S3: SUPERVISED SELF-SUPERVISED LEARNING UNDER LABEL NOISE

Anonymous authors
Paper under double-blind review

ABSTRACT

Despite the large progress in supervised learning with Neural Networks, there are significant challenges in obtaining high-quality, large-scale and accurately labeled datasets. In this context, in this paper we address the problem of classification in the presence of label noise and more specifically, both close-set and open-set label noise, that is when the true label of a sample may, or may not belong to the set of the given labels. In the heart of our method is a sample selection mechanism that relies on the consistency between the annotated label of a sample and the distribution of the labels in its neighborhood in the feature space; a relabeling mechanism that relies on the confidence of the classifier across subsequent iterations; and a training strategy that trains the encoder both with a self-consistency loss and the classifier-encoder with the cross-entropy loss on the selected samples alone. Without bells and whistles, such as co-training so as to reduce the self-confirmation bias, and with robustness with respect to settings of its few hyper-parameters, our method significantly surpasses previous methods on both CIFAR10/CIFAR100 with artificial noise and real-world noisy datasets such as WebVision and ANIMAL-10N.

1 INTRODUCTION

It is now commonly accepted that supervised learning with deep neural networks can provide excellent solutions for a wide range of problems, so long as there is sufficient availability of labeled training data and computational resources. However, these results have been mostly obtained using well-curated datasets in which the classes are balanced and the labels are of high quality. In the real-world, it is often costly to obtain high quality labels especially for large-scale datasets. A common approach is to use semi-automatic methods to obtain the labels (e.g. "webly-labeled" images where the images and labels are obtained by web-crawling). While such methods can greatly reduce the time and cost of manual labeling, they also lead to low quality noisy labels.

To deal with noisy labels, earlier approaches tried to improve the robustness of the model using robust loss functions(Ghosh et al., 2017; Zhang & Sabuncu, 2018; Wang et al., 2019) or robust regularizations (Srivastava et al., 2014; Zhang et al., 2017; Pereyra et al., 2017). Goldberger & Ben-Reuven (2016) tried to model the noise transition matrix between classes while Han et al. (2019); Patrini et al. (2017); Hendrycks et al. (2018) proposed to correct the losses of noisy samples. More recently, sample selection methods became perhaps the dominant paradigm for learning with noisy labels. Most of the recent sample selection methods do so, by relying on the predictions of the model classifier, for example on the per-sample loss (Arazo et al., 2019; Li et al., 2020a) or model prediction (Song et al., 2019; Malach & Shalev-Shwartz, 2017). By separating clean samples and noisy samples and subsequently performing supervised training on the clean set, or semi-supervised training on both, sample selection methods achieved the state-of-the-art results in synthetic and real-world noisy datasets.

However, there are three main issues with current sample selection methods. Firstly, the model classifier based sample selection will be inevitably biased if trained with noisy labels – this is immediately apparent in the case that the sample selection is based on the loss of the model itself (Arazo et al., 2019; Li et al., 2020a; Yu et al., 2019; Han et al., 2018). This phenomenon is aggravated by the accumulation of errors induced by the iterative nature of self-train methods where labels are corrected based on the model's predictions – the so-called self-confirmation bias. Secondly, in su-

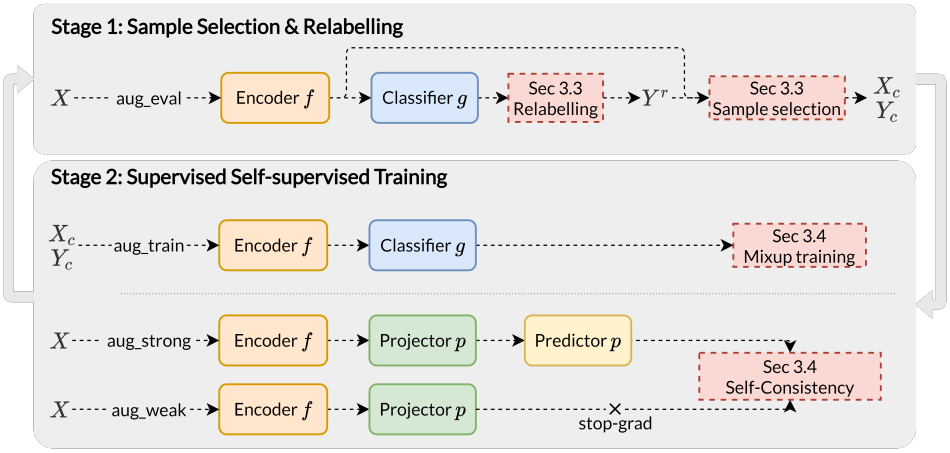

Figure 1: S3 consists of two iterative stages: sample selection&relabelling and supervised self-supervised training.

pervised classification problems, noisy samples usually come from two main categories: closed-set noise where the true labels belong to one of the given classes and open-set noise where the true labels do not belong to the set of labels of the classification problem. Most of the works in the literature, including works that estimate the probabilities of label-exchange between pairs of classes (Goldberger & Ben-Reuven, 2016; Patrini et al., 2017), that do relabeling based on the model's predictions (Song et al., 2019; Han et al., 2019) or works that adopt semi-supervised approaches where samples identified as noisy are considered unlabeled (Li et al., 2020a; Ortego et al., 2021) deal with the former and not directly address the latter. However, this is a considerable source of noise in real world scenarios, e.g., when training from web-crawled data, where there is less control over the collection of the dataset. Finally, current approaches usually require extensive hyperparameters tuning, often even on a per-dataset basi – this is unrealistic in scenarios where there is little knowledge about the types of noise. This is partly due to the complexity of the applied semi-supervised learning method, and partly because of the complicated methods that are employed, such as model pretraining (Zheltonozhskii et al., 2021) and model cotraining (Han et al., 2018; Yu et al., 2019; Li et al., 2020a), so as to deal with self-confirmation bias.

In this paper, we address the problem of training under different types of noise with a simple method–**S3**, that aims at learning a robust feature space, relying on both noisy and clean samples, and identifying in the feature space reliably, clean samples based on the distribution of the labels of the samples in their neighborhood. In the proposed scheme, learning of the feature space is based two factors: First, the consistency loss between the representations of augmented version of the sample (as in (Chen & He, 2021)) – this avoids false-negatives that are inherent in contrastive-learning with instance discrimination; second a classical cross-entropy loss applied only on clean samples – this avoids treating samples identified as noisy as if they belonged to one of the given classes as most methods (Arazo et al., 2019; Li et al., 2020a; Ortego et al., 2021; Wu et al., 2021) do. The noisy sample selection mechanism relies on a measure of confidence what we define using the ground truth label of the sample in question and an estimate of distribution of the labels of its neighbours – in order to deal with noisy samples, we adopt a scheme in which the distribution is calculated based both on the ground truth labels and on consistently (over subsequent iterations) confident estimates of the labels. Our method is embedded into a standard MixUp and data augmentation framework, and without bells and whistles, such as co-training of multiple models, it achieves state-of-the-art results in both synthetic and realistic noise patterns in CIFAR10, CIFAR100, ANIMAL-10N and WebVision datasets.

2 Related Work

Leaning with noisy data by sample selection Some works focused on sample selection to filter out noisy samples. Jiang et al. (2018) introduced a pretrained mentor network to guide the selection

of a student network.Song et al. (2019) evaluates the per-sample losses and identify as clean the top $r\%$ of the samples – the precise ratio $r\%$ depends is either predefined, or is an estimate of the noise level in the specific dataset. Arazo et al. (2019) proposed to model per-sample losses with a Beta Mixture Model (BMM) and split the dataset according to which of the components of the mixture each sample belongs. In a very similar approach, Li et al. (2020a) extended upon Arazo et al. (2019) by introducing semi-supervised learning to fully utilize the dataset. Related to our work, Bahri et al. (2020); Ortego et al. (2021); Wang et al. (2018) also utilized the feature space for sample selection. Bahri et al. (2020) applied KNN for sample selection for closed-set noisy dataset while Ortego et al. (2021) further proposed to relabel samples based on the KNN voting. Wang et al. (2018) proposed to reweight samples based on its probability of being outliers in open-set noisy dataset.

Self-supervised learning Self-supervised methods attempt to learn good representation without human annotations. In the recent years, the dominant method is contrastive learning with instance discrimination task. MoCo (He et al., 2020) is an important baseline for current contrastive learning methods, which reuses the memory bank since samples in a single mini-batch may lead to insufficient negative pairs, and proposes a momentum encoder to update the memory bank in real-time to avoid outdated data representation. SimCLR (Chen et al., 2020) is another important baseline which found that setting mini-batch size to be large enough can eliminate the need for memory bank. More recently, SimSiam (Chen & He, 2021) and BYOL (Grill et al., 2020) proposed a non-contrastive learning framework which enforce the perturbation consitency between different views, and avoid mode collapse by applying stop-gradient and an extra predictor between two representation vectors.

3 METHOD

3.1 PROBLEM FORMULATION

Let us denote with $\boldsymbol{X} = \{\boldsymbol{x}_i\}_{i=1}^N, \boldsymbol{x}_i \in R^d$, a training set with the corresponding one-hot vector labels $\boldsymbol{Y} = \{\boldsymbol{y}_i\}_{i=1}^N, \boldsymbol{y}_i \in \{0,1\}^K$, where K is the number of classes and N is the number of samples. For convenience, let us also denote the index where the one-hot vector $\boldsymbol{y}_i$ is one as the label $l_i \in \{0, ..., K\}$. Finally, let us denote the true labels with $\boldsymbol{Y}' = \{\boldsymbol{y}'_i\}_{i=1}^N$. Clearly, for an open-set noisy label it is the case that $\boldsymbol{y}'_i \neq \boldsymbol{y}_i, \boldsymbol{y}'_i \notin \{0,1\}^K$, while for closed-set noisy samples $\boldsymbol{y}'_i \neq \boldsymbol{y}_i, \boldsymbol{y}'_i \in \{0,1\}^K$. Our goal is to train an encoder f and a classifier g such that $\boldsymbol{f}_i \triangleq f(\boldsymbol{x}_i) \in R^D$ and $\boldsymbol{g}_i \triangleq g(f(\boldsymbol{x}_i)) \in R^K$ are, respectively, a feature representation and a prediction logit of the class of the sample $\boldsymbol{x}$. Typically, the loss for a sample $\boldsymbol{x}_i$ is defined using the cross-entropy loss:

$$L_i = -\sum_{c=1}^{K} \boldsymbol{y}_i^c log \boldsymbol{z}_i^c, \tag{1}$$

where $\boldsymbol{z}_i \triangleq Softmax(\boldsymbol{g}_i) \in R^K$ is the prediction probability vector. We can easily decompose the back-propagating gradient as:

$$\frac{\partial L_i}{\partial \boldsymbol{g}_i} = \boldsymbol{z}_i - \boldsymbol{y}_i \tag{2}$$

From Eq. 2 it becomes apparent that the prediction logit will be optimized to correct class for a correctly labeled sample ($\boldsymbol{y}_i = \boldsymbol{y}'_i$) while for a noisy sample ($\boldsymbol{y}_i \neq \boldsymbol{y}'_i$), the prediction logit will be optimized towards a wrong class. This results to corrupted models when learning under noise.

3.2 OVERVIEW OF PROPOSED METHOD

Aiming to deal with potential concurrence of both open-set and close-set noise, we view the classification network as an encoder f that extracts a feature representation and a classification head g that deals with the classification problem in question. The proposed method, named Supervised Semi-Supervised(S3) learning, attempts to decouple their training so as to deal with possible noise in the labels, by adopting a two stage, iterative scheme, as outlined in Fig. 1.

In the first stage, we utilize a novel sample selection and a novel relabeling mechanism (top block in Fig. 1) that prepares the set based on which the classifier g should be trained in Stage 2. The selection mechanism is based on the assumption of smoothness of labels in the feature space, and more

specifically, on a consistency measure that we defined based on the annotated label of the sample in question and the distribution of the labels in its neighborhood in the feature space. Relabeling is performed on samples for which the classifier gives confident predictions consistently across subsequent iterations. Clearly, the mechanism relies on the quality of the features extracted by the encoder f and should reject samples whose true labels are not in the class set (open-set noise). This stage is explained in Section 3.3.

In the second stage (bottom block in Fig. 1), training is performed with two objectives/losses. First, a cross entropy loss on the output of the classifier g, (i.e., on $g(f(.))$) on the samples selected in Stage 1, that updates both the encoder f and the classifier head g. Second, a self-supervision loss that enforces consistency between the representations of different augmentations of the same sample, and which utilizes all samples, that is both noisy and clean – this updates the encoder f and helps learning a strong feature space on which the selection mechanism of Stage 1 can rely. This is in contrast to other methods, which base the selection mechanism on the classifier itself, and are therefore more susceptible to self-confirmation bias, or need to rely on complex schemes (e.g. co-training) to reduce it. This stage is explained in Section 3.4.

3.3 Sample selection & relabelling

Clean sample selection by normalized neighboring voting Let us denote the similarity between the representations $\boldsymbol{f}_i$ and $\boldsymbol{f}_j$ of any two samples $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ by $\boldsymbol{s}_{ij}, i, j = 1, ..., N$. In our implementation we used the cosine similarity, that is, $\boldsymbol{s}_{ij} \triangleq \frac{\boldsymbol{f}_i^T \boldsymbol{f}_j}{\|\boldsymbol{f}_i\|_2 \|\boldsymbol{f}_j\|_2}$. Let us also denote by N_i the index set of the k nearest neighbors of sample $\boldsymbol{x}_i$ in X based on the calculated similarity. Then, for each sample $\boldsymbol{x}_i$, we calculate the normalized label distribution $\boldsymbol{p}_i \in R^K$ in its neighborhood in the feature space, as the normalized sum of its neighbors' labels. Note, that we used at each epoch t, we use the labels y_n^t (Eq. 6) that a relabeling mechanism provides. More specifically,

$$\boldsymbol{p}_i = \boldsymbol{\pi}^{-1} \boldsymbol{p}'_i, \tag{3}$$

where $\boldsymbol{p}'_i = \frac{1}{k} \sum_{n \in N_i} \boldsymbol{y}_n^t$ and $\boldsymbol{\pi} = \sum_{i=1}^{N} \boldsymbol{p}'_i$. With a slight abuse of notation, we denote by $\boldsymbol{\pi}^{-1}$ the vector whose entries are the inverses of the entries of the vector $\boldsymbol{\pi}$ of the class probabilities in the whole dataset – in this way we compensate for global class imbalances. Once the normalized distribution p_i of the labels in the neighborhood of x_i is estimated, we define a consistency measure c_i as

$$c_i = \frac{\boldsymbol{p}_i(l_i)}{\max_l \boldsymbol{p}_i(l)}, \tag{4}$$

that is the ratio of the value of the distribution p_i at the label l_i divided by the value of its highest peak $\max_l \boldsymbol{p}_i(l)$. Roughly speaking, a high consistency measure c_i at a sample i means that its neighbors agree with the given annotation l_i – this indicates that l_i is likely to be correct. By setting a threshold θ_s to c_i, a clean subset $(\boldsymbol{X}_c, \boldsymbol{Y}_c)$ can be extracted from the noisy dataset.

Noisy sample relabelling by classifier thresholding The normalized distribution p_i of the labels of the neighbors of a sample i constructed by Eq. 3 is to some degree affected by the noisy labels of its neighbors. One option is to use the computed normalized nearest neighbor distribution p_i to relabel the noisy samples – for example, relabel the sample $\boldsymbol{x}_i$ as $\arg\max p_i$ (Ortego et al., 2021). However, this process will introduce self-confirmation bias as we relabel and select samples by relying in both cases on the feature space.

In this paper, we propose to decouple sample selection and sample relabeling. More specifically, given that the classifier is trained in subsequent iterations with relatively clean samples (selected by the mechanism described above), we propose to used its predictions to identify samples for which it has high confidence in its prediction, and relabel them if the confident prediction does not agree with the annotated label. More specifically, let us denote by $\boldsymbol{z}_i^t \triangleq g(f(\boldsymbol{x}_i))$ the prediction at epoch t for sample $\boldsymbol{x}_i$. By keeping track of the most recent L predictions, we calculate

$$\boldsymbol{q}_i = \frac{1}{L} \sum_{t'=t-L}^{t} \boldsymbol{z}_i^{t'}, \tag{5}$$

where $\boldsymbol{q}_i$ is the average prediction for the sample $\boldsymbol{x}_i$. We then modify all the labels $l_i, i = 1, ..., N$ by thresholding $\boldsymbol{q}_i$ at t-th epoch as:

$$l_i^t = I(\max_l \boldsymbol{q}_i(l) > \theta_r) * \arg\max_l \boldsymbol{q}_i(l) + (1 - I(\max_l \boldsymbol{q}_i(l) > \theta_r)) * l_i \tag{6}$$

Please note, that similarly to Section 3.1, we denote the one-hot label corresponding to l_i^t as $\boldsymbol{y}_i^t$ – this will be used in Eq. 3. By setting a high θ_r, a highly confident sample $\boldsymbol{x}_i$ will be relabeled – this can in turn further enhance the quality of sample selection. Note, that we avoid mis-relabelling open-set noise samples as those tend not to have highly confident average predictions.

3.4 SUPERVISED SELF-SUPERVISED TRAINING

Mixup training of the encoder-classifier using the clean subset With two random samples $\boldsymbol{x}_1, \boldsymbol{y}_1$ and $\boldsymbol{x}_2, \boldsymbol{y}_2$ in the clean subset $(\boldsymbol{X}_c, \boldsymbol{Y}_c)$, a mixed new sample $\boldsymbol{x}_m, \boldsymbol{y}_m$ will be generated by Mixup method (Zhang et al., 2017) as:

$$\lambda \sim Beta(\alpha, \alpha), \lambda' = \max(\lambda, 1 - \lambda), \boldsymbol{x}_m = \lambda' \boldsymbol{x}_1 + (1 - \lambda')\boldsymbol{x}_2, \boldsymbol{y}_m = \lambda' \boldsymbol{y}_1 + (1 - \lambda')\boldsymbol{y}_2 \tag{7}$$

We then apply the normal cross-entropy loss for the new virtual mixed sample:

$$L_{ce} = -\sum_{c=1}^{K} \boldsymbol{y}_m^c log \boldsymbol{z}_m^c \tag{8}$$

where $\boldsymbol{z}_m = Softmax(g(f(\boldsymbol{x}_m)))$. Instead of direct training with samples from the clean subset, we expect that virtual samples generated by Mixup are further away from the dataset samples thus can alleviate the noise memorization effect (Zhang et al., 2017). This loss is back-propagated so as to update both the encoder f and the classification head g.

Self-consistency regularization for training the encoder To fully utilize all the samples, we applied a self-consistency loss motivated by recent non-contrastive self-supervised learning methods (Chen & He, 2021; Grill et al., 2020). With a projector head p and prediction head q, we minimize the negative cosine similarity between two different augmented views from the same sample x_i. Denoting the two output vectors as $\boldsymbol{q}_i \triangleq \boldsymbol{q}(p(f(\boldsymbol{x}_{i1})))$ and $\boldsymbol{p}_i \triangleq \boldsymbol{p}(f(\boldsymbol{x}_{i2}))$:

$$L_{sc} = -\frac{\boldsymbol{q}_i^T \boldsymbol{p}_i}{\|\boldsymbol{q}_i\|_2 \|\boldsymbol{p}_i\|_2}, \tag{9}$$

where $\boldsymbol{x}_{i1}, \boldsymbol{x}_{i2}$ denotes two different augmented views. L_{sc} bears similarity to the commonly used consistency regularization in semi-supervised learning methods, however, the consistency is enforced between the projected features rather than the predictions. This allows us to utilize open-set noise also for training whose true labels are not in the label set. Also, we applied gradient stopping and an extra predictor so as to avoid mode collapse. This loss is back-propagated so as to update the encoder f and contributes to learning a strong feature space from both clean and noisy samples.

Data augmentations Strong augmentations, such as (Cubuk et al., 2020; 2019), have shown to be effective in both supervised and semi-supervised learning (Berthelot et al., 2019; Sohn et al., 2020) and recently, Nishi et al. (2021) validated the benefits of strong augmentations within the DivideMix (Li et al., 2020a) framework. In this work, we define and use three types of augmentations: the original image itself (augmentation type 'none') is used for testing, random cropping and horizontal flipping (augmentation type 'weak'), and the augmentation policy proposed in Cubuk et al. (2019) (augmentation type 'strong'). In the model training phase, by default, we apply 'strong' augmentation for $\boldsymbol{x}_{i1}$, 'weak' augmentation for $\boldsymbol{x}_{i2}$ in L_{sc}, and 'strong' augmentation for $\boldsymbol{x}_1, \boldsymbol{x}_2$ in L_{ce}. In the sample selection and relabelling phase, we apply 'weak' augmentation so as to introduce more variance and alleviate the accumulation of error – this is in contrast to most works that do not apply augmentations at similar stages.

Balanced sampler In order to address possible class imbalances in the noise-agnostic dataset, we normalized the local distribution of class labels, with the inverse of the global class probabilities vector $\boldsymbol{\pi}$ during sample selection (Eq. 3). Similarly, the extracted clean subset $(\boldsymbol{X}_c, \boldsymbol{Y}_c)$ might also potentially suffer from class imbalances. To deal with this, we also use a balanced sampler for training with L_{ce} by oversampling the minority class.

The overall training objective is to minimize a weighted sum of L_{ce} and L_{sc}.

$$L = L_{ce} + wL_{sc} \tag{10}$$

while for all experiments, we fix $w = 1$.

4 EXPERIMENTS

4.1 OVERVIEW

In this section, we conduct extensive experiments on two standard benchmarks with artificial label noise, CIFAR-10 and CIFAR-100, and two real-world datasets, WebVision and ANIMAL-10N. We begin by describing the datasets and some details of the implementations and the experimental settings. In Section 4.2, we conducted extensive ablation experiments to show the robustness of our method w.r.t its hyperparameters with different noise types, noise ratios and dataset. In Section 4.3, we conducted extensive ablation studies to validate the benefits of different modules in our method. In Section 4.4 and 4.5, we compared with the state-of-the-art in synthetic noisy datasets and real-world noisy datasets.

Datasets Following the standard practice, for CIFAR-10 and CIFAR-100, we experimented with two types of artificial noise: *symmetric noise* by randomly replacing labels of all samples using a uniform distribution; and *asymmetric noise* by randomly exchanging labels of visually similar categories, such as Horse $\leftrightarrow$ Deer and Dog $\leftrightarrow$ Cat.

WebVision (Li et al., 2017) is a large-scale dataset of 1000 classes of images crawled from the Web. Following previous work (Jiang et al., 2018; Li et al., 2020a; Ortego et al., 2021), we compare baseline methods on the top 50 classes from Google images Subset of WebVision. The noise ratio that is estimated to be around 20%. ANIMAL-10N (Song et al., 2019) is a smaller and recently proposed real-world dataset consists of 10 classes of animals, that are manually labeled with an error rate that is estimated to be approximately 8%. ANIMAL-10N has similar size characteristics to the CIFAR datasets, with 50000 train images and 10000 test images.

Implementation details We used a PresActResNet-18 (He et al., 2016) as the backbone for all CIFAR10/100 experiments following previous works. Unlike previous methods that use specific warmup settings for CIFAR10/CIFAR100, we train the model from scratch with a linear raising θ_s from 0 to 1 in 20 epochs. θ_r is fixed as 0.8 and the prediction track length L is set to 10 for all CIFAR experiments except in the corresponding ablation part. We train all modules with the same SGD optimizer for 300 epochs with a momentum of 0.9 and a weight decay of 5e-4. The initial learning rate is 0.02 and is controlled by a cosine annealing scheduler by Pytorch. The batchsize is fixed as 128. We set $\alpha = 4$ for all noise settings in Mixup training.

For WebVision, we used a standard resnet18 following Ortego et al. (2021) due to the hardware limitation. We train the network with SGD optimizer for 150 epochs with momentum of 0.9 and a weight decay of 1e-4. The initial learning rate is 0.02 and is controlled by a cosine annealing scheduler. The batchsize is fixed as 64. For ANIMAL-10N, we used VGG-19 (Simonyan & Zisserman, 2014) with batch-normalization following (Song et al., 2019). We train the network with SGD optimizer for 100 epochs with momentum of 0.9 and weight decay of 5e-4. The initial learning rate is 0.02 and and is also controlled by a cosine annealing scheduler. The batchsize is fixed as 128. For both datasets, we train the model from scratch with a linear raising θ_s from 0 to 1 in 10 epochs, while θ_r is fixed as 0.9 and the prediction track length L is set to 10. We set $\alpha = 1$ for Mixup. We report averages of at least two runs on a single Nvidia RTX 3090 GPU card.

4.2 HYPERPARAMETERS ROBUSTNESS

We aim at building a framework which is robust in noise-agnostic dataset scenario and with minimal number of hyperparameters. In this section, we conducted extensive ablation experiments to show the robustness of the few hyperparameters with different noise types, noise ratios and dataset.

Relabelling quality VS relabelling proportion The choice of θ_r and L controls the sample relabelling quality and proportion. Roughly speaking, the lower the θ_r and L, the more samples will be

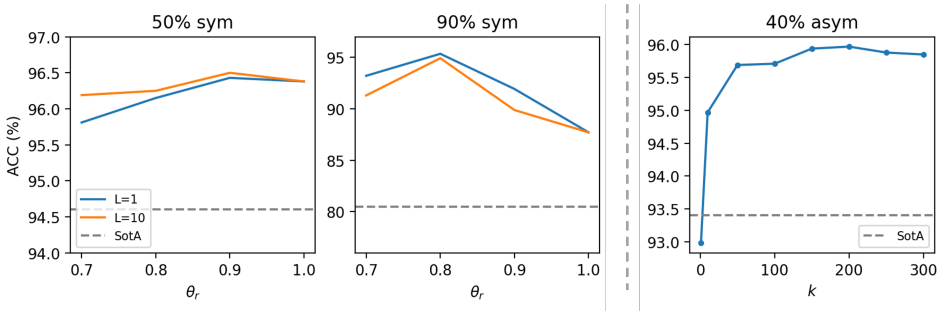

Figure 2: Classification accuracy with different θ_r, L and k of synthetic CIFAR10 datasets. Left: $\theta_r = [0.7, 0.8, 0.9, 1]$, $L = [1, 10]$; Right: $k = [1, 10, 50, 100, 150, 200, 250, 300]$.

relabeled in the training process, which also means that possibly more errors will be introduced. In Fig. 2 we reported performance with different combinations of θ_r and L on the synthetic CIFAR10 noisy dataset. Genrally, our method achieved superior performance than the state-of-the-art with different θ_r and L. For example in CIFAR10 dataset with 90% sym noise, the lowest accuracy is 87.79% – this surpasses the state-of-the-art by $\sim$7%.

Robustness w.r.t k in sample selection The number k controlled the neighborhood size in the sample selection phase. In Fig. 2, we report results with different k for the CIFAR10 dataset with 40% asym noise. Except for too small k which is more sensitive to the noisy samples, the performance is stable and consistently higher than the state-of-the-art.

4.3 Ablations Study

In this section, we conducted extensive ablation studies to validate the benefits of different modules.

Effect of augmentations Table. 1 shows that with stronger augmentations in the model training stage we can achieve consistent better performance at different noise ratio/noise type. For comparison, we also refer to the state-of-the-art result when strong augmentation applied with DivideMix (Nishi et al., 2021). It is clear that we are much better in all augmentation conditions despite the simplicity of our method. Please note that we don't apply cotraining.

Dataset	50% sym CIFAR10						90% sym CIFAR10					
SSR	N		W		S		N		W		S	
SST	W	S	W	S	W	S	W	S	W	S	W	S
ACC (%)	96.18	96.45	96.41	**96.48**	96.14	96.22	93.46	**95.13**	93.70	94.92	90.92	93.85
AugDesc (Nishi et al., 2021)	95.6						91.9					

Table 1: Classification accuracy for different augmentations. N denote 'none' augmentation, W denote 'weak' augmentation, S denote 'strong' augmentation; SSR denote sample selection&relabelling, SST denote supervised self-supervised training.

Effect of self-consistency regularization and Mixup In Table.2 we report the effect of self-contrastive regularization and Mixup. Removing Mixup decreases the performance, especially in high noise ratio and removing self-consistency also lead to degradation. Both help prevent memorization of wrong selections and to explore all samples so as to improve the model robustness.

Effect of balancing strategies To alleviate possible dataset class imbalance, we proposed two dataset balancing strategies in sample selection and model training phase, respectively. In Table. 3 we investigate its effect with controlled asymmetric noise in CIFAR10, a type of noise that is known to create class imbalanced dataset.

Dataset	50% sym	90% sym	40% asym
S3	**96.25**	**94.92**	**95.97**
W/O Mixup	94.13	83.11	93.99
W/O Self-consistency	95.80	93.31	95.54

Table 2: Classification accuracy W/O Mixup and Self-consistency.

Dataset	20% asym	40% asym
S3	**96.54**	**95.97**
W/O balanced selection	96.30	95.70
W/O balanced sampler	96.50	95.84

Table 3: Effect of balancing strategies.

4.4 SYNTHETIC NOISY DATASETS EVALUATION

In this section, we compared our method to several state-of-the-art methods: DivideMix Li et al. (2020a), LossModelling Arazo et al. (2019), Coteaching+ Yu et al. (2019), Mixup Zhang et al. (2017), F-correction Patrini et al. (2017), SELFIE Song et al. (2019), PLC Zhang et al. (2021), PENCIL Yi & Wu (2019), ELR Liu et al. (2020), NCT (Chen et al., 2021), MOIT+ (Ortego et al., 2021), NGC (Wu et al., 2021), ProtoMix (Li et al., 2020b). We show, that the proposed method achieves consistent improvements in all datasets and at all noise types and ratios.

Dataset	CIFAR10					CIFAR100			
Noise type	Symmetric				Assymetric	Symmetric			
Noise ratio	20%	50%	80%	90%	40%	20%	50%	80%	90%
Cross-Entropy	86.8	79.4	62.9	42.7	85.0	62.0	46.7	19.9	10.1
Co-teaching+	89.5	85.7	67.4	47.9	-	65.6	51.8	27.9	13.7
F-correction	86.8	79.8	63.3	42.9	87.2	61.5	46.6	19.9	10.2
Mixup	95.6	87.1	71.6	52.2	-	67.8	57.3	30.8	14.6
PENCIL	92.4	89.1	77.5	58.9	88.5	69.4	57.5	31.1	15.3
LossModelling	94.0	92.0	86.8	69.1	87.4	73.9	66.1	48.2	24.3
DivideMix	96.1	94.6	93.2	76.0	93.4	77.3	74.6	60.2	31.5
ELR	95.8	94.8	93.3	78.7	93.0	77.6	73.6	60.8	33.4
ProtoMix	95.8	94.3	92.4	75.0	91.9	79.1	74.8	57.7	29.3
NGC	95.9	94.5	91.6	80.5	90.6	79.3	75.9	62.7	29.8
S3	**96.6**	**96.3**	**95.7**	**94.9**	**96.0**	**79.7**	**77.2**	**70.4**	**56.8**

Table 4: Evaluation on CIFAR-10 and CIFAR-100 with closed-set noise. Results of other models are copied from Li et al. (2020a); Wu et al. (2021).

Evaluation with controlled closed-set noise We compared to the most competitive works recently. Please note that several of the methods require a per-dataset finetuning of the hyperparameters due to complicated structures and extra modules like model cotraining. Here we report their best results. Table 4 shows results on CIFAR10 and CIFAR100, where it is clear that our method performs clearly better without the use of the additional cotraining techniques. Especially at high noise ratio(90% sym noise), S3 can achieve $\sim$ 95% on CIFAR10 and $\sim$57% on CIFAR100, surpassing the state-of-the-art by 14.5% and 23.4%, respectively. Meanwhile, unlike several previous methods work better in symmetric noise, our method works very well also at the more realistic asymmetric synthetic noise. This indicates that our method is more robust to real-world noise.

Evaluation with combined open-set noise and closed-set noise Tab. 5 shows the performance of our method in a more complex combined noise scenario. Previous methods that are specially designed for open-set noise (Lee et al., 2019; Wang et al., 2018) which degrade rapidly when the open-set noise ratio is decreased from 1 to 0.5. The performance of method without considering open-set noise like Li et al. (2020a) will decrease when the open-set noise ratio is increased. Sachdeva et al. (2021) modified the method of Li et al. (2020a) to deal with combined noise, however report results that are considerably lower than ours.

Methods	Noise ratio	0.3		0.6	
	Open ratio	0.5	1	0.5	1
ILON	Best	87.4	90.4	80.5	83.4
	Last	80.0	87.4	55.2	78.0
RoG	Best	89.8	91.4	84.1	88.2
	Last	85.9	89.8	66.3	82.1
DivideMix	Best	91.5	89.3	91.8	89.0
	Last	90.9	88.7	91.5	88.7
EDM	Best	94.5	92.9	93.4	90.6
	Last	94.0	91.9	92.8	89.4
S3	Best	**96.34**	**96.05**	**94.97**	**94.01**
	Last	**96.13**	**95.95**	**94.81**	**93.54**

Table 5: Evaluation on CIFAR10 with conbined noise. The closed-set noise are generated as symmetric noise while the open-set noise are random samples from CIFAR100. Noise ratio denotes the total noise ratio while the open ratio denotes the proportion of open-set noise. Results of other methods are copied from Sachdeva et al. (2021).

4.5 Real-world noisy datasets evaluation

Method	Network Structure	ACC (%)
DivideMix	Inception-ResNet-v2	77.32
ELR+	Inception-ResNet-v2	77.78
NGC	Inception-ResNet-v2	79.16
Mixup	ResNet-18	74.96
DivideMix	ResNet-18	76.08
ELR	ResNet-18	73.00
MOIT+	ResNet-18	78.76
S3	ResNet-18	**80.08**

Table 6: Testing accuracy on Webvision. Results of other methods are copied from Ortego et al. (2021); Wu et al. (2021).

Finally, in Table. 6 and Table. 7 we show results on the WebVision and ANIMAL-10N datasets respectively. To summarize, our method achieves better performance compared to the current state-of-the-art in both large-scale web-crawled dataset and small-scale human annotated noisy dataset.

Cross-Entropy	SELFIE	PLC	NCT	S3
79.4 ± 0.1	81.8 ± 0.1	83.4 ± 0.4	84.1 ± 0.1	$\mathbf{88.5 \pm 0.1}$

Table 7: Testing accuracy on ANIMAl-10N. Results of other methods are from Chen et al. (2021).

5 Conclusions

In this paper we proposed a method for learning with noisy labels, that relies on a sample selection mechanism, a relabeling mechanism and a training strategy with multi-objective losses that enable us to learn robust features from both noisy and clean samples, and a classifier from only clean or robustly relabeled ones. The proposed method is a simple framework, does not utilize complicated mechanisms such as co-training to deal with self-confirmation bias, and is shown with extensive experiments and ablation studies to be robust to the values of its few hyper-parameters, and to consistently and by large surpass the state-of-the art in both open-set and close-set noise.

REFERENCES

Eric Arazo, Diego Ortego, Paul Albert, Noel O'Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. In *International Conference on Machine Learning*, pp. 312–321. PMLR, 2019.

Dara Bahri, Heinrich Jiang, and Maya Gupta. Deep k-nn for noisy labels. In *International Conference on Machine Learning*, pp. 540–550. PMLR, 2020.

David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *arXiv preprint arXiv:1911.09785*, 2019.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.

Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15750–15758, 2021.

Yingyi Chen, Xi Shen, Shell Xu Hu, and Johan AK Suykens. Boosting co-teaching with compression regularization for label noise. *arXiv preprint arXiv:2104.13766*, 2021.

Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 113–123, 2019.

Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 702–703, 2020.

Aritra Ghosh, Himanshu Kumar, and PS Sastry. Robust loss functions under label noise for deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.

Jacob Goldberger and Ehud Ben-Reuven. Training deep neural-networks using a noise adaptation layer. 2016.

Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.

Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *arXiv preprint arXiv:1804.06872*, 2018.

Jiangfan Han, Ping Luo, and Xiaogang Wang. Deep self-learning from noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5138–5147, 2019.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016.

Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9729–9738, 2020.

Dan Hendrycks, Mantas Mazeika, Duncan Wilson, and Kevin Gimpel. Using trusted data to train deep networks on labels corrupted by severe noise. *arXiv preprint arXiv:1802.05300*, 2018.

Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International Conference on Machine Learning*, pp. 2304–2313. PMLR, 2018.

Kimin Lee, Sukmin Yun, Kibok Lee, Honglak Lee, Bo Li, and Jinwoo Shin. Robust inference via generative classifiers for handling noisy labels. In *International Conference on Machine Learning*, pp. 3763–3772. PMLR, 2019.

Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. *arXiv preprint arXiv:2002.07394*, 2020a.

Junnan Li, Caiming Xiong, and Steven Hoi. Learning from noisy data with robust representation learning. 2020b.

Wen Li, Limin Wang, Wei Li, Eirikur Agustsson, and Luc Van Gool. Webvision database: Visual learning and understanding from web data. *arXiv preprint arXiv:1708.02862*, 2017.

Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. *arXiv preprint arXiv:2007.00151*, 2020.

Eran Malach and Shai Shalev-Shwartz. Decoupling" when to update" from" how to update". *arXiv preprint arXiv:1706.02613*, 2017.

Kento Nishi, Yi Ding, Alex Rich, and Tobias Hollerer. Augmentation strategies for learning with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8022–8031, 2021.

Diego Ortego, Eric Arazo, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Multi-objective interpolation training for robustness to label noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6606–6615, 2021.

Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1944–1952, 2017.

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.

Ragav Sachdeva, Filipe R Cordeiro, Vasileios Belagiannis, Ian Reid, and Gustavo Carneiro. Evidentialmix: Learning with combined open-set and closed-set noisy labels. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 3607–3615, 2021.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685*, 2020.

Hwanjun Song, Minseok Kim, and Jae-Gil Lee. Selfie: Refurbishing unclean samples for robust deep learning. In *International Conference on Machine Learning*, pp. 5907–5915. PMLR, 2019.

Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.

Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, and Shu-Tao Xia. Iterative learning with open-set noisy labels. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8688–8696, 2018.

Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 322–330, 2019.

Zhi-Fan Wu, Tong Wei, Jianwen Jiang, Chaojie Mao, Mingqian Tang, and Yu-Feng Li. Ngc: A unified framework for learning with open-world noisy data. *arXiv preprint arXiv:2108.11035*, 2021.

Kun Yi and Jianxin Wu. Probabilistic end-to-end noise correction for learning with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7017–7025, 2019.

Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *International Conference on Machine Learning*, pp. 7164–7173. PMLR, 2019.

Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.

Yikai Zhang, Songzhu Zheng, Pengxiang Wu, Mayank Goswami, and Chao Chen. Learning with feature-dependent label noise: A progressive approach. *arXiv preprint arXiv:2103.07756*, 2021.

Zhilu Zhang and Mert R Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. *arXiv preprint arXiv:1805.07836*, 2018.

Evgenii Zheltonozhskii, Chaim Baskin, Avi Mendelson, Alex M Bronstein, and Or Litany. Contrast to divide: Self-supervised pre-training for learning with noisy labels. *arXiv preprint arXiv:2103.13646*, 2021.