

# MixNN: Combating Noisy Labels in Deep Learning by Mixing with Nearest Neighbours

Yangdi Lu, Wenbo He

**Abstract**—Noisy labels are ubiquitous in real-world datasets, especially in the ones derived from crowdsourcing and web searching. It is challenging to train deep neural networks with noisy datasets since the networks have been shown to overfit the noisy labels during training, resulting in poor performance. During a “early learning” phase, deep neural networks were observed to fit the clean samples before memorizing the mislabeled samples. In this paper, we dig deeper into the representation distributions in the early learning phase and discover that, regardless of their noisy labels, representations of samples from the same classes still congregate together. Inspired by these findings, we propose an algorithm to mitigate the impact of noisy labels. Rather than using existing approaches to identify and remove the mislabeled samples, we propose a weighted mixing strategy to create new synthetic samples by combining original samples with their nearest neighbours, wherein the weights are calculated using a mixture model learning from the per-sample loss distribution. To enhance the performance in the presence of extreme label noise, we propose to estimate the soft targets by gradually correcting the noisy labels. Our intriguing analyses demonstrate that the estimated soft targets yield a more accurate approximation to ground truth labels and the proposed method produces a superior quality of learned representations with more separated and clearly bounded clusters. The extensive experiments in two benchmarks and two challenging real-world datasets demonstrate that our approach outperforms the existing state-of-the-art methods.

**Index Terms**—Label Noise, Image Classification, Deep Learning

## I. INTRODUCTION

Deep neural networks (DNNs) have achieved remarkable performance in a variety of applications (e.g. image classification [1] and object detection [2]). Despite the use of novel network architectures and efficient optimization algorithms, large-scale training data is always required for these supervised tasks. However, obtaining such high-quality training data with human-annotated labels is extremely expensive and time-consuming in practice. Some non-expert sources, such as Amazon’s Mechanical Turk <sup>1</sup> and online searching engines, have been widely used to lower the high labeling cost. However, due to the limited knowledge and inadvertent mistakes, crowdsourced annotators cannot annotate specific tasks with 100% accuracy, resulting in introducing *noisy labels*. Even the most celebrated and highly-curated datasets, such as ImageNet [3], are famously containing noisy labels.

Unlike traditional supervised learning, which assumes that all label information is correct, we consider the training data contains a certain percentage of samples with incorrect labels. Training DNNs on such unreliable datasets is known to be

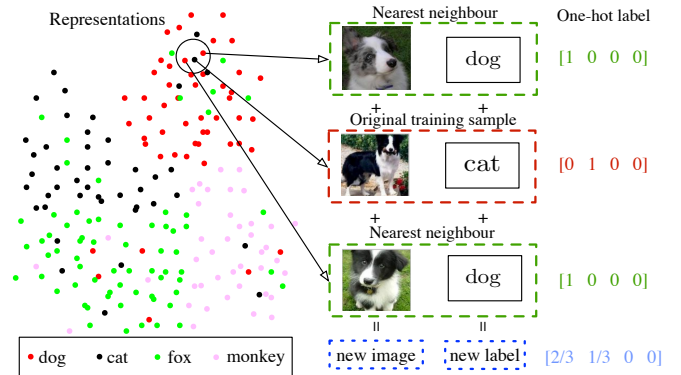


Fig. 1: A synthetic training sample is created by mixing itself with its nearest neighbours based on the learned representations.

highly affected as the significant number of model parameters render DNNs even overfit to noisy labels [4]. Zhang et al. [5] have empirically demonstrated that DNNs can easily fit an entire training dataset with any percentage of corrupted labels, and result in poor generalization capacity on a clean test set. Zhu et al. [6] have observed that the performance drop caused by label noise is more substantial than by other noises, such as feature noise. Therefore, it is crucial to develop learning algorithms that achieve superior generalization capability in the presence of noisy labels.

Given a training set consisting of clean samples and mislabeled samples, a common approach to mitigate the negative impact of noisy labels is to detect the mislabeled samples and eliminate them in the first stage, then train a new classifier with the remaining clean samples in the second stage [7]–[9]. However, the first stage’s filter mechanism for distinguishing the mislabeled samples from the others is critical to the second stage’s classification performance. Assume the filtering mechanism only removes a few mislabeled samples, the unfiltered mislabeled samples still affect the (supervised) loss and deteriorate the classification performance. On the other hand, if too many samples, including clean samples, are eliminated, the remaining data may not be rich enough to generalize to held-out data effectively in the second stage. Therefore, is it possible to robustly train DNNs on noisy data without discarding the informative training samples?

In this paper, we propose an algorithm for combating the negative influence of noisy labels. Our main idea is to generate new synthetic samples that effectively hide the information from the noisy labels, allowing robustly training the DNNs.

<sup>1</sup><https://www.mturk.com/>

Specifically, for each training sample, we search its  $K$  nearest neighbours based on the learned representations. Then we linearly combine the samples with their nearest neighbours in terms of their images and labels to create synthetic samples. By training DNNs with the synthetic samples, our approach mitigates the noisy supervision caused by mislabeled samples and enhances the task performance. Take the case in Fig. 1 as an example – Suppose we simply average the three images and their labels, then the new image is a mixed dog image and the new label is a smoothed label with the probability of  $2/3$  for dog class and  $1/3$  for cat class. Intuitively, the new image appears to have a similar effect to the original image when cutout augmentation is used [10], while the new label is similar to applying label smoothing [11] on the original one-hot label. It is clear to observe that training DNNs with the synthetic samples is more reasonable than train them with the original ones (e.g. a dog image with a wrong label cat).

Instead of simply averaging the samples with its nearest neighbours, we assign dynamic weights to these selected samples and convexly combine them to generate the synthetic samples. Ideally, the weights for clean samples should be large to preserve the correct information while small for mislabeled samples to suppress the wrong supervision. Previous work [12] has observed that DNNs learn the clean pattern before memorizing the complex noisy pattern during training. Specifically, DNNs learn from clean samples at ease and receive inconsistent error supervision from the wrong samples before over-fitting to the entire dataset. Therefore, the predictions and given labels are likely to be consistent on clean samples and inconsistent on wrong samples, resulting in the separation of their loss values in the early learning phase. Based on this observation, we propose to estimate the weights by fitting a two-component Gaussian Mixture Model [13] to the per-sample loss distribution and calculate the posterior probability of the loss value to measure whether a sample is clean or not, allowing the mixing strategy to be dynamic.

To enhance the robustness of our approach, we propose an exponential moving average strategy to estimate targets distribution based on model predictions and given noisy labels. The hard noisy labels are substituted with soft targets, which has been proven to effectively improve the generalization in knowledge distillation [14]. In our scenario, the label quality is improved as the wrong labels can be gradually corrected. Subsequently, the performance of dynamic weight estimation becomes more accurate, providing a more stable supervisory signal in the training procedure.

In summary, our learning framework leverages the mixed samples to prevent the DNNs from overfitting to noisy labels and improve generalization performance. Our main contributions are summarized as follows:

- We provide insights into the memorization procedure and representation distributions of learning with noisy labels. Based on these findings, we propose to generate synthetic training samples to robustly train DNNs, by mixing the original samples with their nearest neighbours.
- We propose to use the dynamic weights for supporting the creation of mixed samples. The weights are proportional to the clean probability of samples, thereby maintaining

correct information while eliminating the wrong information in the mixed samples.

- We propose an exponential moving average strategy to gradually estimate the soft targets, yields enhancing the performance in the presence of extreme label noise.
- We demonstrate that the proposed method outperforms the state-of-the-art methods on two standard benchmarks with simulated label noise and two real-world noisy datasets. We also provide ablation study and empirical analyses (e.g. gradient analysis and feature representations) to verify the effectiveness of different components.

Besides, our method does not require any prior knowledge of the type or the amount of label noise. It does not require any tuning of hyperparameters based on prior knowledge, making our method applicable to real life.

The remainder of this paper is organized as follows: Section II introduces the related work on classification with noisy labels. Section III explores the early learning phenomenon and explains the failure of using cross-entropy loss when learning with noisy labels. Section IV describes each part of our method with detailed explanations. In Section V, we conduct a case study towards a better understanding of our approach and discuss the possible limitations. Section VI details the experiments including a comparison with other state-of-art methods and several analyses. Finally, we conclude the paper in Section VII.

## II. RELATED WORK

Different approaches have been proposed to combating noisy labels in classification task, and they can be classified into the following categories:

**Robust loss functions.** These studies [15]–[22] focus on developing noise-tolerant loss functions. For example, Ghosh et al. [18] have proven mean absolute error (MAE) is a noise-tolerant loss function, while the performance degrades when MAE is adopted in deep neural networks [19]. Wang et al. [20] boost the cross-entropy (CE) loss symmetrically by adding a reverse cross-entropy loss term to avoid overfitting to noisy labels. GCE [19] applies a Box-Cox transformation to probabilities which behaves like a generalized mixture of MAE and CE. This category of method relies on a specific noise model, such as symmetric or asymmetric noise model [23]. Therefore, it only has a minor improvement on the performance when training DNNs with real-world label noise.

**Loss correction and Label correction.** These approaches either iteratively relabel the noisy labels with their own predictions [24]–[26] or estimate the noise transition matrix [23]. For example, Patrini et al. [23] estimate the noise transition matrix and equally treat all samples to correct the loss. Joint-optim [24] iteratively updates the labels with soft or hard pseudo-labels. PENCIL [26] refines the relabel procedure without using prior information about noisy labels. Reed et al. [27] propose a bootstrapping method which corrects the labels by predictions. D2L [28] improves the bootstrapping method by exploiting the dimensionality of feature subspaces.

**Sample selection by Curriculum Learning** These methods [29]–[35] effectively train DNNs by selecting samples through

a meaningful order (i.e. from easy samples to the hard ones). MentorNet [29] pre-trains a mentor network for selecting clean samples to guide the training of the student network. Co-teaching [30] and Co-teaching+ [33] symmetrically train two networks by selecting small-loss instances in a mini-batch to teach the other. Co-matching [35] uses an unsupervised matching loss to boost the performance under extreme label noise through an ensemble manner.

**Semi-supervised learning and meta-learning.** These methods either apply semi-supervised learning techniques after explicitly differentiating noisy samples from training data [8], [36]–[38] or use meta-learning [39], [40]. For example, Li et al. [37] divide the training data into clean and noisy ones, then train two networks with a semi-supervised algorithm MixMatch [41]. SELF [38] progressively filters out mislabeled samples with a semi-supervised approach.

**Other approaches:** Wang et al. [25] apply a Siamese network with a contrastive loss to iteratively pull noisy samples away from the clean ones. Vahdat et al. [42] model the relationship between noisy and clean labels through a Conditional Random Field. Huang et al. [9] use a heuristic way to adjust the learning rate to prevent DNNs from overfitting to label noise. Lu et al. [43] propose a confidence adaptive regularization to prevent DNNs from memorizing noisy labels. Han et al. [44] use the clustering algorithm to find multiple prototypes for correcting noisy labels.

In contrast to the aforementioned literature, our method trains DNNs on noisy labels without: 1) consulting any clean subset; 2) eliminating training samples; 3) applying augmentation techniques from semi-supervised learning; 4) using any prior information. Specifically, we train DNNs with a new training set created by combining the original samples with their nearest neighbours, which prevents DNNs from overfitting to noisy labels during training.

### III. WEAKNESS OF CROSS ENTROPY LOSS

Our goal is to develop an algorithm for training DNNs classifier that performs well on the clean test set despite the presence of noisy labels in the training set. Before introducing our method, we initially describe the preliminary of classification with noisy labels. Then we provide some insights to the memorization procedure of DNNs and analyze the gradient coefficient to explain the weakness of using cross-entropy loss when learning with noisy labels.

#### A. Preliminary

Consider the  $C$ -class classification problem, we have a noisy training set  $\hat{D} = \{(\mathbf{x}_i, \hat{\mathbf{y}}_i)\}_{i=1}^N$ , where  $\mathbf{x}_i$  is an input and  $\hat{\mathbf{y}}_i \in \{0, 1\}^C$  is the one-hot vector corresponding to  $\mathbf{x}_i$ . Note that the ground truth label  $\mathbf{y}_i$  is unobservable, and the observable noisy label  $\hat{\mathbf{y}}_i$  is of certain probability to be incorrect. The classification model maps each input  $\mathbf{x}_i$  to a  $C$ -dimensional logits using a deep neural network model  $\mathcal{N}_\Theta$  and then feeds the logits into a softmax function to produce  $\mathbf{p}_i$  of the conditional probability of each class.

$$\mathbf{p}_i = \text{softmax}(\mathcal{N}_\Theta(\mathbf{x}_i)) = \frac{e^{\mathcal{N}_\Theta(\mathbf{x}_i)}}{\sum_{c=1}^C e^{(\mathcal{N}_\Theta(\mathbf{x}_i))_c}}. \quad (1)$$




Fig. 2: We train ResNet34 [45] on CIFAR-10 dataset with 60% symmetric label noise using cross-entropy loss. Left: The train and test accuracy vs. the number of training epochs. Right: The gradient coefficient  $\mathbf{p}_i - \hat{\mathbf{y}}_i$  of clean and mislabeled samples vs. the number of training epochs.

$\Theta$  denotes the parameters of the DNNs and  $(\mathcal{N}_\Theta(\mathbf{x}_i))_c$  denotes the  $c$ -th entry of logits  $\mathcal{N}_\Theta(\mathbf{x}_i)$ . Traditionally, the model  $\mathcal{N}_\Theta$  is trained via the cross-entropy (CE) loss to measure how well the model fits the training samples.

$$\mathcal{L}_{ce}(\hat{D}, \Theta) = -\frac{1}{N} \sum_{i=1}^N \{\ell_{ce}\}^i = -\frac{1}{N} \sum_{i=1}^N \hat{\mathbf{y}}_i^T \log(\mathbf{p}_i). \quad (2)$$

#### B. Memorization Procedure of CE under Noisy Labels

When trained with noisy labels, the overparameterized DNNs have been observed to first fit the training data with clean labels during an *early learning* stage, before eventually memorizing the examples with wrong labels [12]. Reflected on training and test accuracy in Figure 2, the model achieves maximum test accuracy before achieving the highest training accuracy. During training, the model starts by learning to predict the true labels for correctly labeled training samples. Thus it can predict correct labels for clean test data. However, with the increasing number of training epochs, the model begins making incorrect predictions as it memorizes the mislabeled samples. In addition to the empirical observations, recent study [46] has theoretically proven such an early learning phenomenon also occurs in a simple linear model.

#### C. Gradient Analysis of CE

To further explain why the model memorizes noisy labels eventually, we drive the gradient of cross entropy loss with respect to  $\Theta$  as follows:

$$\nabla \mathcal{L}_{ce}(\hat{D}, \Theta) = -\frac{1}{N} \sum_{i=1}^N \nabla \mathcal{N}_\Theta(\mathbf{x}_i) (\mathbf{p}_i - \hat{\mathbf{y}}_i), \quad (3)$$

where  $\nabla \mathcal{N}_\Theta(\mathbf{x}_i)$  is the Jacobian matrix of the neural network logits for the  $i$ -th input with respect to  $\Theta$ . Assume in clean training data scenario,  $\mathbf{p}_i - \mathbf{y}_i$  of true class entry will always be negative and the rest entries are positive. Therefore, performing stochastic gradient descent increases the probability of true class and reduces the residual probabilities at other classes, which ensures the learning to continue on true class. However, in the noisy label scenario, if  $c$  is the true class, but  $c$ -th entry of noisy label  $(\hat{\mathbf{y}}_i)_c = 0$ , then the contribution of

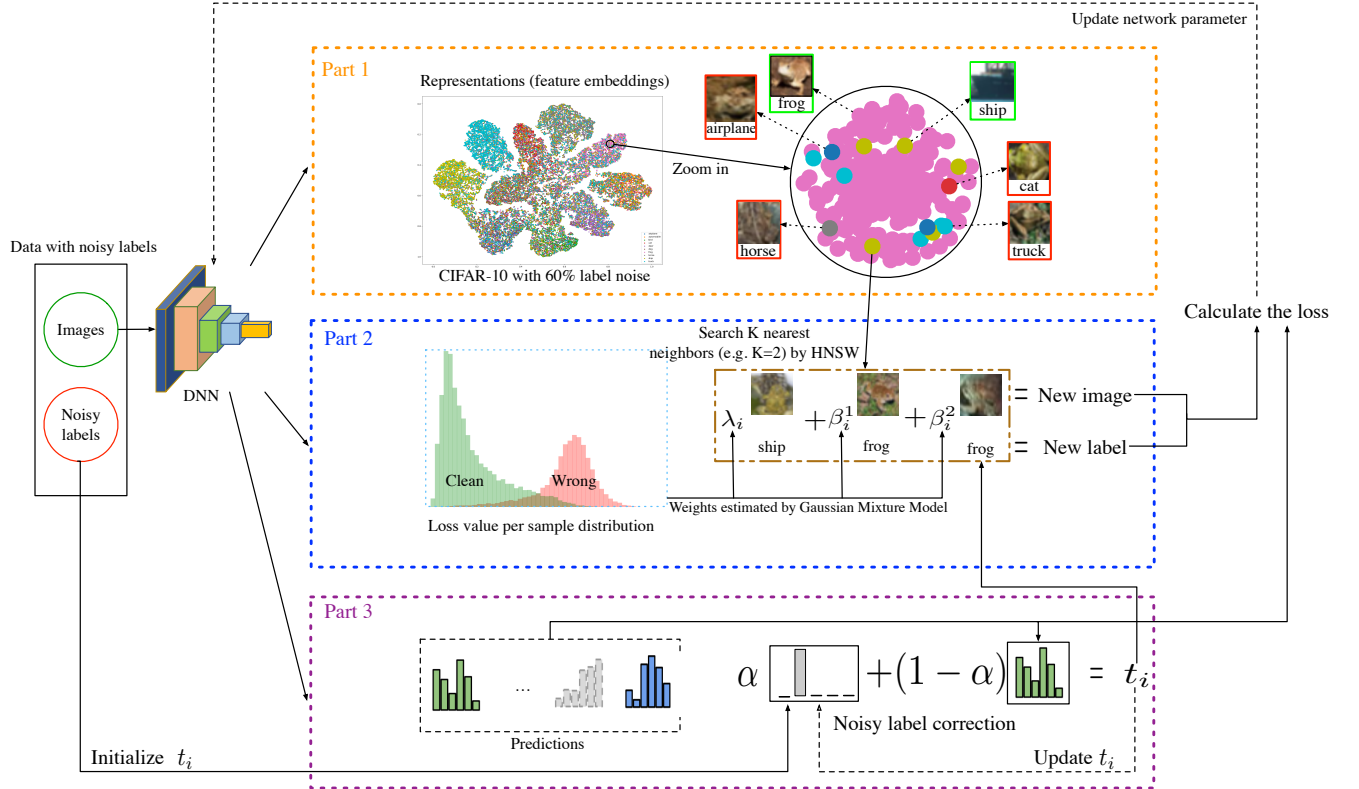


Fig. 3: The proposed method MixNN consists of three parts. Part 1: Based on the feature representations from the penultimate layer, we calculate each training sample’s approximate  $K$ -nearest neighbours by using Hierarchical Navigable Small World (HNSW) graph. Part 2: We combine the original sample with its  $K$ -nearest neighbours by using the dynamic weights estimated from a Gaussian Mixture Model that learned on per-sample loss distribution. Part 3: We estimate the soft targets that gradually correct the noisy labels through an exponential moving average strategy.

the  $i$ -th sample to  $\nabla \mathcal{L}_{ce}(\Theta)$  is reversed (i.e.  $(\mathbf{p}_i - \hat{\mathbf{y}}_i)_c$  should be negative but get positive instead). In the meanwhile, the entry corresponding to the impostor class  $c'$ , is also reversed because  $(\hat{\mathbf{y}}_i)_{c'} = 1$ . Thus, for clean samples, the cross-entropy term  $\mathbf{p}_i - \hat{\mathbf{y}}_i$  tends to vanish (i.e. closer to zero) after the early learning stage because  $\mathbf{p}_i$  is close to  $\mathbf{y}_i$ . For mislabeled samples, the cross-entropy term  $(\mathbf{p}_i - \hat{\mathbf{y}}_i)_c$  is positive, allowing them to dominate the gradient. The right plot in Figure 2 shows the change of gradient coefficient  $(\mathbf{p}_i - \hat{\mathbf{y}}_i)_c$  in training. The gradients of clean samples dominate at the beginning, and then are gradually suppressed by the gradients of mislabeled samples. Therefore, performing stochastic gradient descent eventually results in the memorization of whole training data including mislabeled samples.

#### IV. METHODOLOGY

We name our framework **MixNN** as it **Mixes** each training sample with its **Nearest Neighbours** to train the DNNs. A diagram of our framework is shown in Fig. 3. Generally, MixNN consists of three parts (boxed with different colors). In this section, we first explore the representation distributions in the early learning stage. Then we describe the details in MixNN, including mixing with nearest neighbours,  $K$ -approximate nearest neighbour search, weight estimation, and

noisy labels correction. We also provide the pseudo-code of MixNN for reproducing the experiments conveniently.

##### A. Representation Distributions

In the noisy label scenario, the training data consist of clean samples and mislabeled samples. The goal of MixNN is to prevent the model from memorizing mislabeled samples while continually learning from clean samples. To study whether the learned representations are corrupted due to label noise, we plot the t-SNE graph [47] of learned representations (i.e. the embeddings from penultimate layer) in the early learning stage in Fig. 3 Part 1. As we can see, the learned representations of majority clean samples still congregate in their true classes, while the representations of mislabeled samples disperse in all classes. For clear illustration, we zoom in a random region in frog (pink) class and display the samples in the right as an example. In this region, we observe that most of the frog images have learned correct representations, only a few samples from other classes have learned ‘corrupted’ representations (e.g. the ship image shouldn’t have a frog representation). Besides, most of the mislabeled samples in frog class are also frog images though attached wrong labels from other classes. In our experiments, we find that 88.96% and 82.8% mislabeled samples’ representations still congregate in their true class for CIFAR-10 with 40% and 60% label noise respectively.

This observation motivates us to consider whether we can refer the correct information from the mislabeled samples’ nearest neighbours. To achieve our goal, we generate a new training set where each sample is mixed with its  $K$  nearest neighbours. Ideally, the correct image features are most likely to be preserved while the negative influence of noisy labels can be mitigated, yields achieving robustness to label noise.

### B. Mixing with Nearest Neighbors

Our main idea is to use the correct knowledge from nearest neighbours to mitigate the detrimental impact of noisy labels. For each training sample  $\mathbf{x}_i$  in a mini-batch, we generate a synthetic training sample  $\tilde{\mathbf{x}}_i$  by convex linearly combining it with its  $K$  nearest neighbours. We denote  $K$  nearest neighbours of a training sample  $(\mathbf{x}_i, \hat{\mathbf{y}}_i)$  as  $\Phi^K(\mathbf{x}_i) = \{(\tilde{\mathbf{x}}_i^k, \tilde{\mathbf{y}}_i^k)\}_{k=1}^K = \{(\tilde{\mathbf{x}}_i^1, \tilde{\mathbf{y}}_i^1), (\tilde{\mathbf{x}}_i^2, \tilde{\mathbf{y}}_i^2), \dots, (\tilde{\mathbf{x}}_i^K, \tilde{\mathbf{y}}_i^K)\}$ . Hence, we define the image of a synthetic training sample as

$$\tilde{\mathbf{x}}_i = \lambda_i \mathbf{x}_i + \sum_{k=1}^K \beta_i^k \tilde{\mathbf{x}}_i^k, \quad \sum_{k=1}^K \beta_i^k = 1 - \lambda_i, \quad (4)$$

where  $\lambda_i$  is a dynamic scalar value denoting the weight of original training sample.  $\beta_i^k$  denotes the weight of  $k$ -th nearest neighbour. We have  $\lambda_i + \sum_{k=1}^K \beta_i^k = 1$  to ensure the synthetic training sample still follows the same distribution of original sample after normalization. Similarly, we calculate the new label  $\tilde{\mathbf{y}}_i$  of the new sample by

$$\tilde{\mathbf{y}}_i = \lambda_i \hat{\mathbf{y}}_i + \sum_{k=1}^K \beta_i^k \tilde{\mathbf{y}}_i^k, \quad \sum_{k=1}^K \beta_i^k = 1 - \lambda_i. \quad (5)$$

We then train our model with the synthetic training set  $\tilde{D} = \{(\tilde{\mathbf{x}}_i, \tilde{\mathbf{y}}_i)\}_{i=1}^N$ . Based on  $\tilde{D}$ , we use the cross-entropy loss as the measure of how well the model fits the  $\tilde{D}$ . We denote

$$\tilde{\mathbf{p}}_i = \text{softmax}(\mathcal{N}_\Theta(\tilde{\mathbf{x}}_i)) = \frac{e^{\mathcal{N}_\Theta(\tilde{\mathbf{x}}_i)}}{\sum_{c=1}^C e^{(\mathcal{N}_\Theta(\tilde{\mathbf{x}}_i))_c}}. \quad (6)$$

Thus, our new loss becomes

$$\begin{aligned} \mathcal{L}(\tilde{D}, \Theta) &= -\frac{1}{N} \sum_{i=1}^N \tilde{\mathbf{y}}_i^T \log(\tilde{\mathbf{p}}_i) \\ &= -\frac{1}{N} \sum_{i=1}^N (\lambda_i \hat{\mathbf{y}}_i + \sum_{k=1}^K \beta_i^k \tilde{\mathbf{y}}_i^k)^T \log(\tilde{\mathbf{p}}_i) \\ &= -\frac{1}{N} \sum_{i=1}^N \lambda_i \hat{\mathbf{y}}_i^T \log(\tilde{\mathbf{p}}_i) \\ &\quad - \frac{1}{N} \sum_{i=1}^N \left( \sum_{k=1}^K \beta_i^k \tilde{\mathbf{y}}_i^k \right)^T \log(\tilde{\mathbf{p}}_i). \quad (7) \end{aligned}$$

Compared to CE loss in Eq. (2), the new loss is more resistant to noisy labels. In Eq. (7), the first term is similar to the CE loss, except it is weighted by  $\lambda_i$ . The target in the second term is generated by disparate distributions from  $K$  nearest neighbours. It makes the prediction  $\tilde{\mathbf{p}}_i$  of mixed sample hard to overfit the noisy label  $\hat{\mathbf{y}}_i$ , resulting in mitigating the impact of noisy labels on training DNNs.

### C. Approximate Nearest Neighbour Search

In our method, we search the  $K$ -Nearest Neighbors (KNN) for each training sample based on the learned representation. Assume the learned representation of a training sample is a query vector. A naive approach to performing exact KNN search is to directly compute the distances (e.g. Euclidean distance and Cosine distance) between the query and every element in the training set. Hence, the complexity of the naive approach is  $O(dN)$ , where  $N$  is the size of training set and  $d$  is the dimension of representation vector.

Previous study [48] has demonstrated that exact KNN search solutions may offer a substantial search speedup only in the case of relatively low dimensional data (e.g.  $d < 20$ ) due to ‘‘curse of dimensionality’’. For instance, the complexity of KNN search in KD-tree [49] is  $O(2^d \log(N))$  which is exponential to the dimension of representations. In our case  $d$  can be large. For example, the dimension of representation in penultimate layer for ResNet34 [45] is 512. Therefore, it is inefficient to directly calculate the exact KNN.

To overcome this problem, a concept of Approximate Nearest Neighbours Search (ANNS) [48] was proposed, which relaxes the condition of the exact search by allowing a small number of errors. The quality of an inexact search is defined as the ratio between the number of found true nearest neighbours and  $K$ . In this paper, we adopt the Hierarchical Navigable Small World (HNSW) graph [50] as our search index. It is a fully graph based incremental ANNS structure that can offer a superior logarithmic complexity scaling. The search index in HNSW is a multi-layered structure where each layer is a proximity graph. Each node in the graph corresponds to one of the representations. A nearest neighbour search in HNSW adopts a ‘‘zooming-in’’ style. It starts at an entry point node in the uppermost layer and recursively performs a greedy graph traversal in each layer until it reaches a local minimum in the bottommost one. The maximum number of connections per element in all layers can be made a constant, thus allowing a logarithmic complexity scaling of routing in a navigable small world graph. In this paper, we use Euclidean distance as the measure of similarity and the overall search complexity scaling is  $O(\log(N))$ . When  $K = 4$ , the original exact KNN search costs 5.25 millisecond per sample in our experiment. When using HNSW, the search time is reduced to 0.034 millisecond per sample, which is quite efficient.

### D. Weight Estimation

In mixing functions Eq. (4) and Eq. (5), the weight  $\lambda_i$  indicates how confidently we can trust the original sample, whereas  $\beta_i^1, \dots, \beta_i^K$  indicates how much knowledge is referred from these nearest neighbours. Ideally, we want to preserve correct information from the clean samples while dampening the wrong information from the mislabeled samples. In other words, the weights should be able to indicate the ‘probability’ of a training sample being correctly labeled or not.

Due to the early learning phenomenon, the samples with small-loss values are more likely to be correctly labeled [30]. Therefore, existing sample selection methods [30], [33] select the clean samples according to the magnitude of loss values. In

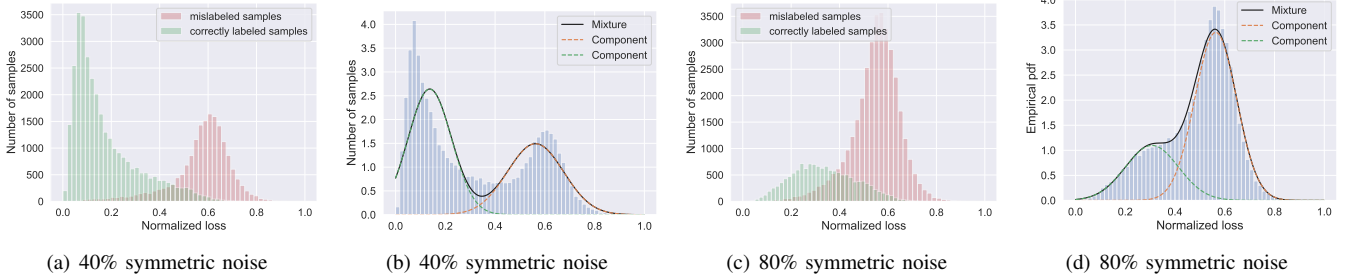


Fig. 4: Train on CIFAR-10 with 40% and 80% label noise after 10 epochs with cross-entropy loss. Plots (a) and (c): The ground truth normalized loss distribution. Plots (b) and (d): The pdf of mixture model and two components after fitting a two-component GMM to per-sample loss distribution.

this paper, we investigate the per-sample loss distribution and find that there is a separation between the loss distribution of correctly labeled samples and loss distribution of mislabeled samples. As shown in Fig. 4 (a) and (c), the normalized loss values of the clean samples are in expectation smaller than the mislabeled ones. Even in the case of extreme label noise (e.g. 80% labels are incorrect), their loss distributions can be differentiated. To estimate the probability of a sample being clean, we introduce a two-component Gaussian Mixture Model (GMM) [13] to fit the per-sample loss distribution as shown in Fig. 4 (b) and (d). The probability density function (pdf) of GMM with  $M$  components on the per sample loss value  $\ell$  can be defined as

$$P(\ell) = \sum_{m=1}^M \pi_m \mathcal{G}(\ell | \mu_m, \sigma_m^2), \quad \sum_{m=1}^M \pi_m = 1, \quad (8)$$

where  $\pi_m$  is the coefficient for the linear convex combination of each individual pdf  $\mathcal{G}(\ell | \mu_m, \sigma_m^2)$ . In our case, we use an Expectation-Maximization (EM) algorithm to estimate the  $\pi_m$ ,  $\mu_m$  and  $\sigma_m^2$ . Therefore, we obtain the probability of a sample being clean or mislabeled through the posterior probability:

$$P(m | \ell) = \frac{P(m)P(\ell | m)}{P(\ell)} = \frac{\pi_m \mathcal{G}(\ell | \mu_m, \sigma_m^2)}{\sum_{m=1}^M \pi_m \mathcal{G}(\ell | \mu_m, \sigma_m^2)} \quad (9)$$

where  $m = 0(1)$  indicate correct (wrong) labels. Note that we always calculate the cross-entropy loss to estimate the clean probability for all samples after every epoch. But we use our loss defined in Eq. (7) for training the model which contains multiple loss terms to deal with label noise.

While mislabeled samples benefit from combining with clean ones, clean samples are contaminated by mislabeled ones, whose training objective is incorrectly modified. The goal of mixing strategy in Eq. (4) and Eq. (5) is to use the dynamic weights to reduce the contribution of mislabeled samples when they are combined with correctly labeled ones. We denote the per sample loss value of  $\mathbf{x}_i$  as  $\ell(\mathbf{x}_i)$ . Thus the dynamic weights are calculated by

$$\lambda_i = \frac{P(m=0 | \ell(\mathbf{x}_i))}{P(m=0 | \ell(\mathbf{x}_i)) + \sum_{k=1}^K P(m=0 | \ell(\tilde{\mathbf{x}}_i^k))}, \quad (10)$$

$$\beta_i^k = \frac{P(m=0 | \ell(\tilde{\mathbf{x}}_i^k))}{P(m=0 | \ell(\mathbf{x}_i)) + \sum_{k=1}^K P(m=0 | \ell(\tilde{\mathbf{x}}_i^k))}. \quad (11)$$

We then use the above weights to guide the generation of synthetic training sample  $(\tilde{\mathbf{x}}_i, \tilde{\mathbf{y}}_i)$ . Consider  $K=1$ , there are totally four mixing cases: clean-clean, clean-wrong, wrong-clean, and wrong-wrong. By using dynamic weights, it largely avoids generating the confusing input to the network in clean-wrong and wrong-clean cases, while retaining the strengths for clean-clean and wrong-wrong combinations. More discussion on these four cases is in Section V.

### E. Noisy Labels Correction

Despite that the synthetic training samples set  $\tilde{D} = \{(\tilde{\mathbf{x}}_i, \tilde{\mathbf{y}}_i)\}_{i=1}^N$  is better than directly using given noisy training dataset  $D = \{(\mathbf{x}_i, \hat{\mathbf{y}}_i)\}_{i=1}^N$ . Nevertheless, using the noisy labels  $\hat{\mathbf{y}}_i$  in Eq. (5) and Eq. (7) may be less effective as  $\hat{\mathbf{y}}_i$  is likely to be incorrect, especially when the noise rate is extremely high. Therefore, we need a better estimation of ground truth label  $\mathbf{y}_i$ . In Fig. 2, we observe that most predictions in early learning stage are correct. Based on this observation, we propose an exponential moving average strategy to gradually estimate the soft target  $\mathbf{t}_i$  by using the noisy label  $\hat{\mathbf{y}}_i$  and model prediction  $\mathbf{p}_i$ . We update  $\mathbf{t}_i$  in each epoch  $E$  by

$$\mathbf{t}_i = \begin{cases} \hat{\mathbf{y}}_i & \text{if } E < E_s \\ \alpha \mathbf{t}_i + (1 - \alpha) \mathbf{p}_i & \text{if } E \geq E_s \end{cases} \quad (12)$$

where  $E_s$  is the epoch that starts performing label correction and  $0 \leq \alpha < 1$  is the momentum. In this paper, we fix  $E_s = 60$  and  $\alpha = 0.9$  by default. We then replace the noisy label  $\hat{\mathbf{y}}_i$  in Eq. (5) and Eq. (7) with the estimated soft target  $\mathbf{t}_i$ . Consequently, using a better  $\mathbf{t}_i$  facilitate the model to memorize more correctly labeled samples and to generate a better new mixed label  $\tilde{\mathbf{y}}$ . The correction accuracy will be discussed in Section VI-F. Overall, put all parts together, our algorithm is described in Algorithm 1.

## V. DISCUSSION AND LIMITATION

To explain how MixNN works, we discuss the possible cases when mixing with one nearest neighbour ( $K=1$ ). We provide an example to illustrate each case in Fig. 5. The synthetic samples  $(\tilde{\mathbf{x}}_i, \tilde{\mathbf{y}}_i)$  are generated by the following four ways.

**Algorithm 1: MixNN**


---

**Input:** Networks  $\mathcal{N}_\Theta$  with parameters  $\Theta$ , training set  $\hat{D}$ , number of nearest neighbors  $K$ , batch size  $B$ , learning rate  $\eta$ , number of total training epochs  $E_{\max}$ ,  $E_s = 60$  momentum  $\alpha = 0.9$ ;

- 1  $\Theta = \text{Warmup}(\hat{D}, \Theta, \mathcal{L}_{ce})$ ; // warmup with cross-entropy loss.
- 2 **Initialize** the target  $\mathbf{t}_i = \hat{\mathbf{y}}_i$  for all samples in  $\hat{D}$ ;
- 3 **for**  $e = 1, 2, \dots, E_{\max}$  **do**
- 4      $P(m = 0 | \ell_{ce}(\mathbf{x}_i)) = \text{GMM}(\hat{D}, \ell_{ce}, \Theta)$ ;  
      // fit GMM to per-sample loss distribution.
- 5     **Obtain**  $K$  Approximate Nearest Neighbors  
       $\Phi^K(\mathbf{x}_i) = \{(\tilde{\mathbf{x}}_i^k, \tilde{\mathbf{y}}_i^k)\}_{k=1}^K$  for each  $\mathbf{x}_i$  by using HNSW;
- 6     **Shuffle**  $\hat{D}$  into  $\frac{|\hat{D}|}{B}$  mini-batches ;
- 7     **for**  $n = 1, 2, \dots, \frac{|\hat{D}|}{B}$  **do**
- 8         **Fetch**  $n$ -th mini-batch  $\hat{D}_n$  from  $\hat{D}$  ;
- 9         **Obtain**  $\mathbf{p}_i$  for each sample in  $\hat{D}_n$  by Eq. (1);
- 10        **if**  $e \geq E_s$  **then**
- 11            **Update**  
             $\mathbf{t}_i = \alpha \mathbf{t}_i + (1 - \alpha) \mathbf{p}_i$  for all sample in  $\hat{D}_n$ ;
- 12            **Estimate** the weights  $\lambda_i, \beta_i^1, \dots, \beta_i^K$  for each sample in  $\hat{D}_n$  by Eq. (10) and Eq. (11).;
- 13            **Generate**  $\tilde{D}_n = \{(\tilde{\mathbf{x}}_i, \tilde{\mathbf{y}}_i)\}_{i=1}^B$  by Eq. (4), Eq. (5) and  $\Phi^K(\mathbf{x}_i)$ ;
- 14            **Obtain**  $\tilde{\mathbf{p}}_i$  for each sample in  $\tilde{D}_n$  by Eq. (6);
- 15            **Calculate** the loss  
             $\mathcal{L}(\tilde{D}_n, \Theta) = -\frac{1}{B} \sum_{i=1}^B \lambda_i \mathbf{t}_i^T \log(\tilde{\mathbf{p}}_i) - \frac{1}{B} \sum_{i=1}^B (\sum_{k=1}^K \beta_i^k \tilde{\mathbf{y}}_i^k)^T \log(\tilde{\mathbf{p}}_i)$ ;
- 16            **Update**  $\Theta = \Theta - \eta \nabla \mathcal{L}(\tilde{D}_n, \Theta)$  ;
- 17 **Output**  $\Theta$ .

---

**Case 1:** Both  $(\mathbf{x}_i, \hat{\mathbf{y}}_i)$  and its nearest neighbour  $(\tilde{\mathbf{x}}_i^1, \tilde{\mathbf{y}}_i^1)$  are clean samples. Based on the analysis of representation distributions in the early learning phase, both  $\mathbf{x}_i$  and  $\tilde{\mathbf{x}}_i^1$  are most likely to be the images from the same class. In this case, the new mixed input is a convex linear combination of two similar images, i.e.,  $\tilde{\mathbf{x}}_i = \lambda_i \mathbf{x}_i + \beta_i^1 \tilde{\mathbf{x}}_i^1$ . And its corresponding label is  $\tilde{\mathbf{y}}_i = \lambda_i \hat{\mathbf{y}}_i + \beta_i^1 \tilde{\mathbf{y}}_i^1$ . Since  $\hat{\mathbf{y}}_i = \tilde{\mathbf{y}}_i^1$  and  $\lambda_i + \beta_i^1 = 1$ , then the mixed label  $\tilde{\mathbf{y}}_i = \hat{\mathbf{y}}_i = \tilde{\mathbf{y}}_i^1$ . Therefore, the new sample  $(\tilde{\mathbf{x}}_i, \tilde{\mathbf{y}}_i)$  is similar to the sample after applying augmentation strategy [51] which encourages the model to behave linearly in-between training samples, resulting in reducing the number of undesirable oscillations when predicting hard samples. The DNNs trained with such samples will be better calibrated [52]. In other words, the prediction softmax scores are much better indicators of the actual likelihood of predictions, which avoids producing the overconfident wrong predictions and improves the estimation of dynamic weights in Section IV-D.

**Case 2:**  $(\mathbf{x}_i, \hat{\mathbf{y}}_i)$  is clean sample and its nearest neighbour  $(\tilde{\mathbf{x}}_i^1, \tilde{\mathbf{y}}_i^1)$  is mislabeled sample. Similarly, the new mixed input is most likely to be a combination of two images from the same class, while their labels are inconsistent. In Fig. 5 case 2, the

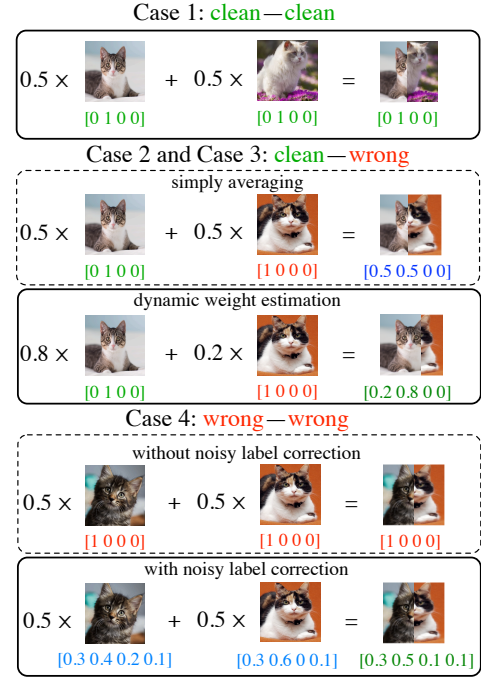


Fig. 5: An example for illustrating the different cases in MixNN. Here, we use the one-hot label vector and each entry indicates four different classes (i.e. dog, cat, fox and monkey).

mixed label becomes an ambiguous target if simply averaging the samples. However, with estimated dynamic weights, the mixed label is determined and correct. Our method to generate the mixed label is similar to label smoothing [11] which scales and translates the original noisy label  $\hat{\mathbf{y}}_i$  to  $(1 - \gamma)\hat{\mathbf{y}}_i + \gamma/C$ , but preserves the label with maximal probability when  $\gamma < 1$ . Different from the label smoothing that uses the fixed uniform distribution to scale the noisy labels, our approach adopts the dynamic weights, where  $\lambda, \beta^1, \dots, \beta^K$  are learned from data, to adaptively adjust the mixed label for better performance.

**Case 3:**  $(\mathbf{x}_i, \hat{\mathbf{y}}_i)$  is mislabeled sample and its nearest neighbour  $(\tilde{\mathbf{x}}_i^1, \tilde{\mathbf{y}}_i^1)$  is clean sample. This case is similar to case 2, so we do not further discuss it.

**Case 4:** Both  $(\mathbf{x}_i, \hat{\mathbf{y}}_i)$  and its nearest neighbour  $(\tilde{\mathbf{x}}_i^1, \tilde{\mathbf{y}}_i^1)$  are mislabeled samples. In this case, the mixed input is combined with two images from the same class, but their labels are both incorrect. Thus the mixed sample is not promised to improve the generalization capacity as shown in Fig. 5. However, the noisy labels are very likely to be corrected after using label correction (more discussions in Section VI-F). Therefore, the mixed label can be a reasonable target compared to the original noisy one.

**Limitations of MixNN** Suppose we have encountered the worst case, i.e., the label correction fails, the new mixed input is a mixture of two images from different classes, and their labels are both incorrect and inconsistent. It is the main limitation of our approach as this case definitely degrades the performance. However, this worst case rarely occurs. In our experiments, we observe that the worst case occurs in only 4.3% of the training data when trained on CIFAR-10 with

60% label noise.

## VI. EXPERIMENTS

In this section, we first test the efficacy of the proposed method on two benchmark datasets with simulated label noise. Then we evaluate the performance of MixNN on the real-world datasets which contain more sophisticated label noise. We also provide ablation study and qualitative results to investigate the effect of different components. We conduct several empirical analyses to gain a better understanding of our approach, including learning stability, gradient analysis, the choice of  $K$  in mixing function, effectiveness of label correction, and feature representations. All experiments are implemented in Pytorch and run on a single NVIDIA A100 GPU.

### A. Effectiveness on Simulated Label Noise

We conduct the experiments with simulated label noise on the following two datasets.

- *CIFAR-10* [53] consists of 60,000 natural colour images, each of size  $32 \times 32$  pixels. These images are classified into 1 of 10 classes, such as dog, cat, automobile, or ship. The training set contains 50,000 images, while the test set contains 10,000 images.
- *CIFAR-100* [53] is similar with CIFAR-10, except it has 100 classes containing 600 images each. The 100 classes are grouped into 20 superclasses. Each image comes with a “fine” label (the class to which it belongs) and a “coarse” label (the superclass to which it belongs).

**Label Noise Simulation** Given CIFAR-10 and CIFAR-100 are initially clean, we follow the way in [23] to corrupt these two benchmarks manually by label transition matrix  $\mathbf{Q}$ , where  $\mathbf{Q}_{ij} = \Pr[\hat{y} = j \mid y = i]$  denotes the probability that noisy label  $\hat{y}$  is flipped from clean label  $y$ . Generally, the matrix  $\mathbf{Q}$  has two representative label noise models. (1) Symmetric noise [17] is generated by uniformly flipping labels in each class to one of the other class labels with probability  $\varepsilon$ . (2) Asymmetric noise [23] is a simulation of fine-grained classification with noisy labels in the real world, where the annotators are more likely to make mistakes only within very similar classes. Fig. 6 shows an example of label transition matrix  $\mathbf{Q}$  for above two label noise models. In this paper, the asymmetric noisy labels are generated by flipping *truck*  $\rightarrow$  *automobile*, *bird*  $\rightarrow$  *airplane*, *deer*  $\rightarrow$  *horse* and *cat*  $\leftrightarrow$  *dog* for CIFAR-10. For CIFAR-100, the noise flips each class into the next, circularly within super-classes. **Data Preprocessing** We apply normalization and regular image augmentations (i.e. random crop and horizontal flip) on training data. The crop size for CIFAR-10 and CIFAR-100 remains 32 after padding with 4 pixels.

**Network and Optimizer** We use ResNet34 [45] for both datasets, and train them using SGD with a momentum of 0.9, a weight decay of 0.001, and a batch size of 128. The networks are trained for 300 epochs. We use the cosine annealing learning rate [54] where the maximum number of epoch for each period is 10, the maximum and minimum learning rate is set to 0.02 and 0.001 respectively. We warm up our networks 10 epochs for CIFAR-10 and 30 epochs for CIFAR-100 with

Symmetric-40%					Asymmetric-40%				
60%	10%	10%	10%	10%	100%	0%	0%	0%	0%
10%	60%	10%	10%	10%	0%	60%	0%	40%	0%
10%	10%	60%	10%	10%	0%	40%	60%	0%	0%
10%	10%	10%	60%	10%	0%	0%	0%	100%	0%
10%	10%	10%	10%	60%	0%	40%	0%	0%	60%

Fig. 6: Example of label transition matrix  $\mathbf{Q}$  (taking 5 classes and noise ratio  $\varepsilon = 0.4$  as an example).

cross-entropy loss. We do not perform early stopping [55] since we don’t assume the presence of clean validation data. All test accuracies are recorded from the last epoch of training. The reason that we train the model 300 epochs is to fully evaluate whether the model will memorize the mislabeled samples, avoiding the interference caused by early stopping [4] (i.e. the model may not start memorizing mislabeled samples when the number of training epochs is small). We set  $K = 1$  since we find that a larger  $K$  worsens the performance. More discussions on choice of  $K$  can be found in Section VI-E.

**Baselines** We compare our method to the following baselines from different categories. (1) CE directly uses the standard cross-entropy loss to train the DNNs on noisy training data. (2) F-correction [23] and Bootstrap [27] belong to loss correction category. (3) GCE [19], SCE [20], NFL+MAE [22], and NCE+RCE [22] belong to robust loss function category. (4) Joint Optim [24], PENCIL [26], RoG+D2L [56], M-correction [57], SEAL [58] and LRT [59] belong to label correction category. (4) Decoupling [31], Co-teaching [30], MentorNet [29] and Iterative-CV [60] belong to sample selection category. (4) O2U-net [9], NLNL [8], DAC [7], Crust [61] and ODD [62] belong to noisy pruning category. (5) SELF [38] belongs to semi-supervised learning category.

**Results on CIFAR-10 and CIFAR-100** Table I shows the classification test accuracies of our approach on CIFAR-10 and CIFAR-100 with different levels of symmetric and asymmetric label noise. As we can see, MixNN achieves the highest accuracy in most cases, especially in challenging ones. For example, on CIFAR-10 with 80% symmetric label noise, MixNN outperforms the best state-of-the-art method (74.84% of DAC) by more than 11%. In the hardest case (i.e. CIFAR-100 with 80% symmetric label noise), we observe that most existing methods achieve relatively low test accuracies and PENCIL even fails to converge. However, MixNN still achieves the best accuracy up to 48.81%. Note that on CIFAR-10/CIFAR-100 with 20% symmetric label noise, NLNL and DAC obtain superior performance and even outperform our approach. However, these two methods are more complex and perform multiple training stages for different purposes. For example, NLNL performs three training stages including a) Division of training data into either clean or noisy data with a DNN model. b) Training initialized DNN with clean data from the first stage and then updating noisy data’s label following the output of DNN trained with clean data. c) Clean data and



TABLE I: Test Accuracy (%) on CIFAR-10 and CIFAR-100 with different ratios of symmetric and asymmetric label noise. We compare with existing methods under the same backbone ResNet34 [45]. The average accuracy and standard deviation of 3 random runs are reported. *symm/asymm* represent symmetric/asymmetric label noise respectively. **Bold** indicates the best results. - indicates the result is not reported.

Model	Dataset Noise type Method/Noise ratio	CIFAR-10					CIFAR-100				
		symm				asymm	symm				asymm
		20%	40%	60%	80%	40%	20%	40%	60%	80%	40%
ResNet34	CE	86.98 ± 0.12	81.88 ± 0.29	74.14 ± 0.56	53.82 ± 1.04	80.11 ± 1.44	58.72 ± 0.26	48.20 ± 0.65	37.41 ± 0.94	18.10 ± 0.82	42.74 ± 0.61
	F-correction [23]	87.99 ± 0.36	83.25 ± 0.38	74.96 ± 0.65	54.64 ± 0.44	83.55 ± 0.58	39.19 ± 2.61	31.05 ± 1.44	19.12 ± 1.95	8.99 ± 0.58	34.44 ± 1.93
	Bootstrap [27]	86.23 ± 0.23	82.23 ± 0.37	75.12 ± 0.56	54.12 ± 1.32	81.21 ± 1.47	58.27 ± 0.21	47.66 ± 0.55	34.68 ± 1.10	21.64 ± 0.97	45.12 ± 0.57
	GCE [19]	89.83 ± 0.20	87.13 ± 0.22	82.54 ± 0.23	64.07 ± 1.38	76.74 ± 0.61	66.81 ± 0.42	61.77 ± 0.24	53.16 ± 0.78	29.16 ± 0.74	47.22 ± 1.15
	SCE [20]	89.83 ± 0.32	87.13 ± 0.26	82.81 ± 0.61	68.12 ± 0.81	82.51 ± 0.45	70.38 ± 0.13	62.27 ± 0.22	54.82 ± 0.57	25.91 ± 0.44	49.32 ± 0.87
	NFL+MAE [22]	-	83.81 ± 0.06	76.36 ± 0.31	45.23 ± 0.52	77.16 ± 0.10	-	58.18 ± 0.08	46.10 ± 0.50	24.78 ± 0.82	43.51 ± 0.42
	NCE+RCE [22]	-	86.02 ± 0.09	79.78 ± 0.50	52.71 ± 1.90	79.59 ± 0.40	-	59.48 ± 0.56	47.12 ± 0.62	25.80 ± 1.12	46.79 ± 0.96
	Joint Optim [24]	92.25	90.79	86.87	69.16	-	58.15	54.81	47.94	17.18	-
	PENCIL [26]	-	-	-	-	<b>91.01</b>	-	69.12 ± 0.62	57.79 ± 3.86	fail	63.61 ± 0.23
	RoG+D2L [56]	-	87.00	78.00	-	-	-	64.90	40.60	-	-
	M-correction [57]	-	92.30	86.10	74.10	-	-	70.10	59.50	39.50	-
	MentorNet [29]	92.00	91.20	74.20	60.00	-	73.50	68.50	61.20	35.50	-
	O2U-net [9]	-	90.30	-	43.40	-	-	69.20	-	39.40	-
	NLNL [8]	<b>94.23</b>	92.43	88.32	-	89.86	71.52	66.39	56.51	-	45.70
	DAC [7]	92.91	90.71	86.30	74.84	-	73.55	66.92	57.17	32.16	-
	SELF [38]	-	91.13	-	63.59	-	-	66.71	-	35.56	-
	MixNN (Ours)	93.91 ± 0.12	<b>92.89 ± 0.02</b>	<b>91.66 ± 0.07</b>	<b>86.08 ± 1.01</b>	<b>90.25 ± 0.76</b>	<b>74.81 ± 0.14</b>	<b>72.97 ± 0.14</b>	<b>67.56 ± 0.17</b>	<b>48.81 ± 0.06</b>	<b>68.18 ± 0.11</b>

label-updated noisy data are both used for training initialized DNN in the final stage. In contrast, our method conducts an end-to-end learning manner which is much simpler than NLNL and also achieves excellent performance. In summary, MixNN shows a consistently strong performance across all datasets with different types and ratios of simulated label noise.

### B. Effectiveness on Real-world Label Noise

We use the following datasets to evaluate the performance of our approach under the real-world noisy labels settings.

- *Clothing1M* [63] contains 1 million images of clothing obtained from online shopping websites with 14 classes, such as T-shirt, sweater, and so on. The labels are generated by using the surrounding texts of the images that are provided by the sellers and thus contain many wrong labels. The overall accuracy of the labels is around 61.54%, with some pairs of classes frequently confused with each other (e.g. knitwear and sweater). This dataset also contains 50k, 14k, and 10k of clean data for training, validation, and testing, respectively. Note that we do not use the 50k clean data in our training process. We report the classification accuracy on the test set when the performance on the validation set is optimal.
- *Webvision* [64] is a large web images dataset that contains more than 2.4 millions of images crawled from the Flickr and Google Images Search. The label noise level of Webvision is estimated at 20%. Following [60], we compare the baseline methods on the first 50 classes of Google image subset. We not only test the trained model on the human-annotated WebVision validation set, but also test the model on ILSVRC12 validation set to evaluate generalization capability.

**Data Preprocessing** We apply normalization and regular data augmentations (i.e. random crop and horizontal flip) on the training sets. Since real-world images are of different sizes, we perform the cropping consistent with the existing work [60]. Specifically,  $224 \times 224$  for Clothing1M (after resizing to  $256 \times 256$ ), and  $227 \times 227$  for Webvision.

**Network and Optimizer** We use the ResNet-50 [45] pre-trained on ImageNet for Clothing1M. We train the model with batch size 64. The optimization is done using SGD with a momentum of 0.9, and weight decay of 0.001. We use the same cosine annealing learning rate as CIFAR-10 except the minimum learning rate is set to 0.0001 and the total epoch is 200. For each epoch, we randomly sample 2000 mini-batches from the training data ensuring that the classes of the noisy labels are balanced. For Webvision, we use InceptionResNetV2 [65] as the backbone architecture. All other optimization details are the same as for CIFAR-10, except for the weight decay (0.0005), the batch size (32) and maximum learning rate (0.01).

**Results on Clothing1M and Webvision** Table II shows the results on Clothing1M dataset. MixNN consistently outperforms other baselines, slightly superior to Joint-Optim. Table III compares MixNN to state-of-the-art methods trained on the Webvision and evaluated on both the WebVision and ImageNet ILSVRC12 validation sets. On WebVision validation set, MixNN produces superior results in terms of top 1 and top 5 accuracies, demonstrating our method is effective on datasets containing real-world label noise. Furthermore, MixNN outperforms other methods in the ILSVRC12 validation set, demonstrating MixNN’s great generalization capability.

### C. Ablation Study

We study the effect of removing different components to provide insights into what makes MixNN successful. The results are in Table IV. First, we remove the dynamic weight estimation in MixNN. Instead, we average (AVG) the selected samples or use random (RDM) weights from a Beta distribution in mixing function. We observe that merely averaging the selected samples does not perform well in all noisy cases, especially when the noise ratio is large. In comparison, using random weights may result in a surprising performance. For instance, it achieves excellent performance in CIFAR-100 with 40% asymmetric label noise. Since asymmetric label noise is concentrated inside each class, the estimated weights may fail to capture the clean probability of difficult samples,

TABLE II: Comparison with state-of-the-art methods trained on Clothing1M. Results of other methods are taken from original papers. All methods use a ResNet-50 architecture pretrained on ImageNet.


Methods	CE	F-correction [23]	GCE [19]	Co-teaching [30]	SEAL [58]	SCE [20]	LRT [59]	Joint-Optim [24]	MixNN
Test accuracy	69.21	69.84	69.75	70.15	70.63	71.02	71.74	72.16	<b>72.39</b>

TABLE III: Comparison with state-of-the-art methods trained on (mini) WebVision. Results of other baseline methods are taken from original papers. All methods use an InceptionResNetV2 architecture.


		F-correction [23]	Decoupling [31]	D2L [28]	MentorNet [29]	Co-teaching [30]	Iterative-CV [60]	Crust [61]	ODD [62]	MixNN
WebVision	top1	61.12	62.54	62.68	63.00	63.58	65.24	72.40	74.60	<b>75.39</b>
	top5	82.68	84.74	84.00	81.40	85.20	85.34	89.56	90.60	<b>90.71</b>
ILSVRC12	top1	57.36	58.26	57.80	57.80	61.48	61.60	67.36	66.70	<b>70.54</b>
	top5	82.36	82.26	81.36	79.92	84.70	84.98	87.84	86.30	<b>90.31</b>

TABLE IV: Ablation study results in terms of test accuracy (%) on CIFAR-10 and CIFAR-100.


Dataset	Noise type	Noise ratio	CIFAR-10				CIFAR-100		
			symm		asymm		symm		asymm
			40%	80%	40%	40%	80%	40%	
MixNN			<b>92.89 ± 0.02</b>	<b>86.08 ± 1.01</b>	<b>90.25 ± 0.76</b>	<b>72.97 ± 0.14</b>	<b>48.81 ± 0.06</b>	68.18 ± 0.11	
MixNN w/o weight estimation (AVG)			88.79 ± 0.33	58.53 ± 0.15	82.27 ± 0.10	64.21 ± 0.19	27.91 ± 0.31	59.73 ± 0.21	
MixNN w/o weight estimation (RDM)			91.92 ± 0.09	78.72 ± 0.13	89.91 ± 0.27	71.08 ± 0.25	41.14 ± 0.18	<b>72.98 ± 0.47</b>	
MixNN w/o label correction			91.73 ± 0.01	74.31 ± 0.49	88.73 ± 0.11	66.09 ± 0.50	30.18 ± 0.52	63.66 ± 0.43	
MixNN w/o $K$ nearest neighbours			92.19 ± 0.04	76.58 ± 0.25	88.69 ± 0.17	71.46 ± 0.08	42.02 ± 0.48	68.61 ± 0.16	



(a) The accuracy vs. epochs.



(b) The gradient coefficient vs. epochs.



(c) The influence of  $K$  on test accuracy (%).

Fig. 7: For plots (a) and (b), we train MixNN on CIFAR-10 with 60% symmetric label noise. Plot (a) shows the training and test accuracy vs. the number of epochs. Plot (b) shows the gradient coefficient vs. the number of epochs. Plot (c) shows the effect of different  $K$  on performance.

forcing the mixing function to use an average way. However, using the random weights is likely to appropriately assign the weights for the challenging samples, reducing the impact of asymmetric label noise. We also remove the label correction to see how it affects the performance. Without label correction, the performance suffers, especially when the noise ratio is high. Further, we investigate the performance of MixNN when using random samples instead of using  $K$  nearest neighbours. We observe that the performance is marginally poorer in the low ratio of label noise. However, when the training data contains high ratios of noisy labels, the performance suffers a large decline, demonstrating the benefit of using  $K$  nearest neighbours in MixNN.

#### D. Learning Stability and Gradient Analysis

In section III, we have demonstrated the failure of CE when trained DNNs with noisy labels. The training accuracy constantly increases indicates the DNNs eventually overfit the

noisy labels, resulting in a drop of accuracy on the clean test set. To verify the denoising effect of MixNN, we show its learning stability by plotting the training and test accuracy vs. the number of epochs in Fig. 7(a) on the CIFAR-10 with 60% label noise. We observe that the training accuracy stabilizes at around 40% after 60 epochs. It means the DNNs only fit the clean samples, resulting in no drop in test accuracy.

We further investigate the gradient coefficient of clean and mislabeled samples when using MixNN. In Fig. 7(b). We can observe the differences compared to the CE gradient coefficient in Fig. 2. The gradient of mislabeled samples is close to 0, letting the gradient of clean samples dominates the whole gradient throughout the training. Therefore, MixNN forces the model to learn from clean samples rather than mislabeled samples. In addition, we notice that there is an obvious gradient drop for clean samples when starting using noisy label correction (Note that we start label correction from epoch 60), which demonstrates the label correction effectively




Fig. 8: Confusion matrix of corrected labels w.r.t clean labels on CIFAR-10.

TABLE V: Correction accuracy (%) on CIFAR-10 and CIFAR-100 with various levels of label noise injected to training set.

Dataset	CIFAR-10					CIFAR-100				
	Noise type					Noise type				
	Noise ratio					Noise ratio				
	20%	40%	60%	80%	40%	20%	40%	60%	80%	40%
Correction accuracy (%)	96.78	95.25	92.89	86.77	91.81	90.71	84.86	74.93	52.45	72.40

suppresses gradient inversion as mentioned in Section III-C.

### E. Choice of $K$

The mixing functions Eq. 4 and Eq. 5 contain a parameter  $K$  to control the number of nearest neighbours mixed with original sample. To study the effect of  $K$  on our approach, we test different  $K$  on CIFAR-10. As shown in Fig. 7(c), MixNN achieves the worst performance as it reduces to regular training (i.e. CE) when  $K = 0$ . However, the performance of MixNN starts decreasing with the continuing increase of  $K$ . To explore whether the performance drop is due to dynamic weight estimation, we add another set of experiments where we simply use averaging strategy (abbreviated as MixAVG). We observe that MixAVG has a similar tendency on performance with the increase of  $K$ . Therefore, dynamic weight estimation is not the reason causing performance drop. We then conjecture that the reason is the over-mixture of input images. Since the resulting image is the mixture of  $K$  images, larger  $K$  leads to a more complex input. To avoid the influence of noisy labels, we conduct another set of experiments that use MixAVG on clean training data. We observe a similar performance drop when increasing  $K$ . Therefore, a reasonable  $K$  can suit our needs, while a large  $K$  makes it more difficult for the model to learn from complicated inputs, resulting in performance degradation.

### F. Effectiveness of Noisy Label Correction

Recall that we perform noisy label correction in Section IV-E. Since the estimated target  $t_i$  is calculated by an exponential moving average between the given noisy labels and model predictions, our method is able to gradually correct the noisy labels. The correction accuracy can be calculated by  $\frac{1}{N} \sum_i \mathbb{1}\{\arg\max \mathbf{y}_i = \arg\max \mathbf{t}_i\}$ , where  $\mathbf{y}_i$  is the ground truth label of training sample  $x_i$ . We evaluate the correction accuracy on CIFAR-10 and CIFAR-100 with different levels of label noise. As we can see the results in Table V, our

method successfully corrects a huge amount of wrong labels and obtains high correction accuracy in all cases. We also plot the confusion matrix of corrected labels w.r.t the clean labels on CIFAR-10 with 40% and 80% symmetric label noise and 40% asymmetric label noise in Figure 8. Our approach corrects the noisy labels impressively well for all classes under different level of label noise. We also observe that class dog and class cat are the most similar classes in CIFAR-10, which greatly increases the difficulty in the label correction process. As a result, our mixing strategy is more reliable to obtain the accurate mixed labels using the corrected soft labels.

### G. Feature Representations

We further investigate the representations learned by our approach compared to that learned by traditional cross entropy loss and SCE [20]. We extract the high-dimensional representations at the penultimate layer and project them to a 2D embedding by using t-SNE [47]. The projected representations are illustrated in Fig. 9 for 20%, 40%, 60% and 80% symmetric label noise respectively. Under all settings, the feature representations learned by our approach are of significantly better quality than that of CE and SCE with more separated and clearly bounded clusters. We find that SCE always keeps the feature representations of mislabeled samples in their true classes, the same as what CE does in the early learning stage, which prevents the model from memorizing them. However, the boundary formed by the SCE becomes increasingly blurred as the noise ratio rises (see  $\varepsilon = 0.8$  case). In contrast, our approach gradually corrects the noisy labels to clean labels, resulting in most of the feature representations in different classes are corrected.

## VII. CONCLUSION

In this paper, we explore the representation distribution in the early learning phase and propose MixNN for robust

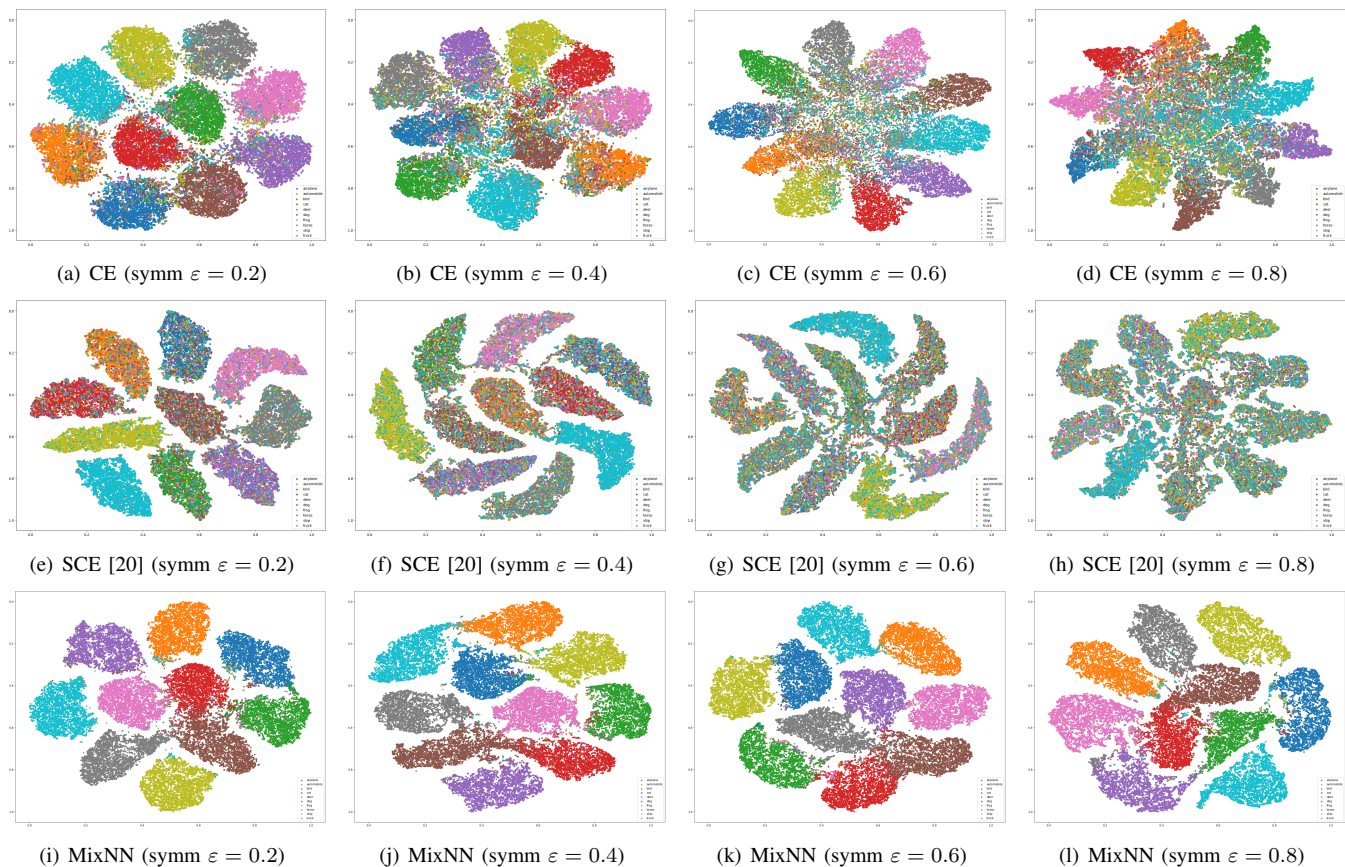


Fig. 9: t-SNE plots of feature representations learned by CE, SCE, and our proposed method MixNN on CIFAR-10 with different ratios of label noise. Different colours represent the different classes in CIFAR-10.

learning with noisy labels. Our approach mitigates the negative influence of noisy labels by training with the synthetic samples obtained by mixing the original training samples with their  $K$  nearest neighbours. The mixing procedure is dynamically adjusted by the learned Gaussian mixture model on per-sample loss distribution. We also propose a strategy that gradually corrects the noisy labels by using an exponential moving average on the given labels and model predictions. Through extensive experiments across multiple datasets with simulated and real-world label noise, we demonstrate that MixNN consistently exhibits substantial performance improvements compared to state-of-the-art methods. Importantly, the proposed approach works with any classifier “out-of-the-box” without any changes to architecture or training procedure. We are interested in adapting MixNN to other domains such as natural language process (NLP) and object detection, and believe MixNN is a promising algorithm for training robust DNNs against noisy labels. We hope that our work will trigger interest in the design of new approaches that provide robustness to label noise in real-world applications.

## REFERENCES

- [1] X. Fang, S. Teng, Z. Lai, Z. He, S. Xie, and W. K. Wong, “Robust latent subspace learning for image classification,” *IEEE transactions on neural networks and learning systems*, vol. 29, no. 6, pp. 2502–2515, 2017.
- [2] Z.-Q. Zhao, P. Zheng, S.-t. Xu, and X. Wu, “Object detection with deep learning: A review,” *IEEE transactions on neural networks and learning systems*, vol. 30, no. 11, pp. 3212–3232, 2019.
- [3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255, Ieee, 2009.
- [4] M. Li, M. Soltanolkotabi, and S. Oymak, “Gradient descent with early stopping is provably robust to label noise for overparameterized neural networks,” in *International Conference on Artificial Intelligence and Statistics*, pp. 4313–4324, PMLR, 2020.
- [5] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, “Understanding deep learning requires rethinking generalization,” 2018.
- [6] X. Zhu and X. Wu, “Class noise vs. attribute noise: A quantitative study,” *Artificial intelligence review*, vol. 22, no. 3, pp. 177–210, 2004.
- [7] S. Thulasidasan, T. Bhattacharya, J. Bilmes, G. Chennupati, and J. Mohd-Yusof, “Combating label noise in deep learning using abstention,” in *International Conference on Machine Learning*, pp. 6234–6243, PMLR, 2019.
- [8] Y. Kim, J. Yim, J. Yun, and J. Kim, “Nlnl: Negative learning for noisy labels,” in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 101–110, 2019.
- [9] J. Huang, L. Qu, R. Jia, and B. Zhao, “O2u-net: A simple noisy label detection approach for deep neural networks,” in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3326–3334, 2019.
- [10] T. DeVries and G. W. Taylor, “Improved regularization of convolutional neural networks with cutout,” *arXiv preprint arXiv:1708.04552*, 2017.
- [11] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.
- [12] 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 *Proceedings of the 34th*

- International Conference on Machine Learning-Volume 70*, pp. 233–242, JMLR. org, 2017.
- [13] H. Permuter, J. Francos, and I. Jermyn, “A study of gaussian mixture models of color and texture features for image classification and segmentation,” *Pattern recognition*, vol. 39, no. 4, pp. 695–706, 2006.
- [14] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *NIPS Workshop*, 2014.
- [15] H. Masnadi-Shirazi and N. Vasconcelos, “On the design of loss functions for classification: theory, robustness to outliers, and savageboost,” in *Advances in neural information processing systems*, pp. 1049–1056, 2009.
- [16] J. P. Brooks, “Support vector machines with the ramp loss and the hard margin loss,” *Operations research*, vol. 59, no. 2, pp. 467–479, 2011.
- [17] B. Van Rooyen, A. Menon, and R. C. Williamson, “Learning with symmetric label noise: The importance of being unhinged,” in *Advances in Neural Information Processing Systems*, pp. 10–18, 2015.
- [18] A. Ghosh, H. Kumar, and P. Sastry, “Robust loss functions under label noise for deep neural networks,” in *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [19] Z. Zhang and M. Sabuncu, “Generalized cross entropy loss for training deep neural networks with noisy labels,” in *Advances in neural information processing systems*, pp. 8778–8788, 2018.
- [20] Y. Wang, X. Ma, Z. Chen, Y. Luo, J. Yi, and J. 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.
- [21] X. Wang, Y. Hua, E. Kodirov, and N. M. Robertson, “Imae for noise-robust learning: Mean absolute error does not treat examples equally and gradient magnitude’s variance matters,” *arXiv preprint arXiv:1903.12141*, 2019.
- [22] X. Ma, H. Huang, Y. Wang, S. Romano, S. Erfani, and J. Bailey, “Normalized loss functions for deep learning with noisy labels,” in *International Conference on Machine Learning*, pp. 6543–6553, PMLR, 2020.
- [23] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1944–1952, 2017.
- [24] D. Tanaka, D. Ikami, T. Yamasaki, and K. Aizawa, “Joint optimization framework for learning with noisy labels,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5552–5560, 2018.
- [25] Y. Wang, W. Liu, X. Ma, J. Bailey, H. Zha, L. Song, and S.-T. Xia, “Iterative learning with open-set noisy labels,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8688–8696, 2018.
- [26] K. Yi and J. Wu, “Probabilistic end-to-end noise correction for learning with noisy labels,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7017–7025, 2019.
- [27] S. E. Reed, H. Lee, D. Anguelov, C. Szegedy, D. Erhan, and A. Rabinovich, “Training deep neural networks on noisy labels with bootstrapping,” in *ICLR (Workshop)*, 2015.
- [28] X. Ma, Y. Wang, M. E. Houle, S. Zhou, S. M. Erfani, S.-T. Xia, S. Wijewickrema, and J. Bailey, “Dimensionality-driven learning with noisy labels,” *arXiv preprint arXiv:1806.02612*, 2018.
- [29] 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 *International Conference on Machine Learning*, pp. 2304–2313, PMLR, 2018.
- [30] 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 *Advances in neural information processing systems*, pp. 8527–8537, 2018.
- [31] E. Malach and S. Shalev-Shwartz, “Decoupling” when to update” from” how to update”,” in *Advances in Neural Information Processing Systems*, pp. 960–970, 2017.
- [32] B. Han, I. W. Tsang, L. Chen, P. Y. Celina, and S.-F. Fung, “Progressive stochastic learning for noisy labels,” *IEEE transactions on neural networks and learning systems*, vol. 29, no. 10, pp. 5136–5148, 2018.
- [33] X. Yu, B. Han, J. Yao, G. Niu, I. W. Tsang, and M. Sugiyama, “How does disagreement help generalization against label corruption?,” *arXiv preprint arXiv:1901.04215*, 2019.
- [34] H. Wei, L. Feng, X. Chen, and B. 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*, pp. 13726–13735, 2020.
- [35] Y. Lu, Y. Bo, and W. He, “Co-matching: Combating noisy labels by augmentation anchoring,” *arXiv preprint arXiv:2103.12814*, 2021.
- [36] Y. Ding, L. Wang, D. Fan, and B. Gong, “A semi-supervised two-stage approach to learning from noisy labels,” in *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1215–1224, IEEE, 2018.
- [37] J. Li, R. Socher, and S. C. Hoi, “Dividemix: Learning with noisy labels as semi-supervised learning,” *arXiv preprint arXiv:2002.07394*, 2020.
- [38] T. Nguyen, C. Mummadi, T. Ngo, L. Beggel, and T. Brox, “Self: learning to filter noisy labels with self-ensembling,” in *International Conference on Learning Representations (ICLR)*, 2020.
- [39] J. Li, Y. Wong, Q. Zhao, and M. S. Kankanhalli, “Learning to learn from noisy labeled data,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5051–5059, 2019.
- [40] B. Han, G. Niu, J. Yao, X. Yu, M. Xu, I. Tsang, and M. Sugiyama, “Pumpout: A meta approach for robustly training deep neural networks with noisy labels,” 2018.
- [41] D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, and C. A. Raffel, “Mixmatch: A holistic approach to semi-supervised learning,” in *Advances in Neural Information Processing Systems*, pp. 5050–5060, 2019.
- [42] A. Vahdat, “Toward robustness against label noise in training deep discriminative neural networks,” in *Advances in Neural Information Processing Systems*, pp. 5596–5605, 2017.
- [43] Y. Lu, Y. Bo, and W. He, “Confidence adaptive regularization for deep learning with noisy labels,” *arXiv preprint arXiv:2108.08212*, 2021.
- [44] J. Han, P. Luo, and X. Wang, “Deep self-learning from noisy labels,” in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5138–5147, 2019.
- [45] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- [46] S. Liu, J. Niles-Weed, N. Razavian, and C. Fernandez-Granda, “Early-learning regularization prevents memorization of noisy labels,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [47] L. Van der Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [48] S. Har-Peled, P. Indyk, and R. Motwani, “Approximate nearest neighbor: Towards removing the curse of dimensionality,” *Theory of computing*, vol. 8, no. 1, pp. 321–350, 2012.
- [49] J. L. Bentley, “Multidimensional binary search trees used for associative searching,” *Communications of the ACM*, vol. 18, no. 9, pp. 509–517, 1975.
- [50] Y. A. Malkov and D. A. Yashunin, “Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 4, pp. 824–836, 2018.
- [51] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” *arXiv preprint arXiv:1710.09412*, 2017.
- [52] S. Thulasidasan, G. Chennupati, J. A. Bilmes, T. Bhattacharya, and S. Michalak, “On mixup training: Improved calibration and predictive uncertainty for deep neural networks,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [53] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- [54] I. Loshchilov and F. Hutter, “Sgdr: Stochastic gradient descent with warm restarts,” *arXiv preprint arXiv:1608.03983*, 2016.
- [55] L. Prechelt, “Early stopping-but when?,” in *Neural Networks: Tricks of the trade*, pp. 55–69, Springer, 1998.
- [56] K. Lee, S. Yun, K. Lee, H. Lee, B. Li, and J. Shin, “Robust inference via generative classifiers for handling noisy labels,” in *International Conference on Machine Learning*, pp. 3763–3772, PMLR, 2019.
- [57] E. Arazo, D. Ortego, P. Albert, N. O’Connor, and K. McGuinness, “Unsupervised label noise modeling and loss correction,” in *Proceedings of the 36th International Conference on Machine Learning*, pp. 312–321, 2019.
- [58] P. Chen, J. Ye, G. Chen, J. Zhao, and P.-A. Heng, “Beyond class-conditional assumption: A primary attempt to combat instance-dependent label noise,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, pp. 11442–11450, 2021.
- [59] S. Zheng, P. Wu, A. Goswami, M. Goswami, D. Metaxas, and C. Chen, “Error-bounded correction of noisy labels,” in *International Conference on Machine Learning*, pp. 11447–11457, PMLR, 2020.
- [60] P. Chen, B. B. Liao, G. Chen, and S. Zhang, “Understanding and utilizing deep neural networks trained with noisy labels,” in *International Conference on Machine Learning*, pp. 1062–1070, 2019.

- [61] B. Mirzasoleiman, K. Cao, and J. Leskovec, "Coresets for robust training of neural networks against noisy labels," *Neural Information Processing Systems (NeurIPS)*, 2020.
- [62] J. Song, Y. Dauphin, M. Auli, and T. Ma, "Robust and on-the-fly dataset denoising for image classification," in *European Conference on Computer Vision*, pp. 556–572, Springer, 2020.
- [63] T. Xiao, T. Xia, Y. Yang, C. Huang, and X. Wang, "Learning from massive noisy labeled data for image classification," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2691–2699, 2015.
- [64] W. Li, L. Wang, W. Li, E. Agustsson, and L. V. Gool, "Webvision database: Visual learning and understanding from web data.," *CoRR*, 2017.
- [65] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, inception-resnet and the impact of residual connections on learning," in *Thirty-first AAAI conference on artificial intelligence*, 2017.