

Do We Really Need Gold Samples for Sample Weighting under Label Noise?

Aritra Ghosh

University of Massachusetts Amherst
Amherst, MA, USA

arighosh@cs.umass.edu

Andrew Lan

University of Massachusetts Amherst
Amherst, MA, USA

andrewlan@cs.umass.edu

Abstract

Learning with labels noise has gained significant traction recently due to the sensitivity of deep neural networks under label noise under common loss functions. Losses that are theoretically robust to label noise, however, often makes training difficult. Consequently, several recently proposed methods, such as Meta-Weight-Net (MW-Net), use a small number of unbiased, clean samples to learn a weighting function that downweights samples that are likely to have corrupted labels under the meta-learning framework. However, obtaining such a set of clean samples is not always feasible in practice. In this paper, we analytically show that one can easily train MW-Net without access to clean samples simply by using a loss function that is robust to label noise, such as mean absolute error, as the meta objective to train the weighting network. We experimentally show that our method beats all existing methods that do not use clean samples and performs on-par with methods that use gold samples on benchmark datasets across various noise types and noise rates.

1. Introduction

Although deep neural networks (DNNs) have achieved impressive performance across many applications, they remain vulnerable to noisy labels. The expressive power of DNNs enables them to learn arbitrary smooth non-linear functions; however, it also means that DNNs can easily overfit to noisy labels, which negatively affects their generalization ability. Learning under label noise is of utmost importance in the big-data regime since large amounts of labels are collected using crowd-sourcing, resulting in significant label noise [48, 62, 11, 58, 48]; see [16, 1, 20] for detailed surveys and recent results on learning under label noise. Methods for mitigating label noise can be broadly categorized into two groups based on whether they use a small number of clean samples in the learning process.

In the first group, most algorithms do not assume access to any clean samples; instead, they use loss functions that

are robust to label noise to handle noisy labels implicitly [57, 17, 70]. Unfortunately, the common cross entropy loss used to train DNNs is particularly sensitive to noisy labels. As an alternative, in [17], the authors have shown that the non-convex mean absolute error (MAE) loss is robust to label noise under some mild assumptions. However, training DNNs under the MAE loss on large datasets is often difficult. In [70], the authors proposed generalized cross-entropy loss, a generalization of MAE loss and cross-entropy loss, that provides an effective trade-off between theoretical guarantees and the ease of DNN training [70]. However, the performance of these theoretically robust losses on large-scale image datasets is often limited.

In the second group, most methods use a small number of clean samples to compute the noise transition matrix, predict which samples are noisy, and correct the loss functions by learning a sample re-weighting strategy [46, 48, 54]. One method is to use a loss correction approach [42] when the noise rates are known or can be accurately estimated from data. Another approach is to use sample re-weighting to downweight samples for which the classifier incurs large losses on and are thus likely to be noisy. Many early methods use heuristics to compute sample weights. However, a pre-determined weighting scheme is not very effective and cannot leverage real data. Therefore, many recent methods proposed to adaptively learn the sample weighting function from the additional clean data [54, 49, 30]. Meta-Weight-Net (MW-Net) [54], a highly effective method, uses recent advances in *meta-learning* to jointly learn the classifier network and the weighting network. MW-Net uses a small number of clean samples as meta samples to learn the sample weighting strategy used for the classifier network via the meta objective. Experimentally, MW-Net achieves impressive performance on real-world, large-scale image datasets under label noise [54]. However, the most apparent drawback of MW-Net and other methods in this group is that we may not have access to clean samples in real-world applications.

Contributions We make a surprising observation that it is very easy to adaptively learn sample weighting functions, even when we do not have access to any clean samples; we

can use *noisy* meta samples to learn the weighting function if we simply change the meta loss function. We theoretically show that the meta-gradient direction for optimizing the weighting network remains the same when using noisy meta samples with meta losses that are robust to label noise such as the MAE [17]. Thus, we can optimize the classifier network using the cross-entropy loss and optimize the weighting network using the MAE loss, both with noisy samples. Although the MAE loss is difficult to optimize for deep classification networks [70], they are easy to optimize for shallow networks that are used as the weighting network [17]. Our method closes the gap on learning the weighting function adaptively without the need to access clean meta samples; this setup closely resembles real-world problem settings where clean samples are often missing.

First, we theoretically show that it is *possible* to learn the weighting function even without any access to clean meta samples. Then, we experimentally show that our proposed method (using noisy meta samples) performs on par with MW-Net (using clean meta samples) and beats existing methods on several benchmark image datasets. Thus, the simple observation that the meta-gradient direction remains the same for noisy meta samples under meta losses that are robust to label noise alleviates the need for clean samples when learning sample re-weighting strategy under label noise.¹

2. Problem Setup and Label Noise

We consider the classification problem setting with a training set $\{(\mathbf{x}_i^{\text{train}}, \mathbf{y}_i^{\text{train}})\}_{i=1}^N$, a meta (or validation) set $\{(\mathbf{x}_j^{\text{meta}}, \mathbf{y}_j^{\text{meta}})\}_{j=1}^M$, and a test set $\{(\mathbf{x}_o^{\text{test}}, \mathbf{y}_o^{\text{test}})\}_{o=1}^O$ where \mathbf{x}_i is the feature vector of the i^{th} sample and $\mathbf{y}_i \in \{0, 1\}^K$ is the class label vector with K total classes. We denote the classifier network prediction as $f(\mathbf{x}, \mathbf{w})$ where \mathbf{w} is the classifier network weights. We denote the clean (unknown) label for an instance \mathbf{x}_i as $\mathbf{y}_i^{\text{true}}$. We consider label noise where feature vectors are clean but the labels vectors are corrupted [42]. We consider three types of label noise. The uniform (or symmetric) noise with rate $\eta \in [0, 1]$, corrupts the true label with probability η uniformly distributed among all classes. Thus, the true class label remains the same with probability $(1 - \eta) + \frac{\eta}{K}$ and the true label is corrupted to all other classes $c \neq \mathbf{y}_i^{\text{true}}$ with probability $\frac{\eta}{K}$. The flip noise model, with noise rate η , corrupts the true labels with probability η to a random single class. Thus, the correct class label remains the same with probability $(1 - \eta)$ and the true label is corrupted to a single class $c \neq \mathbf{y}_i^{\text{true}}$ with probability η . The flip2 noise model, with noise rate η , corrupts the true labels with probability η to two random other classes. Thus, the correct class label remains the same with probability $(1 - \eta)$ and the true label is corrupted to two other classes

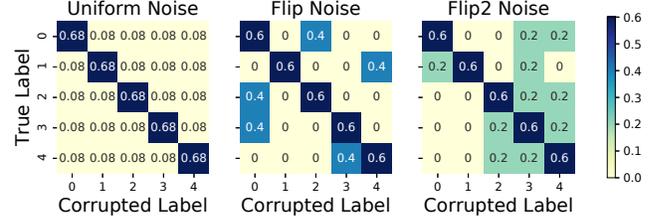


Figure 1. Confusion matrices for the noise models ($K=5, \eta=0.4$).

$c_1, c_2 \neq \mathbf{y}_i^{\text{true}}$ with probability $\frac{\eta}{2}$. Figure 1 shows example confusion matrices for these three noise models.

We assume that the training dataset $\{(\mathbf{x}_i^{\text{train}}, \mathbf{y}_i^{\text{train}})\}_{i=1}^N$ is corrupted with an unknown noise model, but the test set $\{(\mathbf{x}_o^{\text{test}}, \mathbf{y}_o^{\text{test}})\}_{o=1}^O$ remains clean. Deep networks are often optimized with the cross-entropy (CE) loss; we obtain the optimal classifier network parameter \mathbf{w}^* by minimizing CE loss on the training dataset as

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w}} \frac{1}{N} \sum_{i=1}^N \ell_{\text{CE}}(\mathbf{y}_i^{\text{train}}, f(\mathbf{x}_i^{\text{train}}, \mathbf{w}))$$

where ℓ_{CE} is the CE loss.

In sample weighting methods, a common approach is to introduce weighting function $\mathcal{W}(\mathbf{x}_i^{\text{train}}, \mathbf{y}_i^{\text{train}}, \mathbf{w}; \Theta)$, which is (optionally) parameterized by Θ to decide the weight for i^{th} training sample. In MW-Net, the function \mathcal{W} takes sample loss as the input [54]; thus, we can represent the sample weighting function as $\mathcal{W}(\ell_{\text{CE}}(\mathbf{y}_i^{\text{train}}, f(\mathbf{x}_i^{\text{train}}, \mathbf{w})); \Theta)$ where $\ell_{\text{CE}}(\mathbf{y}_i^{\text{train}}, f(\mathbf{x}_i^{\text{train}}, \mathbf{w}))$ is the current sample loss and Θ is the (meta) weight network parameter. For simplicity, throughout the paper, we use $\ell(\mathbf{y}_i^{\text{train}}, f(\mathbf{x}_i^{\text{train}}, \mathbf{w}))$ as $\ell^{i, \text{train}}(\mathbf{w})$ and $\ell(\mathbf{y}_j^{\text{meta}}, f(\mathbf{x}_j^{\text{meta}}, \mathbf{w}))$ as $\ell^{j, \text{meta}}(\mathbf{w})$. Thus, we can calculate classifier network weights \mathbf{w}^* under label noise with the sample weighting parameter Θ as [54]:

$$\operatorname{argmin}_{\mathbf{w}} \mathcal{L}^{\text{train}}(\mathbf{w}; \Theta) \triangleq \frac{1}{N} \sum_{i=1}^N \mathcal{W}(\ell_{\text{CE}}^{i, \text{train}}(\mathbf{w}); \Theta) \ell_{\text{CE}}^{i, \text{train}}(\mathbf{w}).$$

3. Methodology

We first discuss the basic setup of MW-Net [54]. We then show that it is possible to train a robust model without access to clean meta samples (thus, $\mathbf{y}_j^{\text{meta}}$ can be noisy as well).

Meta-Weight-Net. MW-Net learns an auxiliary neural network with parameter Θ for parameterizing the weighting function [54]. Many previous heuristics employ tricks such as removing training samples with a high loss to minimize the effect of noisy samples. Instead of using some fixed heuristics, we can learn a data-driven adaptive non-linear weighting function. MW-Net is an effective way to learn the weighting function using ideas from meta-learning. The idea is to use a clean meta dataset (or validation dataset) $\{(\mathbf{x}_j^{\text{meta}}, \mathbf{y}_j^{\text{meta}})\}_{j=1}^M$ ($M \ll N$) to learn the optimal weighting

¹Our code is available at <https://github.com/arghosh/RobustMW-Net>.

function Θ^* in a bilevel optimization problem [15].

$$\begin{aligned} \min_{\Theta} \mathcal{L}^{\text{meta}}(\mathbf{w}^*(\Theta)) &\triangleq \frac{1}{M} \sum_{j=1}^M \ell^{j,\text{meta}}(\mathbf{w}^*(\Theta)) \\ \text{s.t. } \mathbf{w}^*(\Theta) &= \underset{\mathbf{w}}{\text{argmin}} \frac{1}{N} \sum_{i=1}^N \mathcal{W}(\ell_{\text{CE}}^{i,\text{train}}(\mathbf{w}); \Theta) \ell_{\text{CE}}^{i,\text{train}}(\mathbf{w}) \end{aligned} \quad (1)$$

Thus, the inner objective is to optimize \mathbf{w}^* on $\mathcal{L}^{\text{train}}$ and the outer objective is to optimize for Θ^* on $\mathcal{L}^{\text{meta}}$ with nested loops. In the outer level optimization, the weighting network computes meta loss $\mathcal{L}^{\text{meta}}$ using the loss function $\ell^{\cdot,\text{meta}}$; for brevity, we refer $\ell^{\cdot,\text{meta}}$ as the *loss function of the weighting network* or the *meta loss function* throughout the paper.

In [54], the authors proposed a bilevel online optimization strategy to iteratively optimize both \mathbf{w} and Θ [15]. The gradient for Θ , in Eq. 1 at time step t , requires the function $\hat{\mathbf{w}}(\Theta)$. The classifier weight function $\hat{\mathbf{w}}^t(\Theta)$ (note, this is a function of Θ , not the weight parameter) is computed as:

$$\hat{\mathbf{w}}^t(\Theta) = \mathbf{w}^t - \alpha \frac{1}{n} \sum_{i=1}^n \mathcal{W}(\ell_{\text{CE}}^{i,\text{train}}(\mathbf{w}^t); \Theta) \nabla_{\mathbf{w}} \ell_{\text{CE}}^{i,\text{train}}(\mathbf{w}) \Big|_{\mathbf{w}^t}$$

where n is the minibatch size for the training dataset, and α is the SGD optimizer parameter. Using the function $\hat{\mathbf{w}}^t(\Theta)$, we can take a gradient step for Θ^t based on the loss on the meta dataset in Eq. 1.

$$\Theta^{t+1} = \Theta^t - \beta \frac{1}{m} \sum_{j=1}^m \nabla_{\Theta} \ell^{j,\text{meta}}(\hat{\mathbf{w}}^t(\Theta)) \Big|_{\Theta^t} \quad (2)$$

where β is the weighting network SGD parameter and m is the minibatch size for the meta dataset. Using the updated Θ^{t+1} , we can take a step towards optimal \mathbf{w}^* using simple stochastic gradient descent step

$$\mathbf{w}^{t+1} = \mathbf{w}^t - \alpha \frac{1}{n} \sum_{i=1}^n \mathcal{W}(\ell_{\text{CE}}^{i,\text{train}}(\mathbf{w}^t); \Theta^{t+1}) \nabla_{\mathbf{w}} \ell_{\text{CE}}^{i,\text{train}}(\mathbf{w}) \Big|_{\mathbf{w}^t}.$$

In essence, MW-Net takes gradient descent step for Θ^t towards Θ^* and takes gradient descent step for \mathbf{w}^t towards \mathbf{w}^* using the updated Θ^{t+1} in an online fashion. In [54], the authors use a single layer (with 100 hidden nodes) feed-forward networks as the weighting network. Note, the input to the weight network is the scalar loss value $\ell_{\text{CE}}^{i,\text{train}}(\mathbf{w}(\Theta))$ and the output of the weight network is the scalar sample weight value. The choice of weighting network is effective since a single hidden layer MLP is a universal approximator for any continuous smooth functions. Moreover, the weighting function is mostly a monotonically decreasing function under label noise.

Meta-Weight-Net with noisy meta samples. We can now analyze MW-Net and show how it can work without access to clean meta labels. We first discuss the gradient descent direction of the weighting network with clean meta samples. We denote the term

$$\mathcal{G}(\hat{\mathbf{w}}) = \frac{1}{m} \sum_{j=1}^m \frac{\partial \ell^{j,\text{meta}}(\hat{\mathbf{w}})}{\partial \hat{\mathbf{w}}} \Big|_{\hat{\mathbf{w}}^t} \quad (3)$$

as the average *meta-gradient* for \mathbf{w} on the meta dataset. Taking derivative of $\ell^{j,\text{meta}}(\hat{\mathbf{w}}^t(\Theta))$ w.r.t Θ in Eq. 2, the Θ gradient descent direction is,

$$\begin{aligned} \nabla_{\Theta} \sum_{j=1}^m \ell^{j,\text{meta}}(\hat{\mathbf{w}}^t(\Theta)) \\ = -\frac{\alpha}{n} \sum_{i=1}^n \left(\mathcal{G}(\hat{\mathbf{w}})^{\top} \frac{\partial \ell_{\text{CE}}^{i,\text{train}}(\mathbf{w})}{\partial \mathbf{w}} \Big|_{\mathbf{w}^t} \right) \frac{\partial \mathcal{W}(\ell_{\text{CE}}^{i,\text{train}}(\mathbf{w}^t); \Theta)}{\partial \Theta} \Big|_{\Theta^t} \end{aligned} \quad (4)$$

where m and n is the minibatch size for meta dataset and training dataset respectively. We can understand this update direction as a sum of weighted gradient updates for each training samples. The term $\frac{\partial \mathcal{W}(\ell_{\text{CE}}^{i,\text{train}}(\mathbf{w}^t); \Theta)}{\partial \Theta}$ represents the gradient direction for Θ on training sample i . The weight for training sample i is computed using $\mathcal{G}(\hat{\mathbf{w}})^{\top} \frac{\partial \ell_{\text{CE}}^{i,\text{train}}(\mathbf{w})}{\partial \mathbf{w}} \Big|_{\mathbf{w}^t}$; thus, the weighting network effectively puts more weights for training samples having a gradient direction similar to the average meta-gradient direction \mathcal{G} .

For corrupted meta samples, we denote the meta loss function as, $\mathcal{L}^{\text{noisy-meta}}(\mathbf{w}^*(\Theta)) \triangleq \frac{1}{M} \sum_{j=1}^M \ell^{j,\text{noisy-meta}}(\mathbf{w}^*(\Theta))$, where we use $\ell^{j,\text{noisy-meta}}$ to denote the loss on the corrupted j^{th} meta sample (and $\ell^{j,\text{meta}}$ to denote the loss for the j^{th} clean meta sample). Therefore, the key idea is that if we need to use noisy meta dataset (training dataset can be corrupted arbitrarily and independently), we need to ensure

$$\sum_{j=1}^m \frac{\partial \ell^{j,\text{meta}}(\hat{\mathbf{w}})}{\partial \hat{\mathbf{w}}} \approx C \sum_{j=1}^m \frac{\partial \ell^{j,\text{noisy-meta}}(\hat{\mathbf{w}})}{\partial \hat{\mathbf{w}}}. \quad (5)$$

C can be any positive constant; we only care about gradient direction in SGD. Note in Eq. 4, only the average meta-gradient ($\mathcal{G}(\hat{\mathbf{w}})$) term changes if we use noisy meta samples; gradient for training samples remains the same.

We show that when the outer meta loss function is MAE (thus $\ell^{j,\text{meta}}$ is $\ell_{\text{MAE}}^{j,\text{meta}}$), the noisy (under uniform noise) average meta-gradients remain the same as the average meta-gradients on clean meta dataset. In [54], the CE loss function is used with clean meta samples ($\ell^{j,\text{meta}}$ is $\ell_{\text{CE}}^{j,\text{meta}}$). Note, the classifier network still uses cross-entropy loss (we still use $\ell_{\text{CE}}^{i,\text{train}}$); we need to maintain average meta-gradient direction for meta samples only. We denote the classifier output (softmax probability vector) $f(\mathbf{x}_j, \mathbf{w})$ as $\mathbf{u}_j \in \Delta^{K-1}$ where Δ^{K-1} is the $K-1$ dimensional simplex; MAE and CE losses

are computed as, $\ell_{MAE}(\mathbf{y}_j, \mathbf{u}_j) = \sum_k |\mathbf{u}_{j,k} - \mathbf{y}_{j,k}|$ and $\ell_{CE}(\mathbf{y}_j, \mathbf{u}_j) = -\sum_k \mathbf{y}_{j,k} \log \mathbf{u}_{j,k}$. Note that MAE loss ℓ_{MAE} is bounded and has a symmetric property [17] as:

$$\sum_{c=1}^K \ell(c, f(\mathbf{x}, \mathbf{w})) = \text{constant}, \forall \mathbf{x}, \text{ and } \forall \mathbf{w}. \quad (6)$$

Theorem 1. *Let the meta samples are corrupted with uniform noise with a rate of $\eta < 1$. Suppose the weighting network minimize loss function ℓ on the meta samples and the loss function ℓ satisfies symmetric property in Eq. 6. Then the expected meta-gradient on corrupted meta samples (for updating the weighting network in Eq. 4) is exactly the same as (up to a proportional constant) the meta-gradient on clean meta samples irrespective of the noise model in the training dataset and the loss function in the classifier network.*

Proof. We can compute the expected gradient for \mathbf{w} on a batch of samples from the noisy meta dataset under uniform label noise (with the rate η) as:

$$\begin{aligned} & \mathbb{E} \left[\sum_{j=1}^m \frac{\partial \ell^{j, \text{noisy-meta}}(\hat{\mathbf{w}})}{\partial \hat{\mathbf{w}}} \right] \\ &= \sum_{j=1}^m \left[(1-\eta) \frac{\partial \ell(\mathbf{y}_j, f(\mathbf{x}_j, \hat{\mathbf{w}}))}{\partial \hat{\mathbf{w}}} + \frac{\eta}{K} \sum_{c=1}^K \frac{\partial \ell(c, f(\mathbf{x}_j, \hat{\mathbf{w}}))}{\partial \hat{\mathbf{w}}} \right] \\ &= \sum_{j=1}^m \left[(1-\eta) \frac{\partial \ell(\mathbf{y}_j, f(\mathbf{x}_j, \hat{\mathbf{w}}))}{\partial \hat{\mathbf{w}}} + \frac{\eta}{K} \frac{\partial}{\partial \hat{\mathbf{w}}} \sum_{c=1}^K \ell(c, f(\mathbf{x}_j, \hat{\mathbf{w}})) \right] \\ &= (1-\eta) \sum_{j=1}^m \left[\frac{\partial \ell(\mathbf{y}_j, f(\mathbf{x}_j, \hat{\mathbf{w}}))}{\partial \hat{\mathbf{w}}} \right] = C \sum_{j=1}^m \frac{\partial \ell^{j, \text{meta}}(\hat{\mathbf{w}})}{\partial \hat{\mathbf{w}}} \end{aligned}$$

since $\frac{\partial}{\partial \hat{\mathbf{w}}} \sum_{c=1}^K \ell(c, f(\mathbf{x}_i, \hat{\mathbf{w}})) = 0$ when ℓ is a symmetric loss (such as MAE). Thus, under uniform label noise on the meta dataset, the expected meta-gradient on the noisy meta samples remains the same as the meta-gradient on the clean meta dataset when the meta loss function is symmetric. \square

Theorem 1 shows that symmetric losses on the corrupted meta samples for optimizing the weighting network has the same expected meta-gradient direction as the clean samples irrespective of the classifier network loss function and arbitrarily corrupted training datasets. Moreover, we can show convergence of the weighting network with finite mini batch size of the corrupted meta datasets under some mild condition. Similar to [54], we make the following assumptions.

1. The meta loss ℓ and the classifier network loss $\ell^{\cdot, \text{train}}$ is Lipschitz smooth with constant L and have ρ -bounded gradients.
2. The weighting function $\mathcal{W}(\cdot)$ has bounded gradient and twice differential with bounded Hessian.
3. The classifier network learning rate α_t satisfies $\alpha_t = \min\{1, \frac{k}{T}\}$ for some $k < T$, $k > 0$. The learning rate of the weighting network β_t satisfies $\beta_t = \min\{\frac{1}{L}, \frac{b}{\hat{\sigma}\sqrt{T}}\}$ for some $b > 0$ such that $\frac{\hat{\sigma}\sqrt{T}}{b} \geq L$ where $\hat{\sigma}^2$ is the

variance of drawing a minibatch (possibly corrupted with uniform noise).

Theorem 2. *Under assumptions 1-3, the weighting network converges for both clean and (uniformly) corrupted meta datasets when the meta loss satisfies symmetric condition:*

$$\begin{aligned} & \min_{0 \leq t \leq T} \mathbb{E}[\|\nabla \mathcal{L}^{\text{meta}}(\Theta^t)\|_2^2] \leq \mathcal{O}\left(\frac{\sigma}{\sqrt{T}}\right), \text{ and} \\ & \min_{0 \leq t \leq T} \mathbb{E}[\|\nabla \mathcal{L}^{\text{noisy-meta}}(\Theta^t)\|_2^2] \leq \mathcal{O}\left(\frac{\hat{\sigma}}{(1-\eta)\sqrt{T}}\right), \end{aligned}$$

where σ^2 is the variance of drawing uniformly mini-batch sample at random, $\hat{\sigma}^2 = \sigma^2 + \frac{2\eta\rho^2}{m}$ is the variance adjusted with the uniformly corrupted meta samples, η is the uniform noise rate and m is the minibatch size of the meta dataset².

Remark. Using the convergence rate in Theorem 2, we would expect to approximate the meta-gradient under noisy labels to be similar to the gradient under clean labels with a relatively large mini-batch size [50]. Thus, we can use corrupted meta dataset as long as the meta loss function is symmetric [17]. The classifier network can use CE loss; the weighting network only needs to use symmetric loss for computing meta-gradients from the corrupted meta samples. MAE loss is difficult to optimize for deep networks [70]; however, we note that the weighting network is simple with a single input (the loss value) and a single layer with 100 hidden nodes as used in [54]. We find that the weighting network with MAE loss for the meta samples is amiable to SGD; we also find MAE loss to be better than cross-entropy for the weighting network under varying noise rates and noise models. Although, MAE loss can provide a guarantee for the meta dataset corrupted with uniform label noise; the training datasets do not require any such condition; we can potentially handle training datasets with *instance-dependent label noise* also. We experimentally found that MAE loss significantly outperforms CE loss in the weighting network for uniform label noise. Nevertheless, we also observe that CE loss in the weighting network performs relatively well for noisy meta datasets. We believe as the weighting network is a shallow network with single input and output and equipped with l_2 regularizer (in the form of weight decay), $\frac{\eta}{K} \mathbb{E} \left[\frac{\partial \sum_{c=1}^K \ell_{CE}(c, f(\mathbf{x}, \hat{\mathbf{w}}))}{\partial \hat{\mathbf{w}}} \right]$ is relatively small compared to $(1-\eta) \mathbb{E} \left[\frac{\partial \ell_{CE}(\mathbf{y}, f(\mathbf{x}, \hat{\mathbf{w}}))}{\partial \hat{\mathbf{w}}} \right]$ even for CE loss.

Under the flip noise, we randomly corrupt samples from a single class to a single different random class with a probability of η . Under the flip2 noise, we corrupt uniformly between random two fixed classes. Under a flip $\{K-1\}$ noise model, labels are corrupted to $K-1$ other classes, effectively becoming uniform noise. We consider the flip noise first; we denote the random corrupted class labels as

²Derivation is available in supplementary material.

$\mathcal{T}(\mathbf{y})$, $\mathcal{T}(\mathbf{y}) \neq \mathbf{y}$ where samples with true labels \mathbf{y} are corrupted as $\mathcal{T}(\mathbf{y})$ with probability η . However, once $\mathcal{T}(\mathbf{y})$ is drawn randomly as part of the noise model, it is deterministic. Hence, the expectation of gradient on noisy meta samples do not remain the same as the gradient on the clean meta samples for symmetric losses:

$$\frac{\partial \mathcal{L}^{\text{noisy-meta}}}{\partial \hat{\mathbf{w}}} = \sum_{j=1}^m \left[(1-\eta) \frac{\partial \ell(\mathbf{y}_j, f(\mathbf{x}_j, \hat{\mathbf{w}}))}{\partial \hat{\mathbf{w}}} + \eta \frac{\partial \ell(\mathcal{T}(\mathbf{y}_j), f(\mathbf{x}_j, \hat{\mathbf{w}}))}{\partial \hat{\mathbf{w}}} \right].$$

Hence the theoretical guarantees in Eq. 5 only holds for uniform noise model. Nevertheless, taking expectation w.r.t. to $\mathcal{T}(\mathbf{y})$, we get the same gradient as that of clean ones since $\mathbb{E}_{\mathcal{T}(\mathbf{y})} \left[\frac{\partial \mathcal{L}^{\text{noisy-meta}}}{\partial \hat{\mathbf{w}}} \right] = C \sum_{j=1}^m \frac{\partial \ell^{j, \text{meta}}(\hat{\mathbf{w}})}{\partial \hat{\mathbf{w}}}$. Thus, as we flip to more classes, we would expect more uniform label noise behavior resulting in the theoretical robustness of the weighting network using robust losses. However, we found that when flip noise is relatively small, we can still use MW-Net with noisy meta samples. Moreover, we experimentally observe no significant gains for using clean meta samples even for flip noise (where labels are corrupted to a single other class). We extensively experiment the robustness of the MW-Net under the flip and flip2 noise model in Section 4.

4. Experiments

We perform a series of experiments to evaluate the robustness of the weighting network under noisy meta samples and compare our approach with competing methods. We follow the experimental setup in for fair comparison [54].

Dataset We use two benchmark datasets CIFAR-10 and CIFAR-100 for comparing robust learning methods. We use 1000 meta samples for the weighting network. However, contrary to [54], we do not require clean samples for the meta dataset. Thus, we also experiment with corrupted meta samples. In our results, we differentiate between models using clean (denoted with *asterisk**) and noisy meta samples.

Noise Rate We apply the uniform noise model with rates 0, 0.4, and 0.6 and the flip2 noise model with rates 0, 0.2, 0.4. Furthermore, we also compare against conditions under heavily corrupted training samples with a 0.7 uniform label noise rate and a 0.5 flip2 label noise rate. Moreover, we also apply flip noise with a rate of 0.15 and 0.30 and compare with MW-Net³. Note that the flip noise is more challenging than flip2 and uniform noise model; flip/flip2/uniform noise with rate $\eta \geq 0.50/0.67/1.0$ would result in noisy class to be the majority making learning infeasible.

Baseline methods Our analysis shows the weighting network optimized with MAE loss on corrupted meta samples has the same expected gradient direction as of clean meta samples. We denote our model, utilizing corrupted samples and using a robust MAE loss for the weighting network, as

³The flip2 noise model is denoted as flip noise in [54]

Robust-Meta-Noisy-Weight-Network (referred to as RMNW-Net). We denote the MW-Net model utilizing corrupted meta samples as Meta-Noisy-Weight-Network (referred to as MNW-Net); thus, the MNW-Net model trains the weighting network on the noisy meta dataset using cross-entropy loss as the meta loss function. We distinguish models using noisy meta samples and clean meta samples differently. L2RW [49], GLC [25], MW-Net uses clean meta samples and they are denoted as L2RW*, GLC*, and MW-Net* respectively. Other competing methods include BaseModel, trained on corrupted training data, and several robust learning methods such as Reed [48], S-Model [21], SPL [32], Focal Loss [37], Co-teaching [23], D2L [40], MentorNet [30]. We train MW-Net, MNW-Net, RMNW-Net from scratch for all the datasets and the noise models; for the baseline models, we reuse results from [54]. We obtained slightly better performance for MW-Net and its variants with a smaller learning rate than reported in [54].

Network architecture and optimization For uniform noise, we use Wide ResNet-28-10 (WRN-28-10) [68], and for flip2 noise, we use ResNet-32 [24], following [54]. For the flip noise model, we use WRN-28-10 architecture⁴. The weighting network is a single layer neural network with 100 hidden nodes and ReLU activations. We train for MW-Net and its variant for 120 epochs using SGD with momentum 0.9, weight-decay 5×10^{-4} , and an initial learning rate of 0.05. We use an optimizer scheduler to divide the learning rate by 10 after 36 and 38 epoch (for uniform, flip, and flip2 noise with WRN-28-10) and after 40 and 50 epoch (for flip2 noise with ResNet-32), following [54]. The learning rate for the weighting network is fixed to 10^{-3} with weight decay $5e-4$. We do not perform hyper-parameter optimization following [54]. We use a batch size of 100 for both the training samples and the meta samples. We repeat experiments with five different seeds for corrupting samples with label noise and initializing the classifier networks.

4.1. Results

Table 1 lists the average accuracy and standard deviation (std) over five runs for all models on CIFAR-10/100 dataset corrupted with uniform and flip2 noise model. Under 40% and 60% uniform noise, performance of MNW-Net drops by 6% and 12% on CIFAR-10 dataset and drops by 20% and 30% on CIFAR-100 dataset compared to the no noise case. Other baseline models using corrupted meta samples performs worse than MNW-Net. However, using the clean meta samples, performance of MW-Net* only drops by 4% and 9% on CIFAR-10 dataset and drops by 10% and 20% on CIFAR-100 dataset. Only our proposed approach RMNW-Net, even without access to clean meta samples, retains

⁴We also experiment with flip2 noise with WRN-28-10 to show indifference to architecture for MW-Net model and its variants which we detail in supplementary material.

Model	Uniform Noise Rate			Flip2 Noise Rate		
	0%	40%	60%	0%	20%	40%
Dataset CIFAR-10						
L2RW*	92.38 ± 0.10	86.92 ± 0.19	82.24 ± 0.36	89.25 ± 0.37	87.86 ± 0.36	85.66 ± 0.51
GLC*	94.30 ± 0.19	88.28 ± 0.03	83.49 ± 0.24	91.02 ± 0.20	89.68 ± 0.33	88.92 ± 0.24
MW-Net*	95.15 ± 0.13	90.35 ± 0.21	86.52 ± 0.09	92.35 ± 0.29	90.28 ± 0.27	87.04 ± 0.89
BaseModel	95.60 ± 0.22	68.07 ± 1.23	55.12 ± 3.03	92.89 ± 0.32	76.83 ± 2.30	70.77 ± 2.31
Reed-Hard	94.38 ± 0.14	81.26 ± 0.51	73.53 ± 1.54	92.31 ± 0.25	88.28 ± 0.36	81.06 ± 0.76
S-Model	83.79 ± 0.11	79.58 ± 0.33	–	83.61 ± 0.13	79.25 ± 0.30	75.73 ± 0.32
Self-paced	90.81 ± 0.34	86.41 ± 0.29	53.10 ± 1.78	88.52 ± 0.21	87.03 ± 0.34	81.63 ± 0.63
Focal Loss	95.70 ± 0.15	75.96 ± 1.31	51.87 ± 1.19	93.03 ± 0.16	86.45 ± 0.19	80.45 ± 0.97
Co-teaching	88.67 ± 0.25	74.81 ± 0.34	73.06 ± 0.25	89.87 ± 0.10	82.83 ± 0.85	75.41 ± 0.21
D2L	94.64 ± 0.33	85.60 ± 0.13	68.02 ± 0.41	92.02 ± 0.14	87.66 ± 0.40	83.89 ± 0.46
Fine-tuning	95.65 ± 0.15	80.47 ± 0.27	78.75 ± 2.40	93.23 ± 0.23	82.47 ± 3.64	74.07 ± 1.56
MentorNet	94.35 ± 0.42	87.33 ± 0.22	82.80 ± 1.35	92.13 ± 0.30	86.36 ± 0.31	81.76 ± 0.28
MNW-Net	94.98 ± 0.1	88.8 ± 0.44	84.26 ± 0.18	92.45 ± 0.11	90.67 ± 0.18	87.92 ± 0.44
RMNW-Net	95.23 ± 0.13	90.8 ± 0.23	86.31 ± 0.28	92.66 ± 0.13	89.78 ± 0.24	85.38 ± 0.37
Dataset CIFAR-100						
L2RW*	72.99 ± 0.58	60.79 ± 0.91	48.15 ± 0.34	64.11 ± 1.09	57.47 ± 1.16	50.98 ± 1.55
GLC*	73.75 ± 0.51	61.31 ± 0.22	50.81 ± 1.00	65.42 ± 0.23	63.07 ± 0.53	62.22 ± 0.62
MW-Net*	78.06 ± 0.14	70.39 ± 0.16	63.29 ± 0.23	68.92 ± 0.35	64.01 ± 0.37	57.38 ± 0.48
BaseModel	79.95 ± 1.26	51.11 ± 0.42	30.92 ± 0.33	70.50 ± 0.12	50.86 ± 0.27	43.01 ± 1.16
Reed-Hard	64.45 ± 1.02	51.27 ± 1.18	26.95 ± 0.98	69.02 ± 0.32	60.27 ± 0.76	50.40 ± 1.01
S-Model	52.86 ± 0.99	42.12 ± 0.99	–	51.46 ± 0.20	45.45 ± 0.25	43.81 ± 0.15
Self-paced	59.79 ± 0.46	46.31 ± 2.45	19.08 ± 0.57	67.55 ± 0.27	63.63 ± 0.30	53.51 ± 0.53
Focal Loss	81.04 ± 0.24	51.19 ± 0.46	27.70 ± 3.77	70.02 ± 0.53	61.87 ± 0.30	54.13 ± 0.40
Co-teaching	61.80 ± 0.25	46.20 ± 0.15	35.67 ± 1.25	63.31 ± 0.05	54.13 ± 0.55	44.85 ± 0.81
D2L	66.17 ± 1.42	52.10 ± 0.97	41.11 ± 0.30	68.11 ± 0.26	63.48 ± 0.53	51.83 ± 0.33
Fine-tuning	80.88 ± 0.21	52.49 ± 0.74	38.16 ± 0.38	70.72 ± 0.22	56.98 ± 0.50	46.37 ± 0.25
MentorNet	73.26 ± 1.23	61.39 ± 3.99	36.87 ± 1.47	70.24 ± 0.21	61.97 ± 0.47	52.66 ± 0.56
MNW-Net	77.97 ± 0.12	62.95 ± 0.38	55.75 ± 0.56	68.83 ± 0.31	63.69 ± 0.48	56.46 ± 0.28
RMNW-Net	78.27 ± 0.17	70.76 ± 0.20	63.79 ± 0.39	69.69 ± 0.21	65.46 ± 0.51	57.41 ± 0.51

Table 1. Classification accuracy on the clean test set of CIFAR-10/CIFAR-100 dataset under uniform noise and flip2 noise (using WRN-28-10 architecture for uniform noise and ResNet-32 for flip2 noise [54]). L2RW*, GLC*, and MW-Net* use 1000 clean meta samples; all other models use 1000 corrupted meta samples. Best models from each group are **bold**.

Model	CIFAR-10			CIFAR-100		
	Flip Noise Rate			Flip Noise Rate		
	0%	15%	30%	0%	15%	30%
MWNet*	95.15 ± 0.13	93.73 ± 0.2	91.01 ± 0.31	78.06 ± 0.14	73.98 ± 0.32	66.84 ± 0.25
MNW-Net	94.98 ± 0.1	93.75 ± 0.15	92.01 ± 0.23	77.97 ± 0.12	73.79 ± 0.26	66.38 ± 0.51
RMNW-Net	95.23 ± 0.13	93.17 ± 0.12	88.69 ± 0.53	78.27 ± 0.17	74.16 ± 0.25	66.72 ± 0.37

Table 2. Classification accuracy on clean test set of CIFAR-10/CIFAR-100 dataset under flip noise (using WRN-28-10 architecture). Best models from each group are **bold**.

performance similar to MW-Net* on both dataset.

Under flip2 noise, we observe MNW-Net performs slightly better than RMNW-Net on CIFAR-10 dataset while RMNW-Net performs better on CIFAR-100 dataset. In CIFAR-10/CIFAR-100 dataset, MNW-Net/RMNW-Net performs better than all methods using corrupted meta samples. Moreover, we also note that both MNW-Net and RMNW-

Net performs similar to MW-Net without access to the clean meta samples for the flip2 noise model. This observation suggests that even for flip2 noise, we may not need access to clean meta samples. Under clean datasets (0% noise rate), we observe that RMNW-Net performs marginally better than MW-Net* for both WRN-28-10 architecture and ResNet-32 architectures. Improved performance of RMNW-Net com-

Model/ Noise Rate	CIFAR-10						CIFAR-100					
	Uniform		Flip2		Flip		Uniform		Flip2		Flip	
	40%	60%	20%	40%	15%	30%	40%	60%	20%	40%	15%	30%
MW-Net*	<i>0.9714</i>	<i>0.9631</i>	<i>0.9487</i>	<i>0.9647</i>	0.9857	0.9721	<i>0.9644</i>	<i>0.9581</i>	<i>0.8561</i>	<i>0.5962</i>	<i>0.7948</i>	<i>0.5145</i>
MNW-Net	0.9346	0.8874	0.2531	0.5137	0.2692	0.4901	0.9587	0.9390	0.4946	0.5647	0.3407	0.4698
RMNW-Net	0.9848	0.9879	0.9692	0.9694	<i>0.9807</i>	<i>0.9614</i>	0.9749	0.9800	0.9445	0.9001	0.9322	0.7469

Table 3. AUC for detecting noisy training samples. Best models are **bold**, second best models are *italic*.

pared to MW-Net* on clean datasets may suggest MAE loss is suitable for the weighting network for achieving better generalization ability; we leave such studies for future works.

Table 2 lists the average accuracy and std for MW-Net*, MNW-Net, and RMNW-Net for the flip noise model. Similar to flip2 noise, we observe that under flip noise, MNW-Net performs better on CIFAR-10 dataset whereas RMNW-Net performs better on CIFAR-100 dataset. Moreover, we do not observe any significant drop in performance when using corrupted meta samples compared to MW-Net*. Surprisingly, in some cases (30% flip noise on CIFAR-10 dataset), we observe, MNW-Net performs better than MW-Net*; a slightly larger std number might be a reason behind this.

Robustness under heavy label noise Figure 2 shows performances of MW-Net*, MNW-Net, and RMNW-Net under heavy uniform ($\eta = 0.7$) and flip2 ($\eta = 0.5$) noise. We observe that for 70% uniform noise, RMNW-Net performs better than MNW-Net by $\sim 15\%$ and $\sim 1.5\%$ on CIFAR-10 and CIFAR-100 dataset respectively. Under flip2 noise, MNW-Net performs better on CIFAR-10 dataset and RMNW-Net performs better on CIFAR-100 dataset. These observations along with our analytical results suggest RMNW-Net performs the best on when meta samples are corrupted with uniform noise and performs on par with MW-Net* in almost all cases.

Corrupted sample detection We also study whether the weighting network can separate the noisy training samples from the clean training samples. We take the output of the weighting network as the probability of being clean samples and compute the AUC metric on the corrupted training datasets. Table 3 lists the AUC numbers on predicting the clean samples for MW-Net*, MNW-Net, and RMNW-Net for different noise models and noise rates. We run for 120 epochs and test the performances on corrupted training datasets while training; in Table 3, we list the best AUC numbers obtained while training the weighting network and the classifier network. MW-Net* performs better than MNW-Net as expected due to clean meta samples. However, surprisingly, RMNW-Net, with corrupted meta samples, performs better than MW-Net* with clean meta samples. We also observe MW-Net* performs well (in terms of AUC numbers) under uniform noise as also observed in [54]; however, the performance deteriorates significantly for flip2 and flip noise model on CIFAR-100 dataset. Nevertheless, RMNW-Net retains strong performance for all cases. Surprisingly, although RMNW-Net has better performances in terms of

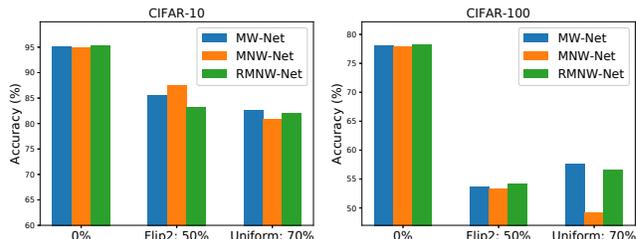


Figure 2. Performance comparison for different models under heavy uniform and flip2 noise (using WRN-28-10 architecture)

AUC numbers, on the primary classification task, we observe MW-Net*, MNW-Net, and RMNW-Net to perform similarly. Further investigations reveal that for RMNW-Net, the importance weights for most samples are in a tighter range whereas, for MW-Net*, the importance weights for most samples spread out over a larger range. A smaller gap in RMNW-Net might be the reason for not having significantly improved classification performance. We note that having a high AUC number on the corrupted sample detection task is not a prerequisite for having the best performance on the primary classification task. However, a high discrimination ability to separate the corrupted samples is indeed important for devising a more powerful robust model. We suspect that RMNW-Net can filter out corrupt samples from the noisy training datasets using a small number of corrupted meta samples with high accuracy, though we leave confirmation to future work.

5. Related Works

Machine learning models are highly vulnerable to label noise in the training datasets; label noise becomes more important for the DNN classifier with enormous memorization capability [3]. Commonly used convex potential losses [5] are not tolerant to label noise; theoretically, noise robustness of convex loss potentials have been analyzed in [53] and experimentally studied in [43]. A class of common approaches looks for models that are inherently tolerant to label noise; thus training with noisy labels should not affect the predictive performance on unbiased test datasets [57, 18, 19]. For asymmetric label noise, an unbiased estimator can be constructed if the noise rates are known [42]. Computing the noise rates from corrupted datasets has been also studied under some strict assumptions [51]. For binary classifiers, a sufficient condition has been derived for a loss functions to be tolerant to label noise in [18]; ramp loss, probit loss, unhinged loss satisfies this condition [18, 57, 4].

For a multi-class DNN classifier, MAE loss is proved to be tolerant to label noise [17]. A generalized cross-entropy loss has been proposed that offers noise robustness with easier optimization procedure [70]. Along the line of robust MAE loss [17], several new loss functions have been recently proposed [39, 38, 63, 60, 10, 45, 14, 12]. Several methods have also focused on adding a noise transition matrix layer before computing the softmax probability distributions [21, 46, 25, 65, 62, 64]. Among these, several methods additionally use a subset of clean gold samples for computing the transition matrix with some guarantees [25, 46, 62]. The core idea of the noise transition layer is to perform Expectation-Minimization (EM) algorithm to iteratively predict the noisy samples and learn the classifier networks [33].

Another prominent approach for learning with label noise is to determine the clean samples in the corrupted training datasets. EM algorithms for label noise also fall under these settings where the expectation steps learn the probable clean samples. The idea of choosing probable clean samples goes farther back; many heuristics have been proposed to determine the clean samples in the training procedure and remove the effect of those samples in earlier research [8, 7]. Recently, there has been a resurgence of methods using ideas along these lines [21, 41, 23, 56, 11, 23, 30, 56, 58, 67, 37, 48, 22]. A similar concept is introduced in curriculum learning where model learns from easy samples first, and then adapts to harder examples [6, 29, 11, 36]. In Self-Paced learning (SPL), the model first optimizes easy samples (probably clean) samples [32, 28]. In Mentornet [30], the model consists of a student-teacher network where the teacher finds clean samples for the students. Co-teaching is another method with two neural networks where one network uses samples predicted as clean (with some threshold) using the other network [23]. Along this line, many recent methods for robust DNN learning use an ensemble (two or more networks) approach to mitigate the effect of label noise [30, 67, 23, 66, 9]. Although one network can be susceptible to noise, two (or more) networks can be possibly used to remove noisy samples. Some methods also augment an unbiased clean dataset; in [67], authors use a different subset of the noisy dataset for multiple networks with the help of clean datasets whereas in [58], authors use some number of clean samples to learn a separate clean sample prediction networks. Susceptibility of DNN under label noise is also due to memorization and overfitting capabilities. Based on the overfitting behavior, in [40], the authors proposed a dimensionality driven measure to determine the epochs where overfitting starts to occur. Along a similar line to [40], in [27], the learning rate is varied between overfitting and underfitting in the training procedure. However, while these methods perform well in practice, they do not guarantee robustness and often require hyper-parameter tuning and closely watching training losses.

Methods, identifying clean samples and retraining the classifiers, effectively reduce the importance weights for probable noisy samples. Thus, many proposed approaches learn the importance weights with or without the help of additional clean samples [49, 37, 35, 54, 34, 26, 59]. The core idea is to assign the importance weights based on loss values; for example, in Focal loss, the model gives more weight to high loss samples [37]. In [48], the authors propose a regularized method to use samples that have labels similar to the model predictions; thus, the model implicitly gives low weights to high loss samples. Many recent methods propose to learn the importance weight function by non-linear functions [49, 54]. Meta-learning methods are suitable for learning the weighting network to learn the classifier network [13, 44, 2]. The meta-learning framework learns a network that can adapt to many tasks with very few samples [55, 44, 47]; model agnostic meta-learning, a powerful approach for meta-learning, optimizes an optimal initializer for many tasks [13]. In [49], the authors use meta-learning methods L2R for learning to weight samples; similar to MW-Net, the model learns a weighting function using clean meta (validation) datasets. Along this line, several proposed methods [61, 54, 35, 69, 31, 52] use meta-learning frameworks.

6. Discussions and Future Works

An optimal sample weighting strategy is of utmost importance for learning robust deep neural networks; learning a weighting strategy alleviates the need for heuristics to assign weights to corrupted samples in the training datasets. However, learning the weighting network comes at a cost; previous research uses gold standard clean datasets for the weighting network. In this paper, we close the gap for learning the weighting network without access to clean meta samples. Theoretical results and extensive experimentation suggest that using robust loss functions in the weighting network can perform similar to the case with clean meta samples. Although MAE loss is difficult to optimize for deep networks, optimizing the weighting (shallow) network with MAE loss is easy and effective. Albeit simple, we show that it is possible to learn robust classifiers by importance weighting without access to clean meta samples.

Interestingly, we observe the weight network with robust losses has much more discernibility than the cross-entropy loss in separating the corrupted samples from the training datasets. Identifying the clean samples (instead of assigning importance weights) using the meta-network with robust losses is an intriguing avenue for future work. Although our proposed approach performs similar to MW-Net even without access to the clean meta samples for all the noise models, our theoretical results only hold for the uniform noise model on the meta dataset. Thus, the degree of robustness for MW-Net and RMNW-Net under the flip and flip2 noise model remains an open problem to tackle.

References

- [1] Görkem Algan and Ilkay Ulusoy. Image classification with deep learning in the presence of noisy labels: A survey. *arXiv preprint arXiv:1912.05170*, 2019.
- [2] Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. Learning to learn by gradient descent by gradient descent. In *NeurIPS*, pages 3981–3989, 2016.
- [3] Devansh Arpit, Stanisław Jastrzembowski, Nicolas Balas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, and Simon Lacoste-Julien. A closer look at memorization in deep networks. In *ICML*, page 233–242, 2017.
- [4] Han Bao, Clayton Scott, and Masashi Sugiyama. Calibrated surrogate losses for adversarially robust classification. In *COLT*, pages 408–451, 2020.
- [5] Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [6] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *ICML*, pages 41–48, 2009.
- [7] Carla E Brodley and Mark A Friedl. Identifying mislabeled training data. *Journal of artificial intelligence research*, 11:131–167, 1999.
- [8] Carla E Brodley, Mark A Friedl, et al. Identifying and eliminating mislabeled training instances. In *AAAI*, pages 799–805, 1996.
- [9] Haw-Shiuan Chang, Erik Learned-Miller, and Andrew McCallum. Active bias: Training more accurate neural networks by emphasizing high variance samples. In *NeurIPS*, pages 1002–1012, 2017.
- [10] Nontawat Charoenphakdee, Jongyeong Lee, and Masashi Sugiyama. On symmetric losses for learning from corrupted labels. In *ICML*, pages 961–970, 2019.
- [11] Xinlei Chen and Abhinav Gupta. Webly supervised learning of convolutional networks. In *ICCV*, pages 1431–1439, 2015.
- [12] Lei Feng, Senlin Shu, Zhuoyi Lin, Fengmao Lv, Li Li, and Bo An. Can cross entropy loss be robust to label noise? pages 2206–2212. *IJCAI*, 2020.
- [13] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, pages 1126–1135, 2017.
- [14] Eduardo Fonseca, Manoj Plakal, Daniel PW Ellis, Frederic Font, Xavier Favory, and Xavier Serra. Learning sound event classifiers from web audio with noisy labels. In *ICASSP*, pages 21–25. IEEE, 2019.
- [15] Luca Franceschi, Paolo Frasconi, Saverio Salzo, Riccardo Grazi, and Massimiliano Pontil. Bilevel programming for hyperparameter optimization and meta-learning. In *ICML*, pages 1568–1577, 2018.
- [16] Benoît Frénay and Michel Verleysen. Classification in the presence of label noise: a survey. *IEEE transactions on neural networks and learning systems*, 25(5):845–869, 2013.
- [17] Aritra Ghosh, Himanshu Kumar, and PS Sastry. Robust loss functions under label noise for deep neural networks. In *AAAI*, pages 1919–1925, 2017.
- [18] Aritra Ghosh, Naresh Manwani, and PS Sastry. Making risk minimization tolerant to label noise. *Neurocomputing*, 160:93–107, 2015.
- [19] Aritra Ghosh, Naresh Manwani, and PS Sastry. On the robustness of decision tree learning under label noise. In *PAKDD*, pages 685–697. Springer, Cham, 2017.
- [20] Github. Awesome-learning-with-label-noise. <https://github.com/subeeshvasu/Awesome-Learning-with-Label-Noise>, 2020.
- [21] Jacob Goldberger and Ehud Ben-Reuven. Training deep neural-networks using a noise adaptation layer. In *ICLR*, 2017.
- [22] Bo Han, Jiangchao Yao, Gang Niu, Mingyuan Zhou, Ivor Tsang, Ya Zhang, and Masashi Sugiyama. Masking: A new perspective of noisy supervision. In *NeurIPS*, pages 5836–5846, 2018.
- [23] 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. In *NeurIPS*, pages 8527–8537, 2018.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [25] Dan Hendrycks, Mantas Mazeika, Duncan Wilson, and Kevin Gimpel. Using trusted data to train deep networks on labels corrupted by severe noise. In *NeurIPS*, pages 10456–10465, 2018.
- [26] Wei Hu, Yangyu Huang, Fan Zhang, and Ruirui Li. Noise-tolerant paradigm for training face recognition cnns. In *CVPR*, pages 11887–11896, 2019.
- [27] Jinchuan Huang, Lie Qu, Rongfei Jia, and Binqiang Zhao. O2u-net: A simple noisy label detection approach for deep neural networks. In *ICCV*, pages 3326–3334, 2019.
- [28] Lu Jiang, Deyu Meng, Shou-I Yu, Zhenzhong Lan, Shiguang Shan, and Alexander Hauptmann. Self-paced learning with diversity. In *NeurIPS*, pages 2078–2086, 2014.
- [29] Lu Jiang, Deyu Meng, Qian Zhao, Shiguang Shan, and Alexander G Hauptmann. Self-paced curriculum learning. In *AAAI*, page 6, 2015.
- [30] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, pages 2304–2313, 2018.
- [31] Youngdong Kim, Junho Yim, Juseung Yun, and Junmo Kim. Nlnl: Negative learning for noisy labels. In *ICCV*, pages 101–110, 2019.
- [32] M Pawan Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. In *NeurIPS*, pages 1189–1197, 2010.
- [33] N Lawrence and Bernhard Schölkopf. Estimating a kernel fisher discriminant in the presence of label noise. In *ICML*, pages 306–306, 2001.
- [34] Kuang-Huei Lee, Xiaodong He, Lei Zhang, and Linjun Yang. Cleannet: Transfer learning for scalable image classifier training with label noise. In *CVPR*, pages 5447–5456, 2018.

- [35] Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S Kankanhalli. Learning to learn from noisy labeled data. In *CVPR*, pages 5051–5059, 2019.
- [36] Yuncheng Li, Jianchao Yang, Yale Song, Liangliang Cao, Jiebo Luo, and Li-Jia Li. Learning from noisy labels with distillation. In *ICCV*, pages 1910–1918, 2017.
- [37] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017.
- [38] Yang Liu and Hongyi Guo. Peer loss functions: Learning from noisy labels without knowing noise rates. In *ICML*, 2020.
- [39] Xingjun Ma, Hanxun Huang, Yisen Wang, Simone Romano, Sarah Erfani, and James Bailey. Normalized loss functions for deep learning with noisy labels. In *ICML*, 2020.
- [40] Xingjun Ma, Yisen Wang, Michael E Houle, Shuo Zhou, Sarah Erfani, Shutao Xia, Sudanthi Wijewickrema, and James Bailey. Dimensionality-driven learning with noisy labels. In *ICML*, pages 3355–3364, 2018.
- [41] Devraj Mandal, Shrisha Bharadwaj, and Soma Biswas. A novel self-supervised re-labeling approach for training with noisy labels. In *WACV*, pages 1381–1390, March 2020.
- [42] Nagarajan Natarajan, Inderjit S Dhillon, Pradeep K Ravikumar, and Ambuj Tewari. Learning with noisy labels. In *NeurIPS*, pages 1196–1204, 2013.
- [43] David F Nettleton, Albert Orriols-Puig, and Albert Fornells. A study of the effect of different types of noise on the precision of supervised learning techniques. *Artificial intelligence review*, 33(4):275–306, 2010.
- [44] Alex Nichol and John Schulman. Reptile: a scalable meta-learning algorithm. *arXiv preprint arXiv:1803.02999*, 2(3):4, 2018.
- [45] Giorgio Patrini, Frank Nielsen, Richard Nock, and Marcello Carioni. Loss factorization, weakly supervised learning and label noise robustness. In *ICML*, pages 708–717, 2016.
- [46] 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 *CVPR*, pages 1944–1952, 2017.
- [47] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017.
- [48] Scott Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. Training deep neural networks on noisy labels with bootstrapping. *ICLR Workshops*, 2015.
- [49] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *ICML*, pages 4334–4343, 2018.
- [50] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- [51] Clayton Scott, Gilles Blanchard, and Gregory Handy. Classification with asymmetric label noise: Consistency and maximal denoising. In *COLT*, pages 489–511, 2013.
- [52] Paul Hongsuck Seo, Geeho Kim, and Bohyung Han. Combinatorial inference against label noise. In *NeurIPS*, pages 1173–1183, 2019.
- [53] Rocco A Servedio. Smooth boosting and learning with malicious noise. *Journal of Machine Learning Research*, 4(Sep):633–648, 2003.
- [54] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. In *NeurIPS*, pages 1919–1930, 2019.
- [55] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, pages 4077–4087, 2017.
- [56] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *CVPR*, pages 5552–5560, 2018.
- [57] Brendan Van Rooyen, Aditya Menon, and Robert C Williamson. Learning with symmetric label noise: The importance of being unhinged. In *NeurIPS*, pages 10–18, 2015.
- [58] Andreas Veit, Neil Alldrin, Gal Chechik, Ivan Krasin, Abhinav Gupta, and Serge Belongie. Learning from noisy large-scale datasets with minimal supervision. In *CVPR*, pages 839–847, 2017.
- [59] Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, and Shu-Tao Xia. Iterative learning with open-set noisy labels. In *CVPR*, pages 8688–8696, 2018.
- [60] Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *ICCV*, pages 322–330, 2019.
- [61] Zhen Wang, Guosheng Hu, and Qinghua Hu. Training noise-robust deep neural networks via meta-learning. In *CVPR*, pages 4524–4533, 2020.
- [62] Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *CVPR*, pages 2691–2699, 2015.
- [63] Yilun Xu, Peng Cao, Yuqing Kong, and Yizhou Wang. L_{dmi}: An information-theoretic noise-robust loss function. 32:6225–6236, 2019.
- [64] Jiangchao Yao, Hao Wu, Ya Zhang, Ivor W Tsang, and Jun Sun. Safeguarded dynamic label regression for noisy supervision. In *AAAI*, volume 33, pages 9103–9110, 2019.
- [65] Kun Yi and Jianxin Wu. Probabilistic end-to-end noise correction for learning with noisy labels. In *CVPR*, pages 7017–7025, 2019.
- [66] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *ICML*, pages 7164–7173, 2019.
- [67] Bodi Yuan, Jianyu Chen, Weidong Zhang, Hung-Shuo Tai, and Sara McMains. Iterative cross learning on noisy labels. In *WACV*, pages 757–765. IEEE, 2018.
- [68] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016.
- [69] Weihe Zhang, Yali Wang, and Yu Qiao. Metacleaner: Learning to hallucinate clean representations for noisy-labeled visual recognition. In *CVPR*, pages 7373–7382, 2019.
- [70] Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. In *NeurIPS*, pages 8778–8788, 2018.