about the data distribution in an unsupervised manner, (2) training of two parallel networks following the same strategy as in [21] [11] to estimate the small loss instances, (3) re-labeling of the large loss samples using the small loss instances and (4) final training of one network using an augmented set of the small loss samples and a portion of the re-labeled large loss samples. Pseudo-code of the proposed approach is provided in Algorithm 1. Now, we describe the different steps in detail.

4.1. Pre-training

To deal with significant amount of label noise, we first want to impart the knowledge of the training data distribution to the proposed network architecture without using the noisy labels. Recently, self-supervised models like [9][41][42][6][27][1] have been used to learn high level convolutional network features in an unsupervised manner. Features learned in a self-supervised way by solving a pretext task such as image colorization [41][42], prediction of relative position of image patches [6][27], ego-motion pre-

dition by using motion between two frames [1] and prediction of 2d rotation of the images [9] have helped to learn unsupervised features which are useful for further downstream computer vision tasks.

Inspired by this, we propose to first use a self-supervised training paradigm to learn a feature representation without using any label information. We have selected the RotNet model [9] for our framework as shown in Figure 2 (marked with light blue shade). We have selected RotNet model [9] for the self-supervision task because studies in [14] have shown that this model gives superior performance in the downstream classification task as compared to the other self-supervision approaches. Here, the pretext task is to correctly predict the rotation operation (such as image rotations of 0, 90, 180, 270 degrees) which have been applied to the input image. The loss function used is the standard cross-entropy loss \mathcal{L}^{ce} with a 4-dim logit layer (for the 4 degrees of orientation i.e., here $K = 4$).

As mentioned, the proposed mCT-S2R is built on the Co-teaching approach [11] and has the same base architecture (Figure 2). For seamless transfer of the learned weights from the self-supervised pretext task, we use the same architecture as in [11] with RotNet objective [9]. To train the model for the pretext task, we feed each image through the image rotation module (shown inset in Figure 2) to generate the training data and its labels, i.e. the orientation degree. This training is performed on the same dataset on which we need to train the classification model later i.e., the downstream task. Once trained, this pre-trained network can be used for building the noise-robust classification model.

Based on the studies in [9][11][21], we need to consider two important things: (1) the best way to transfer the knowl-

edge learned while solving the pretext task and (2) imparting randomness in the two parallel networks, so that they do not converge to the same model. The network architecture for the two parallel networks is identical as shown in Figure 2. Each network has a final dense layer of K -dim, which is randomly initialized, where, K denotes the total number of training categories in the data. The work in [9] has shown that the initial layers learn generic features as compared to the later layers, which are mostly tuned to the pretext task. Thus to impart the knowledge of the data distribution following the studies in [9] [28], the weights of the first 3 convolutional layers of both the networks in the proposed framework are initialized with the corresponding weights of the pre-trained self-supervised model. The randomness of the rest of the layers prevents the two networks from converging to the same model during training.

4.2. Computing the small-loss instances

We train two parallel networks, each one having the architecture as in Figure 2 simultaneously as in [21] [11] by feeding the original image data and the corresponding noisy labels. Let us denote the two networks as p and q and their learnable weights as w_p and w_q . For each network, we denote the output from the final, fully connected (fc) layer logits to be $\{u^p, u^q\}$ of dimension K , and the penultimate fc layer features to be $\{f^p, f^q\}$ (128-dim). The algorithm proceeds as in [11], where, in each mini-batch, the network computes the loss over the available samples in the mini-batch, sorts the loss in ascending order, and selects a portion of the samples to form the small loss instance set $\{\mathcal{D}^p, \mathcal{D}^q\}$ based on the value of $R(T)$. The parameters w_p and w_q are updated by using the information from the peer network i.e., \mathcal{D}^q and \mathcal{D}^p . The loss to train the network is the standard cross-entropy loss $\mathcal{L}^{ce} = -\log\left(\frac{e^{u[i]}}{\sum_{k=1}^K e^{u[k]}}\right)$, where i is the correct class index.

In [11], the two networks are trained until the pre-defined number of iterations T_{\max} . During testing, one of the learned networks (either p or q) is used for classification. In contrast, we train the two parallel networks for an initial T_{update} number of steps and then collect all the small loss $\{S^p, S^q\}$ and large loss $\{L^p, L^q\}$ indices set based on the value of $R(T)$. Unlike [11], we use the two networks only to get an estimate of the small and large loss samples. For subsequent training, we choose one of the networks, thus reducing the computational requirement significantly.

4.3. Re-labeling the large loss samples

Let us consider, that we choose the network p and its index set $\{S^p, L^p\}$. The set of all small loss samples is denoted as \mathcal{D}^s . In this step, first, all the training data is passed through the network p and their features $\mathcal{F} = \{f^p\}$ are extracted. The indices S^p and L^p are used to compute the feature sets for the small loss and large loss samples.

Algorithm 1 Proposed Framework - mCT-S2R

Input: Network structure $\{p, q\}$, learning rate η , ϵ , epoch $T_k, T_{update}, T_{\max}$, total number of mini-batches N_{\max} , number (N) of small loss instances to compute class means, training set \mathcal{D} , threshold value κ .

Output: Trained noise robust classification model p .

Step 1: Pre-training

1. Randomly initialize one network and train it using the self-supervised objective loss function as in [9].
2. Randomly initialize the model weights w_p and w_q for the two networks p and q .
3. The initial three conv layer weights of both p and q are loaded with the weights of the pre-trained self-supervised model, while the remaining layers are randomly initialized.

Step 2: Compute small-loss instances

```

for  $T = 1, 2, \dots, T_{update}$  do
  Shuffle the training set  $\mathcal{D}$ .
  for  $N = 1, \dots, N_{\max}$  do
    Fetch mini-batch  $\mathcal{D}^{mb}$  from  $\mathcal{D}$ .
    Sub-select a portion of the samples (as defined
    by  $R(T)$ ) from  $\mathcal{D}^{mb}$  to form the small loss set  $\mathcal{D}^j$ .
    Hence,  $\mathcal{D}^j = \arg \min_{x: |x| \geq R(T)|\mathcal{D}^{mb}|} \mathcal{L}^{ce}(j, x)$ 
    Update  $w_j = w_j - \eta \nabla \mathcal{L}^{ce}(j, \mathcal{D}^j)$ .
    ( $j = \{p, q\}$ , if  $j = p, \bar{j} = q$ )
    Update  $R(T) = 1 - \min\left(\frac{T}{T_k} \epsilon, \epsilon\right)$ .

```

Step 3: Re-labeling large-loss samples

1. Pick one network randomly (say p). Store the sorted (in ascending order) small and large loss indices at T_{update} epoch as $\{S^p, L^p\}$. The small loss samples are denoted as \mathcal{D}^s .
2. Compute the class-wise means using the the top- N small loss instances (**Algorithm2**).
3. Based on distance of the large loss samples from the means, generate the re-labeled targets \mathcal{Y}^r and their confidence of re-labeling \mathcal{C} (**Algorithm3**).
5. Select the large loss samples whose confidence $c \geq \kappa$ to form \mathcal{D}^r .
6. Construct augmented training set $\mathcal{D}^{aug} = \mathcal{D}^s \cup \mathcal{D}^r$.

Step 4: Final training with augmented data

```

for  $T = T_{update} + 1, T_{update} + 2, \dots, T_{\max}$  do
  Shuffle the training set  $\mathcal{D}^{aug}$ .
  for  $N = 1, \dots, N_{\max}$  do
    Fetch mini-batch  $\mathcal{D}^{mb}$  from  $\mathcal{D}^{aug}$ .
    Obtain  $\mathcal{D}^p = \arg \min_{x: |x| \geq R(T)|\mathcal{D}^{mb}|} \mathcal{L}^{ce}(p, x)$ 
    Update  $w_p = w_p - \eta \nabla \mathcal{L}^{ce}(p, \mathcal{D}^p)$ .
    Update  $R(T) = 1 - \min\left(\frac{T}{T_k} \epsilon, \epsilon\right)$ .

```

Algorithm 2 Feature Extraction & Mean Set Computation

Input: Network p , training data \mathcal{D} , sorted small & large loss index set $\{S^p, L^p\}$, N : number of samples for computing class means.

Output: Features of large loss samples \mathcal{F}^L , class-wise means $\mathcal{M} = \{\mu_1, \dots, \mu_K\}$ for K categories

1. Pass data \mathcal{D} through network p and generate the extracted features.
 2. Use index set S^p and L^p to get the features for large loss and small loss samples as \mathcal{F}^L and \mathcal{F}^S .
 3. Use N small loss samples per category to compute its mean μ_k for class $k \in K$.
-

Algorithm 3 Re-labeling & Confidence Computation

Input: Feature set of large loss samples \mathcal{F}^L , class-wise means $\mathcal{M} = \{\mu_1, \dots, \mu_K\}$

Output: Re-labeled targets $\mathcal{Y}^r = \{y^r\}$, Confidence values $\mathcal{C} = \{c\}$

for $i = 1, \dots, |\mathcal{F}^L|$ **do**

 Compute distance d of the i^{th} sample from \mathcal{M} .

 Generate similarity value \hat{d} by passing $(-d)$ through softmax function. (A smaller distance denotes a higher similarity.)

 Set pseudo label as: $y^r = \arg \max_k \hat{d}[k]$

 Set confidence of re-labeling as: $c = \max_k \hat{d}[k]$

The per class mean feature representation is then computed using the small loss sample features to form the mean set $\mathcal{M} = \{\mu_1, \mu_2, \dots, \mu_K\}$. Instead of using all the samples of each class to compute the means, we use only the top- N small loss samples per category, i.e. the samples with the lowest loss. This step of feature extraction and class-wise mean computation is illustrated in Algorithm 2.

Next, we use the class-wise means \mathcal{M} and the extracted features of the large loss samples to generate their pseudo-labels \mathcal{Y}^r and confidence values \mathcal{C} . The pseudo labels are generated by computing the distance of each large loss sample feature from \mathcal{M} and converting it into a softmax probability score. The index of the highest softmax value gives the pseudo-label y^r , while the highest value is stored as its confidence measure value c . The steps are outlined in Algorithm 3.

4.4. Final training with augmented data

Once we obtain the re-labeled samples and their confidence scores, we sub-select a portion of the large loss samples, whose confidence $c \geq \kappa$. Let us denote this set by \mathcal{D}^r . Finally, we construct the augmented training set as $\mathcal{D}^{aug} = \mathcal{D}^s \cup \mathcal{D}^r$. We use the original provided labels for the small loss instances i.e., \mathcal{D}^s and the predicted labels for the large loss samples i.e., \mathcal{D}^r to further train the chosen network (here p).

5. Experimental Evaluation

Here, we describe in details the extensive experiments performed to evaluate the effectiveness of the proposed framework.

5.1. Dataset details and noise models

We conduct experiments on three standard benchmark datasets used for evaluating models trained with noisy labels, namely MNIST [5], CIFAR-10 [15] and CIFAR-100 [15]. MNIST [5] dataset contains gray-scale images of handwritten digits of size 28×28 spread over 10 different categories. This dataset has 60,000 training and 10,000 testing examples. CIFAR-10 [15] and CIFAR-100 [15] contains a huge collection of 32×32 size color images which has been divided into a training:testing split of 50,000 : 10,000 examples. CIFAR-100 (CIFAR-10) data contains 100 (10) different categories, with each category having 600 (6000) images per class.

We follow the standard protocol as in [30][31] to introduce label noise into the datasets. Specifically, the labels are corrupted using a noise transition matrix P , where $P_{ij} = \text{Prob}(\hat{l} = j | l = i)$, i.e., probability that the label of an example actually belonging to category i ($l = i$) has been corrupted to belong to category j ($\hat{l} = j$). We have used two different noise models: (1) symmetry flipping [36] (symmetry) and (2) pair flipping (pair). In symmetry flipping, the probability of an example being correctly labeled under the noisy setting is $(1 - \epsilon)$, while the probability of being corrupted (i.e. ϵ) is equally distributed over the rest of the categories. In pair flipping, the probability of being corrupted is distributed over a single category. Studies in [11] have shown that it is typically harder to train classification networks under the pairflip noise model as compared to the symmetry noise model, which is further validated by the performance of the algorithms under those settings.

We evaluate our framework under extreme noisy setting [11], with $\epsilon = 0.45$ (for pairflip) and $\epsilon = 0.5$ (for symmetry), where almost half of the training examples have been corrupted and have noisy labels. Studies in [11] have shown that algorithms need additional information to handle more label noise i.e., $\epsilon > 0.5$ in the training data.

5.2. Network structure and implementation details

We implement the proposed framework in PyTorch (version 0.4.1) [29] on a NVIDIA GTX 1080 Ti GPU using the architecture shown in Figure 2. The model is trained with Adam optimizer using momentum of 0.9 and a learning rate

Table 1. Average test accuracy on the MNIST, CIFAR-10 and CIFAR-100 datasets over the last ten epochs in the supervised setting (with 100% clean labels).

	MNIST	CIFAR-10	CIFAR-100
Acc	99.60	89.34	63.41

Table 2. Average test accuracy on the MNIST dataset over the last ten epochs. The results of our approach without (and with) self-supervision pre-training is denoted as mCT-R and mCT-S2R respectively.

Flipping-Rate	Standard	Bootstrap	S-model	F-correction	Decoupling	MentorNet	Co-teaching	mCT-R	mCT-S2R
Pair-45%	56.52 ± 0.55	57.23 ± 0.73	56.88 ± 0.32	0.24 ± 0.03	58.03 ± 0.07	80.88 ± 4.45	87.63 ± 0.21	95.52 ± 0.04	94.72 ± 0.02
Symmetry-50%	66.05 ± 0.61	67.55 ± 0.53	62.29 ± 0.46	79.61 ± 1.96	81.15 ± 0.03	90.05 ± 0.30	91.32 ± 0.06	97.79 ± 0.02	97.82 ± 0.02
Symmetry-20%	94.05 ± 0.16	94.40 ± 0.26	98.31 ± 0.11	98.80 ± 0.12	95.70 ± 0.02	96.70 ± 0.22	97.25 ± 0.03	99.05 ± 0.03	98.86 ± 0.01

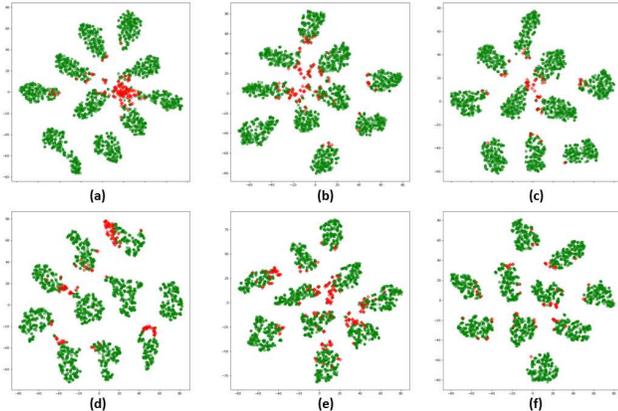


Figure 3. T-sne plots of the large loss samples for MNIST. The small loss samples are used to re-label the large loss instances and the correctly (wrongly) re-labeled data are marked with green (red) respectively. (a, b, c) and (d, e, f) denote the results of mCT-R and mCT-S2R. (Left to right): The three columns correspond to the noise models: pairflip-45%, symmetric-50% and symmetric-20%. (Figure best viewed in color)

of 0.001. First, the model is trained using self-supervision based on the RotNet objective for 25 epochs. For handling large amounts of label noise, the first three convolutional layers of the two parallel networks are initialized with pre-trained weights of the self-supervised model and the remaining layers are initialized with random weights. We use a batch size of 128 and train for 200 epochs, and report the average test accuracy over the last ten epochs. We follow the same protocol as in [11] and fix the other parameters of our model as $R(T) = 1 - \epsilon \cdot \min(\frac{T}{T_k}, 1)$, $T_k = 10$, $T_{max} = 200$. We set $T_{update} = 30$ and the softmax confidence measure $\kappa = 0.90$ for all our experiments. We also perform ablation studies to analyze the performance of our algorithm with the change in T_{update} and κ and report the results later.

We compare the proposed approach against the state-of-art approaches: Bootstrap [31], S-model [10], F-correction [30], Decoupling [21], MentorNet [13] and Co-teaching [11]. The results of all the compared methods are taken from the published results in [11]. We also compute the classification accuracy of the proposed mCT-S2R on the three datasets using all the clean labels. The results in Table 1 represents the upper-bound of classification accuracy that can possibly be achieved by our network.

5.3. Evaluation on MNIST dataset

The results of the proposed mCT-S2R under different noise models on the MNIST dataset is reported in Table 2. We also report the results without self-supervision (denoted by mCT-R) to analyze the importance of the two major components of the proposed framework, namely self-supervision and re-labeling. We observe the following: (1) For all the approaches, the performance for the pairflip noise model is lower as compared to symmetry flip; (2) As expected, the performance of all the algorithms for 20% symmetric noise is better than for 50% noise; (3) The proposed approach significantly outperforms the state-of-the-art models, especially for pairflip noise; (4) The performance of mCT-S2R for the symmetric case is close to the best possible value (shown in Table 1). This might be because MNIST is an easier dataset compared to the others and the data for a particular category in MNIST has comparatively lesser variations. (5) The results of mCT-S2R and mCT-R are very similar, signifying that self-supervision is not particularly beneficial for this dataset.

Now, we analyze the usefulness of re-labeling in the proposed framework. We train the two parallel networks till T_{update} , then randomly select one of them, and use the small loss instances to predict the labels of the large loss instances. The t-sne [20] plots of the large loss sample features are shown in Figure 3. The large loss samples which are correctly re-labeled are marked with *green*, while those which are wrongly re-labeled are marked with *red*. The top (bottom) row denotes the results without (with) self-supervised pre-training. We observe that most of the large loss samples are getting correctly re-labeled in both the cases. Based on a softmax confidence threshold, our framework selects a portion of the re-labeled samples (preferably from the green marked ones) to train the network further. This extra set of correctly re-labeled samples helps our approach to achieve better classification accuracy.

5.4. Evaluation on CIFAR-10 dataset

We report the results of our approach on CIFAR-10 dataset in Table 3. Even for this data, we observe similar trends in the results as with MNIST. But for CIFAR-10, the results of mCT-S2R are significantly better than mCT-R, which clearly shows the effectiveness of self-supervised pre-training for more difficult data.

Table 3. Average test accuracy on the CIFAR-10 dataset over the last ten epochs. The results of our approach without (and with) self-supervision pre-training is denoted as mCT-R and mCT-S2R respectively.

Flipping-Rate	Standard	Bootstrap	S-model	F-correction	Decoupling	MentorNet	Co-teaching	mCT-R	mCT-S2R
Pair-45%	49.50 ± 0.42	50.05 ± 0.30	48.21 ± 0.55	6.61 ± 1.12	48.80 ± 0.04	58.14 ± 0.38	72.62 ± 0.15	78.09 ± 0.11	80.58 ± 0.54
Symmetry-50%	48.87 ± 0.52	50.66 ± 0.56	46.15 ± 0.76	59.83 ± 0.17	51.49 ± 0.08	71.10 ± 0.48	74.02 ± 0.04	77.69 ± 0.24	81.23 ± 0.07
Symmetry-20%	76.25 ± 0.28	77.01 ± 0.29	76.84 ± 0.66	84.55 ± 0.16	80.44 ± 0.05	80.76 ± 0.36	82.32 ± 0.07	84.89 ± 0.09	87.21 ± 0.04

Table 4. Average test accuracy on the CIFAR-100 dataset over the last ten epochs. The results of our approach without (and with) self-supervision pre-training is denoted as mCT-R and mCT-S2R respectively.

Flipping-Rate	Standard	Bootstrap	S-model	F-correction	Decoupling	MentorNet	Co-teaching	mCT-R	mCT-S2R
Pair-45%	31.99 ± 0.64	32.07 ± 0.30	21.79 ± 0.86	1.60 ± 0.04	26.05 ± 0.03	31.60 ± 0.51	34.81 ± 0.07	31.80 ± 0.11	38.67 ± 0.08
Symmetry-50%	25.21 ± 0.64	21.98 ± 6.36	18.93 ± 0.39	41.04 ± 0.07	25.80 ± 0.04	39.00 ± 1.00	41.37 ± 0.08	41.46 ± 0.06	52.87 ± 0.16
Symmetry-20%	47.55 ± 0.47	47.00 ± 0.54	41.51 ± 0.60	61.87 ± 0.21	44.52 ± 0.04	52.13 ± 0.40	54.23 ± 0.08	56.29 ± 0.19	60.49 ± 0.36

The t-sne plots for the different losses in Figure 4 (top row) clearly indicate that for the models trained without self-supervision, distinct clusters of data have not been formed and the number of wrongly marked samples (marked in "red") is quite high. Thus, the re-labeled samples will potentially have many wrongly assigned labels. In the bottom row, self-supervision pre-training clearly helps in forming distinct (10) clusters of data leading to a larger number of correctly re-labeled samples (marked in "green"). These additional correctly labeled data helps the proposed framework to achieve significantly better classification performance.

5.5. Evaluation on CIFAR-100 dataset

We report the results of our algorithm on the CIFAR-100 dataset in Table 4. This is the hardest dataset of all the three due to the presence of 100 different categories, which increases the possibility of incorrect re-labeling of the large loss samples. Indeed, we observe that for pairflip-45%, mCT-R shows poorer performance than [11], possibly due to incorrect re-labeling of the large loss samples. For this dataset, mCT-S2R shows significant improvements as compared to mCT-R. From the t-sne plots for the different losses in Figure 5, we clearly observe how self-supervised pre-training immensely helps to correctly re-label the samples leading to the improved performance of mCT-S2R.

5.6. Ablation studies

Here, we provide additional analysis of the proposed framework. The analysis of why self-supervised pre-training helps to improve the performance of the network is shown in Figure 3, 4, 5. Further analysis for different choice of the hyper-parameters is provided below.

Choice of T_{update} : The T_{update} parameter controls the epoch after which we re-label the large loss samples using

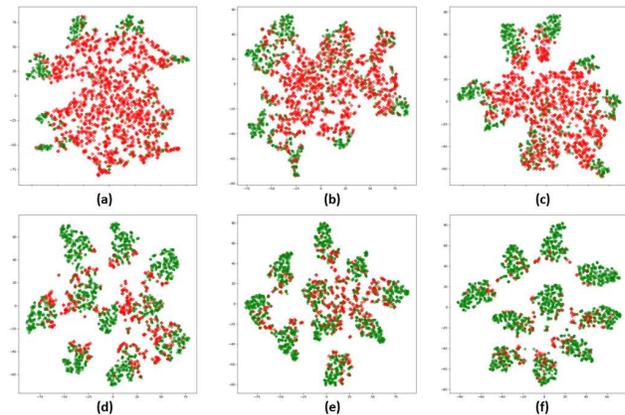


Figure 4. T-sne plots of the large loss samples for CIFAR-10. The small loss samples are used to re-label the large loss instances and the correctly (wrongly) re-labeled data are marked with green (red) respectively. (a, b, c) and (d, e, f) denote the results of mCT-R and mCT-S2R. (Left to right): The three columns correspond to the noise models: pairflip-45%, symmetric-50% and symmetric-20%. (Figure best viewed in color)

the small loss instances. We analyze the performance of the proposed mCT-S2R using different values of $T_{update} = \{10, 20, 30\}$ and report the results in Table 5 on MNIST and CIFAR-100 datasets. We observe that the results are quite consistent, which implies that the proposed framework performs well for a wide range of hyper-parameter T_{update} value. Having a smaller value of T_{update} is advantageous for our algorithm since after re-labeling, we can train only a single network which is computationally much lighter.

Choice of $R(T)$: The value of the dropping rate $R(T)$ is set to be $1 - \epsilon \cdot \min(\frac{T}{T_k}, 1)$ and common profiles of $R(T)$ for different noise levels ϵ is shown in Figure 1. Essentially, after $T \geq T_k$, $R(T) = 1 - \epsilon$, which implies that the training tries to concentrate on the correctly labeled portion of the

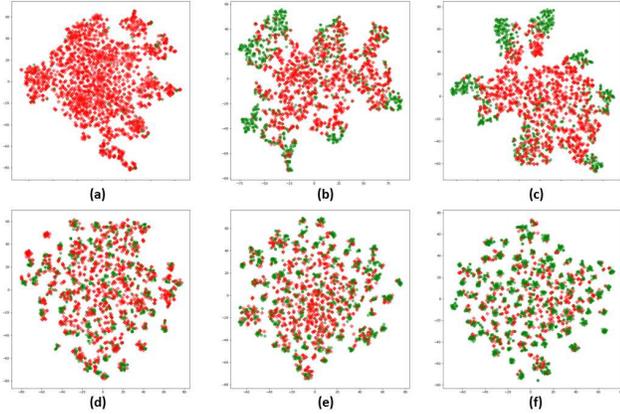


Figure 5. T-sne plots of the large loss samples for CIFAR-100. The small loss samples are used to re-label the large loss instances and the correctly (wrongly) re-labeled data are marked with green (red) respectively. (a, b, c), (d, e, f) denote the results of mCT-R and mCT-S2R. (Left to right): The three columns correspond to noise models: pairflip-45%, symmetric-50% and symmetric-20%. (Figure best viewed in color)

Table 5. Average test accuracy (over the last ten epochs) on MNIST and CIFAR-100 datasets for different values of T_{update} .

MNIST	Symmetry-50%			Pair-45%		
T_{update}	10	20	30	10	20	30
Accuracy	98.62	98.38	97.82	97.41	96.66	94.72
CIFAR-100	Symmetry-50%			Pair-45%		
T_{update}	10	20	30	10	20	30
Accuracy	53.84	53.61	52.87	38.57	39.88	38.67

data. It was observed in [11] that if too many instances are dropped, then the networks might not get sufficient training data and the performance can substantially deteriorate. However, in our work, since we are re-labeling the large loss samples and utilizing the possibly correct ones for training, the performance of the proposed approach may not be so intricately related to the value of ϵ . We analyze the performance of our algorithm in Table 6 by setting $\epsilon = 0.5$ to compute $R(T)$. We choose this value, since, in general, it is difficult to train deep based noise-tolerant models for a larger noise rate ($\epsilon > 0.5$) without any additional information [11]. From Table 6, we observe that the performance of mCT-S2R is quite consistent, which indicates that even without the knowledge of the true noise level ϵ , our framework should perform well.

Choice of the softmax threshold value κ : The parameter κ controls which large loss re-labeled samples are picked for further training of the network. We analyze the performance of our algorithm in Table 7 by setting $\kappa = \{0.80, 0.90, 0.95\}$ and observe that the results are consistent for a wide range of κ values. We considered a high value of κ for our evaluation as a higher value essentially means that we are only using the most confident re-labeled

Table 6. Average test accuracy (over the last ten epochs) on MNIST, CIFAR-10 and CIFAR-100 datasets for $\epsilon = 0.5$ as compared to the recommended value ($\epsilon = \text{actual noise level}$, denoted by “*”) as proposed in [11].

ϵ	MNIST		CIFAR-10		CIFAR-100	
	*	0.5	*	0.5	*	0.5
Pair-45%	94.72 ± 0.02	97.79 ± 0.02	80.58 ± 0.54	83.53 ± 0.10	38.67 ± 0.08	41.06 ± 0.13
Symmetry-20%	98.86 ± 0.01	99.45 ± 0.02	87.21 ± 0.04	83.65 ± 0.07	60.49 ± 0.36	60.76 ± 2.14

Table 7. Average test accuracy (over the last ten epochs) on the MNIST and CIFAR-100 datasets for different values of κ .

MNIST	Symmetry-50%			Pair-45%		
κ	0.80	0.90	0.95	0.80	0.90	0.95
Accuracy	94.34	94.72	94.95	97.91	97.82	97.79
CIFAR-100	Symmetry-50%			Pair-45%		
κ	0.80	0.90	0.95	0.80	0.90	0.95
Accuracy	38.14	38.67	37.95	53.19	52.87	53.04

samples for training the network.

Choice of the network for re-labeling: Here, we analyze the performance of our method when one of the networks (in random) is chosen for re-labeling and subsequent training. We observe from the results in Table 8 that the performance does not vary much and hence our framework is robust to the choice of the network used for re-labeling.

Table 8. Average test accuracy (over the last ten epochs) on the MNIST, CIFAR-10 and CIFAR-100 datasets for different noise levels: (a) Pair-45%, (b) Symmetry-50% and (c) Symmetry-20% with the 1^{st} (first row) and 2^{nd} (second row) network used for re-labeling respectively.

	MNIST			CIFAR-10			CIFAR-100		
	(a)	(b)	(c)	(a)	(b)	(c)	(a)	(b)	(c)
1^{st}	94.72	97.82	98.86	80.58	81.23	87.21	38.67	52.87	60.49
2^{nd}	94.56	97.53	98.88	80.76	81.10	86.99	39.57	53.00	61.23

6. Conclusion

In this work, we present a framework mCT-S2R for training neural networks under different amounts of label noise. The training paradigm uses the concept of small loss instances along with re-labeling to build noise-resistant classification models. In addition, we have also shown how self-supervision pre-training can effectively help to boost the performance further. Our framework works on a wide variety of datasets under different noise models and significantly outperforms the current state-of-the-art, while being computationally lighter. Ablation studies have shown that the proposed model performs very well over a wide range of hyper-parameter values and thus does not need an exclusive clean validation set for tuning of the hyper-parameters for optimum performance.

References

- [1] P. Agrawal, J. Carreira, and J. Malik. Learning to see by moving. In *ICCV*, pages 37–45, 2015.
- [2] R. K. Ando and T. Zhang. Two-view feature generation model for semi-supervised learning. In *ICML*, pages 25–32, 2007.
- [3] D. Arpit, S. Jastrzebski, N. Ballas, D. Krueger, E. Bengio, M. S. Kanwal, T. Maharaj, A. Fischer, A. Courville, Y. Bengio, et al. A closer look at memorization in deep networks. In *ICML*, pages 233–242, 2017.
- [4] A. Blum and T. Mitchell. Combining labeled and unlabeled data with co-training. In *COLT*, pages 92–100, 1998.
- [5] L. Deng. The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- [6] C. Doersch, A. Gupta, and A. A. Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, pages 1422–1430, 2015.
- [7] D. Flatow and D. Penner. On the robustness of convnets to training on noisy labels, 2017.
- [8] B. Frénay and M. Verleysen. Classification in the presence of label noise: a survey. *IEEE Transactions on Neural Networks and Learning Systems*, 25(5):845–869, 2013.
- [9] S. Gidaris, P. Singh, and N. Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, pages 1–16, 2018.
- [10] J. Goldberger and E. Ben-Reuven. Training deep neural networks using a noise adaptation layer. In *ICLR*, pages 1–9, 2017.
- [11] B. Han, Q. Yao, X. Yu, G. Niu, M. Xu, W. Hu, I. Tsang, and M. Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NIPS*, pages 8527–8537, 2018.
- [12] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [13] L. Jiang, Z. Zhou, T. Leung, L.-J. Li, and L. Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, pages 2304–2313, 2018.
- [14] A. Kolesnikov, X. Zhai, and L. Beyer. Revisiting self-supervised visual representation learning. In *CVPR*, pages 1920–1929, 2019.
- [15] A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- [17] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- [18] T. Liu and D. Tao. Classification with noisy labels by importance reweighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(3):447–461, 2015.
- [19] W. Liu, Y.-G. Jiang, J. Luo, and S.-F. Chang. Noise resistant graph ranking for improved web image search. In *CVPR*, pages 849–856, 2011.
- [20] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.
- [21] E. Malach and S. Shalev-Shwartz. Decoupling” when to update” from” how to update”. In *NIPS*, pages 960–970, 2017.
- [22] H. Masnadi-Shirazi and N. Vasconcelos. On the design of loss functions for classification: theory, robustness to outliers, and savageboost. In *NIPS*, pages 1049–1056, 2009.
- [23] A. Menon, B. Van Rooyen, C. S. Ong, and B. Williamson. Learning from corrupted binary labels via class-probability estimation. In *ICML*, pages 125–134, 2015.
- [24] V. Mnih and G. E. Hinton. Learning to label aerial images from noisy data. In *ICML*, pages 567–574, 2012.
- [25] N. Natarajan, I. S. Dhillon, P. K. Ravikumar, and A. Tewari. Learning with noisy labels. In *NIPS*, pages 1196–1204, 2013.
- [26] K. Nigam and R. Ghani. Analyzing the effectiveness and applicability of co-training. In *CIKM*, pages 86–93, 2000.
- [27] M. Noroozi and P. Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, pages 69–84, 2016.
- [28] M. Noroozi, A. Vinjimoor, P. Favaro, and H. Pirsiavash. Boosting self-supervised learning via knowledge transfer. In *CVPR*, pages 9359–9367, 2018.
- [29] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. 2017.
- [30] G. Patrini, A. Rozza, A. Krishna Menon, R. Nock, and L. Qu. Making deep neural networks robust to label noise: A loss correction approach. In *CVPR*, pages 1944–1952, 2017.
- [31] S. Reed, H. Lee, D. Anguelov, C. Szegedy, D. Erhan, and A. Rabinovich. Training deep neural networks on noisy labels with bootstrapping. In *ICLR Workshop*, pages 1–11, 2015.
- [32] M. Ren, W. Zeng, B. Yang, and R. Urtasun. Learning to reweight examples for robust deep learning. In *ICML*, pages 4334–4343, 2018.
- [33] T. Sanderson and C. Scott. Class proportion estimation with application to multiclass anomaly rejection. In *AISTATS*, pages 850–858, 2014.
- [34] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [35] T. W. Tsai, C. Li, and J. Zhu. Countering noisy labels by learning from auxiliary clean labels. *arXiv preprint arXiv:1905.13305*, 2019.
- [36] B. Van Rooyen, A. Menon, and R. C. Williamson. Learning with symmetric label noise: The importance of being unhinged. In *NIPS*, pages 10–18, 2015.
- [37] P. Welinder, S. Branson, P. Perona, and S. J. Belongie. The multidimensional wisdom of crowds. In *NIPS*, pages 2424–2432, 2010.
- [38] X. Yu, B. Han, J. Yao, G. Niu, I. Tsang, and M. Sugiyama. How does disagreement help generalization against label corruption? In *ICML*, pages 7164–7173, 2019.

- [39] X. Yu, T. Liu, M. Gong, K. Batmanghelich, and D. Tao. An efficient and provable approach for mixture proportion estimation using linear independence assumption. In *CVPR*, pages 4480–4489, 2018.
- [40] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR*, pages 1–15, 2017.
- [41] R. Zhang, P. Isola, and A. A. Efros. Colorful image colorization. In *ECCV*, pages 649–666, 2016.
- [42] R. Zhang, P. Isola, and A. A. Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *CVPR*, pages 1058–1067, 2017.
- [43] X. J. Zhu. Semi-supervised learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2005.