

MILD: Modeling the Instance Learning Dynamics for Learning with Noisy Labels

Chuanyang Hu¹, Shipeng Yan¹, Zhitong Gao¹ and Xuming He^{1,2}

¹ShanghaiTech University

²Shanghai Engineering Research Center of Intelligent Vision and Imaging

{huchy3, yanshp, gaozht, hexm}@shanghaitech.edu.cn

Abstract

Despite deep learning has achieved great success, it often relies on a large amount of training data with accurate labels, which are expensive and time-consuming to collect. A prominent direction to reduce the cost is to learn with noisy labels, which are ubiquitous in the real-world applications. A critical challenge for such a learning task is to reduce the effect of network memorization on the falsely-labeled data. In this work, we propose an iterative selection approach based on the Weibull mixture model, which identifies clean data by considering the overall learning dynamics of each data instance. In contrast to the previous small-loss heuristics, we leverage the observation that deep network is easy to memorize and hard to forget clean data. In particular, we measure the difficulty of *memorization* and *forgetting* for each instance via the transition times between being misclassified and being memorized in training, and integrate them into a novel metric for selection. Based on the proposed metric, we retain a subset of identified clean data and repeat the selection procedure to iteratively refine the clean subset, which is finally used for model training. To validate our method, we perform extensive experiments on synthetic noisy datasets and real-world web data, and our strategy outperforms existing noisy-label learning methods.

1 Introduction

Deep neural networks has achieved impressive success in a wide spectrum of machine learning tasks. However, they usually rely on a large amount of data with accurate supervision, which is expensive and time-consuming to obtain in practice. In real-world applications, a common strategy for collecting large-scale data is to use online queries [Blum *et al.*, 2003] and/or crowdsourcing [Yan *et al.*, 2014], which are typically cheap but the resulting annotations are noisy. Nonetheless, recent studies [Arpit *et al.*, 2017] demonstrate that the capacity of deep networks is sufficient to fit even randomly generated labels, which hinders the generalization when the deep networks learn with noisy labels. Therefore, it is essential to develop robust training methods with noisy labels for deep neu-

ral networks, which can benefit a broad scope of applications such as medical diagnosis [Li *et al.*, 2021] and autonomous driving systems [Feng *et al.*, 2020].

In order to alleviate the impact of noisy data, it is critical to identify the clean subset of data, also referred to as sample selection [Song *et al.*, 2022]. By removing mislabeled data as much as possible, we are able to reduce the impact of mislabeled data and robustify the learning process. During this process, perhaps the most important aspect of sample selection is the validity of the data selection metric. There are some existing works exploring the design of metric. Small-loss metric is the most popular selection approach, which considers samples with small loss values to be clean. This is based on the observation that deep networks learn easy patterns first, and then they overfit to the noisy pattern [Arpit *et al.*, 2017]. Recently, FINE [Kim *et al.*, 2021] utilizes the principal component of feature representation and split the data according to the projection to the principal component. However, those criteria depends on the current training state of models for removing corrupted data and are often unstable.

To tackle the above-mentioned issues, we propose a novel sample selection framework for learning from noisy labels. In contrast to prior works that simply utilize current state of a trained model, we adopt a novel perspective that exploits the overall learning dynamic to select clean training data. Our main idea is to design a new metric to differentiate between the clean and corrupted data based on the observation that the network is easy to memorize and hard to forget clean samples. To this end, we introduce memorization and forgetting measurements to represent the overall learning dynamic for each instance, which are integrated into our metric. Moreover, we build a mixture model of the metric distribution to determine the selection threshold, and perform multiple training rounds to iteratively refine the selected clean data.

Specifically, at each round, we distinguish examples based on the metric which takes into account the *memorization* and *forgetting* property of each data during the learning process. To this end, we record the prediction history of each instance during the training as a sequence, and measure the difficulty of memorization by the transition times from being misclassified to memorized and the difficulty of forgetting by the transition times from being memorized to misclassified. We then develop a selection metric by combining the memorization and forgetting measurements, as they play a complementary

role in describing the learning dynamics. To perform selection, we adopt a Weibull mixture model to fit the metric distribution and determine the selection threshold. Finally, to take advantage of the falsely-labeled data, we further extend our sample selection approach by incorporating semi-supervised techniques in the model training.

To validate the effectiveness of our proposed approach, we perform extensive experiments on CIFAR-10 [Krizhevsky *et al.*, 2009], CIFAR-100 [Krizhevsky *et al.*, 2009], Mini-WebVision [Li *et al.*, 2017] and Mini-ImageNet [Jiang *et al.*, 2020] datasets under different noise ratios and types. The experimental results show that our method consistently outperforms other methods, thereby demonstrating its superiority. Besides, we also perform detailed analysis of our approach to provide more insights on the new metric.

Our main contributions can be summarized as follows,

- We propose a novel selection metric by taking advantage of instance learning dynamics to discriminate clean and corrupted data. Our strategy incorporates the learning dynamics from both instance memorization and forgetting perspective.
- We develop a novel sample-selection approach based on a Weibull mixture model on the metric distribution, which refines the clean data selection via an iterative training strategy.
- Our approach achieves the state of the art on five popular noisy image classification benchmarks. We also show that our method can be effectively combined with multiple semi-supervised methods.

2 Related Works

In this section, we briefly review existing relevant works on learning with noisy labels. Noisy data interfere with the learning of clean data [Arpit *et al.*, 2017]. To reduce the influence, the approach can be mainly summarized as follows: sample-selection methods, semi-supervised based methods and other noisy label approaches.

Sample-selection methods To avoid the influence of data noises, sample selection approaches [Song *et al.*, 2022] to detect and filter falsely-labeled data are proposed. It is critical to choose an appropriate criterion to separate clean and corrupted samples when deciding on a sample selection procedure. Small-loss [Arpit *et al.*, 2017] is the most widely utilized criterion. The empirical findings of study [Arpit *et al.*, 2017] indicate that deep neural networks acquire progressively more sophisticated hypotheses and learn the simple patterns of clean data prior to fitting noisy data, supporting the reasonableness of small-loss criterion. Few works [Nguyen *et al.*, 2019; Zhou *et al.*, 2020] adopt the exponential moving average of loss to select clean data. FINE [Kim *et al.*, 2021] filters noisy instances by their eigenvectors which computes the alignment of image feature and the principal component of gram matrix in the learned representation space. [Zhu *et al.*, 2021] uses the number of projected gradient steps to attack a data point to assess a data point’s distance to its class boundary, which can then be used to distinguish between clean and corrupted data.

In addition to the selection criteria, there are also works [Song *et al.*, 2022] that employ several networks to collaborate with one another or multi-round learning to iteratively enhance the selected dataset. MentorNet [Jiang *et al.*, 2018] introduce mentor network to supervise the training of a student network, where MentorNet provides a curriculum for StudentNet’s training to concentrate on the clean samples. JoCoR [Wei *et al.*, 2020] employs two networks to make predictions on the same mini-batch data and improves the agreement of two network via computing a joint loss with Co-Regularization for each training instance. In Co-teaching [Han *et al.*, 2018a], each network selects its small-loss instances and feeds them to its peer network, which can filter various sorts of noise errors because of the distinct each network’s learning ability. Multi-round learning refines the selected subset repeatedly, which keeps improving the quality of selected set as the number of rounds increase. For multi-round learning, INCV [Chen *et al.*, 2019] introduces cross-validation to randomly partition noisy training data in order to determine clean data for the subsequent round training. It also employs Co-teaching strategy to exploit the selected data to train deep networks. ITLM [Shen and Sanghavi, 2019] iteratively minimize the trimmed loss by repeating multiple rounds. For each round, it selects small loss samples to build a more clean dataset and then retrain the deep network using selected data. In order to gain from falsely-labeled data, semi-supervised techniques are also explored to combine with sample selection to further improve performance [Li *et al.*, 2020; Ortego *et al.*, 2021]. In contrast to previous approaches simply considering the present state of the model, we distinguish the clean and corrupted data based on their learning dynamics and then do multi-round learning.

Semi-supervised learning There are some works [Li *et al.*, 2020; Ortego *et al.*, 2021; Song *et al.*, 2019; Zhou *et al.*, 2020] trying to leverage the falsely-labeled data in a semi-supervised(SSL) manner. DivideMix [Li *et al.*, 2020] first perform sample selection to eliminate the samples that are more likely to be corrupted, and utilize the noisy samples as unlabeled data to regularize training in a SSL manner with MixMatch [Berthelot *et al.*, 2019]. MOiT [Ortego *et al.*, 2021] introduces interpolated supervised contrastive loss with MixUp [Zhang *et al.*, 2017a] to interpolate data points in order to enhance model performance. ScanMix [Zhang *et al.*, 2017a] develops from DivideMix and introduce Expectation-Maximization(EM) based semantic clustering to improve the robustness to noisy label. SELFIE [Song *et al.*, 2019] corrects refurbishable example, which can be corrected with high precision. Concretely, it considers the example with consistent label predictions as refurbishable examples. In this study, we also demonstrate that our method can be used effectively with a variety of semi-supervised methodologies.

Other noisy label methods [Ghosh *et al.*, 2017] demonstrates that risk minimization with mean absolute error(MAE) can be more tolerant of noisy labels, but is not appropriate for complicated data. Generalized cross entropy(GCE) [Zhang and Sabuncu, 2018] is a more general kind of noise-robust loss that combines the benefits of categorical cross entropy loss(CCE) and MAE. APL [Ma *et al.*, 2020] further com-

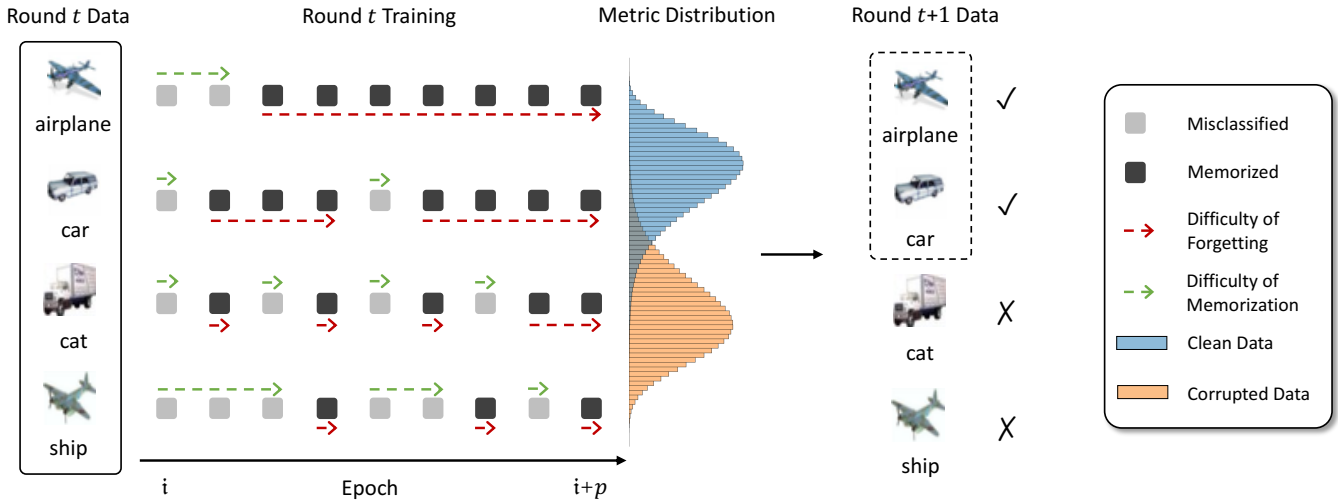


Figure 1: The overview of each sample selection round. The input dataset at left is noisy like the truck image with cat label. To model learning dynamic for each instance, we record the learning status at each epoch as a sequence, where gray and black squares indicate misclassified and memorized. The p is the number of epochs for current round. Note that the total length of green and red arrows indicate the difficulties of memorization and forgetting. Then we compute the selection metric based on the difficulties. Finally, we identify and select the clean data, and then use them for subsequent round.

bines existing robust loss functions that mutually boost each other. Other works utilize implicit/explicit regularization [Liu *et al.*, 2020] or transition matrix learning[?] to improve the robustness of the model. MixUp [Zhang *et al.*, 2017a] training strategy that greatly prevents memorization to noise. Self-training [Arazo *et al.*, 2019] has been employed to generate pseudo label to correct the falsely-labeled data.

3 Method

In this section, we present our approach to the robust learning of image classification tasks, aiming to improve the generalization of classification model. To this end, we propose a sample selection approach leveraging the instance learning dynamics to separate the clean and falsely-labeled data. Below we first present the overall sample selection framework in Sec. 3.1. Then we introduce the design of proposed selection metric using learning dynamics in Sec. 3.2 and the selection strategy based on a Weibull mixture model in Sec. 3.3. Finally, we present the exploitation of corrupted data in Sec. 3.4.

3.1 Multi-round Sample Selection Framework

The noisy training set $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^{|\mathcal{D}|}$ where x_i, y_i is the i -th image and its corresponding label. Note that there are some images that are falsely labeled. During training, our sample selection will repeat for multiple rounds. We adopt the categorical cross entropy loss on current training set \mathcal{D} to update the model. Algorithm 1 describes the selection procedure of one round. As training progress, we continually update selection metric values. At the end of each round, we apply a Weibull mixture model to identify and choose the clean data as the training set \mathcal{D}' for next round.

3.2 Selection Metric Design

Model the instance learning dynamics In this work, we distinguish the corrupted and clean data via their learning dynamics. First, we represent the instance learning dynamic with instance prediction sequence S_i for each instance x_i , which records the instance learning status for each epoch. Concretely, given an image x_i , the model outputs the prediction label $\hat{y} = \arg \max_c p(y = c|x)$. The example x_i is deemed to be *memorized* if the prediction label \hat{y} equals the dataset label y . If not, it is referred to as *misclassified*. Based on the learning status, we can obtain the k -th element of instance prediction sequence $S_i^k = \mathbb{1}_{\hat{y}_i=y_i}$, which corresponds to the epoch k . Besides, the indicator function $\mathbb{1}_{\hat{y}_i=y_i}$ is 1 when $\hat{y}_i = y_i$, otherwise 0. Note that we record instance’s learning status before updating the model with gradient. It is worth noting that as learning progresses, the proportion of memorized examples keep increasing.

Selection metric We can further summarize the instance learning dynamics to distinguish between clean and corrupted data. Note that [Arpit *et al.*, 2017] finds that the model often first learns the patterns of clean data and then overfits to corrupted data. Besides, [Toneva *et al.*, 2018] observes that the corrupted data is more easy to forget compared to clean data. Motivated by these observations, we can hypothesis that clean samples are both easy to learn and hard to forget. Furthermore, we estimate the difficulty of memorization and forgetting based on the instance prediction sequence S_i . Therefore, we propose a novel selection metric by combining them as we believe that memorization and forgetting are complementary to separate clean and corrupted data.

The intuitive solution to estimate the difficulty of memorization is to measure the average transition times from misclassified to memorized. Similarly, the difficulty of forgetting can be estimated via the average transition times

from memorized to misclassified. An instance prediction sequence S_i^k may have several memorization and forgetting events during learning, as shown in Fig. 1. Therefore, we split the instance sequence S_i into continuous segments $G_i = \{g_{i,1}^u, \dots, g_{i,n_i^u}^u\} \cup \{g_{i,1}^l, \dots, g_{i,n_i^l}^l\}$ with same indicator value S_i^k in segment, where the segments g_j are disjoint. It is worth noting that $g_{i,j}^u$ means the segment where the instances have not been memorized. Similarly, $g_{i,j}^l$ refers to the segment where the instances are memorized by the model. n_i^u and n_i^l are the number of misclassified segments g_i^u and memorized segments g_i^l , respectively. Therefore, given instance prediction sequence S_i , we define the difficulty of memorization M_i and forgetting F_i as follows

$$M_i = \frac{\sum_{j=0}^{n_i^u} |g_j^u|}{n_i^u}, \quad F_i = \frac{\sum_{j=0}^{n_i^l} |g_j^l|}{n_i^l} \quad (1)$$

where $|\cdot|$ represents the segment length. Considering that clean instances typically have higher F_i and smaller M_i , we create the selection metric by combining the memorization and forgetting in the way described below

$$C_i^F = M_i - \lambda F_i \quad (2)$$

where λ is the coefficient to adjust the influence of memorization and forgetting. Typically, we take $\lambda = 1$ in experiment. Moreover, our metric can be simplified by omitting the number of segments from equation 1. Empirically we find that the simplified version is very close to the full version C^F . Therefore, our metric after simplification is stated as follows

$$C_i = \tilde{M}_i - \lambda \tilde{F}_i = n_i^u M_i - \lambda n_i^l F_i = \sum_{j=0}^{n_i^u} |g_j^u| - \lambda \sum_{j=0}^{n_i^l} |g_j^l| \quad (3)$$

where \tilde{M}_i , \tilde{F}_i are to measure the difficulty of memorization and forgetting when ignoring the number of segments.

3.3 Selection Strategy

Based on empirical observation, we adopt the Weibull distribution to represent the distribution of our metric scores on the data with clean or noisy labels. The probability density function of Weibull distribution is defined as follows

$$\phi(x|\alpha, \beta) = \frac{\beta}{\alpha} \left(\frac{x}{\alpha}\right)^{\beta-1} e^{-(x/\alpha)^\beta} \quad (4)$$

where x is the metric distribution, α is the scale parameter and β is the shape parameter.

Given the metric distribution F_C , we use the Weibull mixture model which contains two components to fit the metric distributions of clean and falsely-labeled data.

$$P(x|\theta) = k_1 * \phi_c(x|\alpha_1, \beta_1) + k_2 * \phi_f(x|\alpha_2, \beta_2) \quad (5)$$

where $\theta = (\alpha_1, \beta_1, \alpha_2, \beta_2)$, k_1 and k_2 are the coefficients of each component of the mixture model which subjects to $k_1 + k_2 = 1$, ϕ_c and ϕ_f represent the distribution of clean and falsely-labeled data respectively. According to the definition of our metric, the component of the Weibull mixture model

Algorithm 1 Pseudocode of Sample Selection

Input: Current training set D , Model M , Number of Epoch E

```

1: Initialization:  $e \leftarrow 0$ 
2: for  $e = 0; e < E; e++$  do
3:   for Batch data  $d$  in  $D$  do
4:     Prediction  $pred \leftarrow M(d)$ 
5:     Update prediction sequence  $S$  of instances in  $d$ 
6:     Compute loss
7:     Loss back propagation and optimize  $M$ 
8:   end for
9: end for
10: Distribution  $F_C \leftarrow$  calculate metric  $C$  of each sample
11: Fit weibull mixture model to  $F_C$ 
12: Obtain selection threshold  $\tau = \alpha_2$ 
13:  $D' = \{(x_i, y_i) | C_i < \tau\}_{i=1}^{|D|}$ 
    
```

Output: training set D' and Model M for next round

corresponding to the falsely-labeled data is typically has a larger mean value, as the falsely-labeled data is difficult to memorize and easy to forget, resulting in a relatively large metric value. Consequently, we are able to determine which component corresponds to the clean or falsely-labeled data.

Finally, we choose the threshold $\tau = \alpha_2$ and select all the training data instances with a score smaller than τ for the next round. The reason for not choosing the position where the two components have the same probability as the threshold is because the performance will deteriorate due to the sharp drop of recall in multiple rounds. Empirically, we found choosing the threshold τ can be beneficial to maintain a high recall while also reducing the computational complexity.

3.4 Utilizing Data with Incorrect Labels

Note that our sample selection approach discards the identified falsely-labeled data and uses the selected clean data for the model training. Our method can also be extended and combined with existing semi-supervised methods, where we regard the corrupted data as unlabeled data. Concretely, we first apply our sample selection framework on the noisy training set to separate it into clean data and corrupted data. Then we utilize a semi-supervised learning strategy on the combination of the clean data and the unlabeled data derived from the corrupted to further improve the generalization of model.

4 Experiments

In this section, we conduct extensive experiments to validate the effectiveness of our method. Specifically, we evaluate our method on five popularly used datasets including CIFAR-10, CIFAR-100, Mini-ImageNet, Mini-Webvision and CIFAR-N. Due to the page limitation, the results of CIFAR-N dataset which is also the web noise is in the supplementary material. We also perform a series of ablation studies and analysis to provide more insights for our method. We first introduce the experiment setup in Sec. 4.1, followed by the experimental results on synthetic noise datasets in Sec. 4.2. Then we present the experimental results on web noise datasets in Sec. 4.3.

Dataset	CIFAR-10				CIFAR-100			
Noisy Type	Symmetry		Asymmetry		Symmetry		Asymmetry	
Noise Ratio	20	50	80	40	20	50	80	40
Standard	87.0 ± 0.1	78.2 ± 0.8	53.8 ± 1.0	85.0 ± 0.0	58.7 ± 0.3	42.5 ± 0.3	18.1 ± 0.8	42.7 ± 0.6
Bootstrap[Reed <i>et al.</i> , 2014]	86.2 ± 0.2	-	54.1 ± 1.3	81.2 ± 1.5	58.3 ± 0.2	-	21.6 ± 1.0	45.1 ± 0.6
Forward[Patrini <i>et al.</i> , 2017]	88.0 ± 0.4	-	54.6 ± 0.4	83.6 ± 0.6	39.2 ± 2.6	-	9.0 ± 0.6	34.4 ± 1.9
Co-teaching[Han <i>et al.</i> , 2018b]	89.3 ± 0.3	83.3 ± 0.6	66.3 ± 1.5	88.4 ± 2.8	63.4 ± 0.0	49.1 ± 0.4	20.5 ± 1.3	47.7 ± 1.2
Co-teaching+[Yu <i>et al.</i> , 2019]	89.1 ± 0.5	84.9 ± 0.4	63.8 ± 2.3	86.5 ± 1.2	59.2 ± 0.4	47.1 ± 0.3	20.2 ± 0.9	44.7 ± 0.6
TopoFilter[Wu <i>et al.</i> , 2020]	90.4 ± 0.2	86.8 ± 0.3	46.8 ± 1.0	87.5 ± 0.4	66.9 ± 0.4	53.4 ± 1.8	18.3 ± 1.7	56.6 ± 0.5
CRUST[Mirzasoleiman <i>et al.</i> , 2020]	89.4 ± 0.2	87.0 ± 0.1	64.8 ± 1.5	82.4 ± 0.0	69.3 ± 0.2	62.3 ± 0.2	21.7 ± 0.7	56.1 ± 0.5
JoCoR[Wei <i>et al.</i> , 2020] [†]	90.2 ± 0.4	47.9 ± 2.2	24.1 ± 1.9	74.1 ± 0.3	64.8 ± 1.2	51.2 ± 1.4	9.2 ± 2.6	42.1 ± 0.8
FINE[Kim <i>et al.</i> , 2021]	91.0 ± 0.1	87.3 ± 0.2	69.4 ± 1.1	89.5 ± 0.1	70.3 ± 0.2	64.2 ± 0.5	25.6 ± 1.2	61.7 ± 1.0
MILD	93.0 ± 0.2	88.7 ± 0.2	79.1 ± 0.5	89.8 ± 0.3	74.2 ± 0.2	67.3 ± 0.3	36.0 ± 0.3	69.9 ± 0.6

Table 1: Test accuracy (%) on CIFAR-10 and CIFAR-100 under Resnet34 backbone with different noise ratios and types. The average accuracies and standard deviations over three trials are reported. † means the results are reproduced.

Finally, we provide the ablation studies and analysis of our method in Sec. 4.4.

4.1 Experiment Setup

Datasets To systematically evaluate the noisy label methods, we conduct exhaustive experiments on both synthetic and web noise datasets. For synthetic noise, we follow the protocol proposed in [Kim *et al.*, 2021] to generate symmetric and asymmetric noise on CIFAR-10 [Krizhevsky *et al.*, 2009] and CIFAR-100 [Krizhevsky *et al.*, 2009] datasets. Symmetric noise is provided by randomly flipping the label of a fraction of samples. In order to make asymmetric noise, sample labels are changed to a specified other class. Specifically, the mapping relationship for asymmetric noise in CIFAR-10 is as follows: truck → car, bird → airplane, deer → horse, cat → dog and dog → cat. For CIFAR-100 dataset, the asymmetric noise is generated by circularly flipping each class to the next class inside super-classes. For real scenarios with web noises, We adopt Mini-Imagenet, Mini-Webvision and CIFAR-N datasets for evaluation. Besides, Mini-ImageNet [Jiang *et al.*, 2020] collects noisy images from Internet. Specifically, it annotates the correctness of the collected image labels manually. The dataset includes around 50,000 training images and 5,000 val images for 100 classes. Mini-Webvision [Li *et al.*, 2017] is derived from Webvision dataset [Li *et al.*, 2017] by picking top 50 categories of google images. Note that Webvision dataset is built through keyword searches on Flickr and Google Images, with the 1000 ImageNet categories selected as keywords. Mini-Webvision dataset consists of 65,944 images for training and 2500 images for test. CIFAR-N dataset involves realistic annotation noise and more details are in supplementary material.

Comparison methods To demonstrate the effectiveness of our method, we compare our approach with existing sample selection method like FINE [Kim *et al.*, 2021], Jo-

CoR [Wei *et al.*, 2020] and semi-supervised based approaches like DivideMix [Li *et al.*, 2020] and MOIT [Ortego *et al.*, 2021]. It is worth noting that for fair comparison with semi-supervised based approaches, we expand our approach by employing the same semi-supervised techniques. Concretely, MixMatch [Berthelot *et al.*, 2019] used in DivideMix is applied on our approach, named as MILD-MM. MixUp [Zhang *et al.*, 2017a] and contrastive learning used in MOIT are introduced into our approach, named as MILD-MUCL. Note that we take the same configurations with MOIT and DivideMix for the hyper-parameters about semi-supervised.

Implementation details Our implementation is based on PyTorch. For CIFAR-10 and CIFAR-100 dataset, the image size is 32x32. We utilize the ResNet-34 for CIFAR-10 and CIFAR-100 and adopt the same data augmentations as FINE[Kim *et al.*, 2021], including random cropping and random horizontal flipping. We adopt the SGD optimizer with an initial learning rate of 0.01, a momentum of 0.9. For Mini-Webvision and Mini-ImageNet, the image size is resized to 224x224 and 84x84, respectively. We conduct experiments using both Inception-ResNet v2 [Szegedy *et al.*, 2017] and 18-layers ResNet [He *et al.*, 2016] on Mini-Webvision dataset following [Li *et al.*, 2020]. Besides, 18-layers ResNet is used for Mini-Imagenet following [Ortego *et al.*, 2021]. We use the SGD optimizer with an initial learning rate of 0.1, a momentum of 0.9. For all datasets, we use the cosine-annealing learning rate scheduler for each round and the batch size is set to 128. The training epochs of each round is 20, 50, 100, 100 for CIFAR-10, CIFAR-100, Mini-Imagenet and Mini-Webvision, respectively. Following FINE [Kim *et al.*, 2021], we tune the hyper-parameters by an additional validation set for CIFAR-10 and CIFAR-100. Also following DivideMix [Li *et al.*, 2020], the validation set is provided to tune the hyper-parameters for web noise datasets like Mini-Webvision and Mini-ImageNet.

Methods	Semi-Tech	Noise Ratio		
		20%	40%	80%
Mix[Zhang <i>et al.</i> , 2017b]	mixup	54.60	50.40	37.32
DivideMix[Li <i>et al.</i> , 2020]	mixmatch	50.30	50.94	35.42
ELR[Liu <i>et al.</i> , 2020]	temporal ensembling	58.10	50.62	41.68
MOiT[Ortego <i>et al.</i> , 2021] [†]	mixup + contrastive learning	61.62	58.10	43.84
MILD-MM	mixmatch	62.56	59.14	44.78
MILD-MUCL	mixup + contrastive learning	62.46	59.12	44.68

Table 2: Test accuracy (%) at final epoch on Mini-Imagenet dataset. † denotes the results are obtained by running their code.

Methods	Semi-Tech	Backbone	Accuracy	Backbone	Accuracy
Co-teaching[Han <i>et al.</i> , 2018b]	-	Inception-resnet v2	63.58	Resnet18	-
FINE[Kim <i>et al.</i> , 2021]	-		75.24		-
MILD	-	Inception-resnet v2	76.63	Resnet18	73.42
Mix[Zhang <i>et al.</i> , 2017b]	mixup		-		73.76
ELR[Liu <i>et al.</i> , 2020]	temporal ensembling	Inception-resnet v2	76.26	Resnet18	71.88
DivideMix[Li <i>et al.</i> , 2020]	mixmatch		77.32		74.64
MOiT[Ortego <i>et al.</i> , 2021]	mixup + contrastive learning		80.00 [†]		77.76
MILD-MM	mixmatch	Inception-resnet v2	78.88	Resnet18	78.08
MILD-MUCL	mixup + contrastive learning		80.52		78.52

Table 3: Test accuracy (%) at final epoch on Mini-Webvision. † denotes the results are obtained by running their code.

4.2 Experiments on Synthetic Noise Datasets

For synthetic noise datasets like CIFAR-10 and CIFAR-100, we first compare our approach with sample selection approaches, shown in Table. 1. We can see that our method consistently outperforms other selection methods under different types and noise rates on both datasets. It is worth noting that in general, the performance gain produced by our method improves as the noise rate increases. Specifically, compared to FINE, we improve the accuracy from 25.6 to 36.0(+10.4%) under the symmetric 80% noise of CIFAR-100 dataset.

4.3 Experiments on Web Noise Datasets

For web noise dataset, we conduct extensive experiments on Mini-ImageNet and Mini-Webvision. For Mini-ImageNet dataset, we compare our approach with semi-supervised based methods in Table 2. It is shown that our method consistently surpasses other methods under different noise ratios in real noise scenario. Specifically, MILD-MUCL surpasses MOiT with 0.84% under 20% noise ratio, and MILD-MM outperforms than DivideMix with 9.36% accuracy on 40% noise ratio.

For Mini-Webvision dataset, we compare both sample-selection and semi-supervised based approaches in Table 3. Specifically, we compare two different backbones including ResNet18 and Inception-ResNet v2. It is evident that our method consistently performs better than other sample-

selection methods. Specifically, we surpass FINE with 1.4% performance improvement on Inception-ResNet v2 backbone. Besides, using the same semi-supervised technique, our method consistently outperforms than other semi-supervised methods. Concretely, our method(MILD-MM) surpasses DivideMix with 1.56% performance gain using Inception-ResNet v2 backbone and 3.44% improvement on ResNet18 backbone, demonstrating the robustness of our technique to the model architecture. Compared to MOiT, MILD-MUCL achieves 0.52% improvement with ResNet18 backbone, which demonstrates the superiority of our proposed approach. In summary, we can see that our approach can benefit from kinds of semi-supervised techniques.

4.4 Ablation Study and Analysis

Visualization. Figure. 2 shows the histogram of metric values for clean data(blue columns) and falsely-labeled(orange columns) data and the fitting result of mixture model under 20% symmetric noise on CIFAR-100 dataset. We can see that with our proposed metric and mixture model, the clean data and falsely-labeled data can be effectively separated, which can be further utilized to remove falsely-labeled data to alleviate the influence of noises on subsequent training.

Component analysis. To provided more insights on the effectiveness of the components in our framework, we conduct ablation experiments in Table.4. We use another simple se-

Memorization	Forgetting	Mixture Model	CIFAR-10			CIFAR-100		
			Precision	Recall	Accuracy	Precision	Recall	Accuracy
✓			86.75	70.91	72.76	77.85	36.92	30.37
✓	✓		92.73	75.06	77.56	79.61	37.79	33.29
✓	✓	✓	93.37	71.85	79.11	85.41	37.17	36.03

Table 4: Components ablation study on symmetric 80% noise.

metric	CIFAR-10			CIFAR-100		
	Precision	Recall	Accuracy	Precision	Recall	Accuracy
Small-loss[Shen and Sanghavi, 2019]	97.75	95.31	92.30	88.37	88.15	67.79
Eigenvector[Kim <i>et al.</i> , 2021]	99.91	77.05	91.04	99.56	82.43	70.39
MILD	98.35	98.96	93.03	95.74	98.78	74.17

Table 5: Performance (%) comparison with other sample selection methods on symmetric 20% noise.

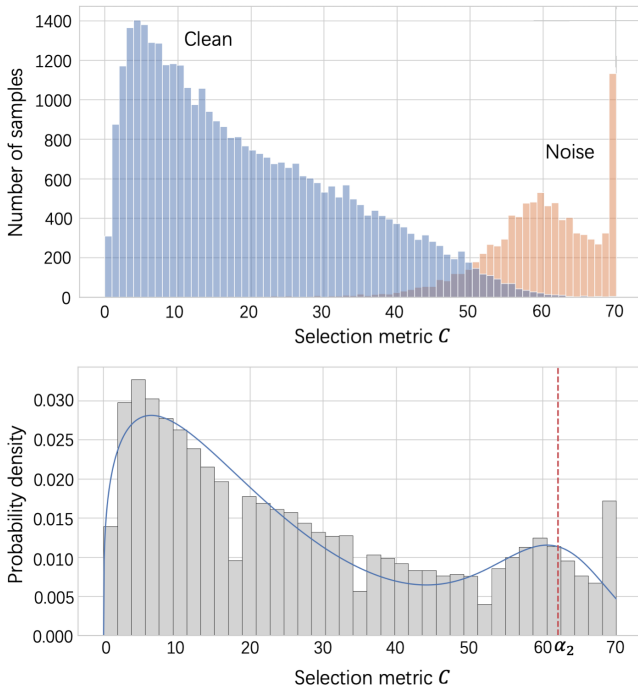


Figure 2: The visualization of metric distribution and mixture model on CIFAR-100 symmetric 20% noise.

lection strategy to compare with the mixture model. Based on the metric value C_i , we select $R|D|$ instances with smallest C_i as the training set for subsequent round, where the selection ratio R is set to 0.9. As Table.4 shows, it is indicated that combining memorization and forgetting can improve the selection quality, which highlights the importance of considering overall learning dynamics. Meanwhile, we can see that the mixture model can further improve the performance, which indicates that the mixture model can benefit the sample selection.

Comparison with sample selection metrics. We also conduct the metric comparison with small-loss and eigenvector-based[Kim *et al.*, 2021] in same sample selection framework. For example, we perform multi-round learning with small loss. Note that in each round, different from existing small-loss methods, we select the best epoch with highest test accuracy for small loss selection. Table.5 compares the precision, recall of the selection data and test accuracy for different methods on 20% symmetric noise of CIFAR datasets. We can see that our method achieves better recall when keeping the comparable precision with Eigenvector. Specifically, we surpass Eigenvector with 16.35% for recall on CIFAR-100.

5 Conclusion

Deep networks easily memorize the patterns of noisy data, which causes performance degradation. To alleviate the influence of falsely-labeled data, it is critical to identify correctly labeled data from a noisy training dataset. In this work, we propose a novel sample selection framework to differentiate clean and corrupted data via their learning dynamics. We measure the difficulty of memorization and forgetting to represent learning dynamics. Concretely, the difficulty of memorization is estimated with the transition times from misclassified to memorized. Besides, the difficulty of forgetting is estimated with the transition times from memorized to misclassified. We combine them to build the selection criterion as memorization and forgetting are complementary to identify clean data. We alternating the model training and selecting clean data to keep improving the quality of selected dataset. To verify the effectiveness of proposed method, we conduct experiments on synthetic noise datasets and web noise datasets. The experimental results show that our method consistently outperforms than other comparison methods. Despite the fact that our method consistently beats other methods, it is worth noting that there is still a considerable room for improvement given the recall of our method at extremely high noise rate, such as 80%. For negative social impact, it is possible to be leveraged by malicious applications for our method to improve their application.

Acknowledgments

This work was supported by Shanghai Science and Technology Program 21010502700 and Shanghai Frontiers Science Center of Human-centered Artificial Intelligence.

References

- [Arazo *et al.*, 2019] 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(ICML)*, 2019.
- [Arpit *et al.*, 2017] Devansh Arpit, Stanisław Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *International conference on machine learning(ICML)*, 2017.
- [Berthelot *et al.*, 2019] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in Neural Information Processing Systems*, 2019.
- [Blum *et al.*, 2003] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *Journal of the ACM (JACM)*, 2003.
- [Chen *et al.*, 2019] Pengfei Chen, Ben Ben Liao, Guangyong Chen, and Shengyu Zhang. Understanding and utilizing deep neural networks trained with noisy labels. In *International Conference on Machine Learning(ICML)*, 2019.
- [Feng *et al.*, 2020] Di Feng, Christian Haase-Schütz, Lars Rosenbaum, Heinz Hertlein, Claudius Glaeser, Fabian Timm, Werner Wiesbeck, and Klaus Dietmayer. Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [Ghosh *et al.*, 2017] 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(AAAI)*, 2017.
- [Han *et al.*, 2018a] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *Advances in neural information processing systems(NeurIPS)*, 2018.
- [Han *et al.*, 2018b] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *Advances in neural information processing systems*, 31, 2018.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [Jiang *et al.*, 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(ICML)*, 2018.
- [Jiang *et al.*, 2020] Lu Jiang, Di Huang, Mason Liu, and Weilong Yang. Beyond synthetic noise: Deep learning on controlled noisy labels. In *International Conference on Machine Learning(ICML)*, 2020.
- [Kim *et al.*, 2021] Taehyeon Kim, Jongwoo Ko, JinHwan Choi, Se-Young Yun, et al. Fine samples for learning with noisy labels. *Advances in Neural Information Processing Systems*, 34, 2021.
- [Krizhevsky *et al.*, 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Technical report, University of Toronto*, 2009.
- [Li *et al.*, 2017] 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.
- [Li *et al.*, 2020] Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. *arXiv preprint arXiv:2002.07394*, 2020.
- [Li *et al.*, 2021] Shuailin Li, Zhitong Gao, and Xuming He. Superpixel-guided iterative learning from noisy labels for medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention(MICCAI)*, 2021.
- [Liu *et al.*, 2020] Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. *Advances in neural information processing systems*, 33:20331–20342, 2020.
- [Ma *et al.*, 2020] Xingjun Ma, Hanxun Huang, Yisen Wang, Simone Romano, Sarah Erfani, and James Bailey. Normalized loss functions for deep learning with noisy labels. In *International Conference on Machine Learning(ICML)*, 2020.
- [Mirzasoleiman *et al.*, 2020] Baharan Mirzasoleiman, Kaidi Cao, and Jure Leskovec. Coresets for robust training of neural networks against noisy labels. *arXiv preprint arXiv:2011.07451*, 2020.
- [Nguyen *et al.*, 2019] Duc Tam Nguyen, Chaithanya Kumar Mummadi, Thi Phuong Nhung Ngo, Thi Hoai Phuong Nguyen, Laura Beggel, and Thomas Brox. Self: Learning to filter noisy labels with self-ensembling. *arXiv preprint arXiv:1910.01842*, 2019.
- [Ortego *et al.*, 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*, pages 6606–6615, 2021.

- [Patrini *et al.*, 2017] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1944–1952, 2017.
- [Reed *et al.*, 2014] 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.
- [Shen and Sanghavi, 2019] Yanyao Shen and Sujay Sanghavi. Learning with bad training data via iterative trimmed loss minimization. In *International Conference on Machine Learning (ICML)*, 2019.
- [Song *et al.*, 2019] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. Selfie: Refurbishing unclean samples for robust deep learning. In *International Conference on Machine Learning (ICML)*, 2019.
- [Song *et al.*, 2022] Hwanjun Song, Minseok Kim, Dongmin Park, Yooju Shin, and Jae-Gil Lee. Learning from noisy labels with deep neural networks: A survey. *IEEE Transactions on Neural Networks and Learning Systems (TNNLS)*, 2022.
- [Szegedy *et al.*, 2017] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-first AAAI conference on artificial intelligence*, 2017.
- [Toneva *et al.*, 2018] Mariya Toneva, Alessandro Sordani, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J Gordon. An empirical study of example forgetting during deep neural network learning. In *International Conference on Learning Representations (ICLR)*, 2018.
- [Wei *et al.*, 2020] Hongxin Wei, Lei Feng, Xiangyu Chen, and Bo An. Combating noisy labels by agreement: A joint training method with co-regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13726–13735, 2020.
- [Wu *et al.*, 2020] Pengxiang Wu, Songzhu Zheng, Mayank Goswami, Dimitris Metaxas, and Chao Chen. A topological filter for learning with label noise. *Advances in neural information processing systems*, 33:21382–21393, 2020.
- [Yan *et al.*, 2014] Yan Yan, Rómer Rosales, Glenn Fung, Ramanathan Subramanian, and Jennifer Dy. Learning from multiple annotators with varying expertise. *Machine learning*, 2014.
- [Yu *et al.*, 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*, pages 7164–7173. PMLR, 2019.
- [Zhang and Sabuncu, 2018] Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems (NeurIPS)*, 2018.
- [Zhang *et al.*, 2017a] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- [Zhang *et al.*, 2017b] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- [Zhou *et al.*, 2020] Tianyi Zhou, Shengjie Wang, and Jeff Bilmes. Robust curriculum learning: from clean label detection to noisy label self-correction. In *International Conference on Learning Representations*, 2020.
- [Zhu *et al.*, 2021] Jianing Zhu, Jingfeng Zhang, Bo Han, Tongliang Liu, Gang Niu, Hongxia Yang, Mohan Kankanhalli, and Masashi Sugiyama. Understanding the interaction of adversarial training with noisy labels. *arXiv preprint arXiv:2102.03482*, 2021.