

# Optimized Gradient Clipping for Noisy Label Learning

Xichen Ye<sup>1</sup>, Yifan Wu<sup>1,2</sup>, Weizhong Zhang<sup>2,4</sup>, Xiaoqiang Li<sup>1\*</sup>, Yifan Chen<sup>3\*</sup>, Cheng Jin<sup>2,5</sup>

<sup>1</sup>Shanghai University

<sup>2</sup>Fudan University

<sup>3</sup>Hong Kong Baptist University

<sup>4</sup>Shanghai Key Laboratory of Intelligent Information Processing

<sup>5</sup>Shanghai Collaborative Innovation Center of Intelligent Visual Computing

{yexichen0930, victorwu, xqli}@shu.edu.cn, {weizhongzhang, jc}@fudan.edu.cn, yifanc@hkbu.edu.hk

## Abstract

Previous research has shown that constraining the gradient of loss function w.r.t. model-predicted probabilities can enhance the model robustness against noisy labels. These methods typically specify a fixed optimal threshold for gradient clipping through validation data to obtain the desired robustness against noise. However, this common practice overlooks the dynamic distribution of gradients from both clean and noisy-labeled samples at different stages of training, significantly limiting the model capability to adapt to the variable nature of gradients throughout the training process. To address this issue, we propose a simple yet effective approach called Optimized Gradient Clipping (OGC), which dynamically adjusts the clipping threshold based on the ratio of noise gradients to clean gradients after clipping, estimated by modeling the distributions of clean and noisy samples. This approach allows us to modify the clipping threshold at each training step, effectively controlling the influence of noise gradients. Additionally, we provide statistical analysis to certify the noise-tolerance ability of OGC. Our extensive experiments across various types of label noise, including symmetric, asymmetric, instance-dependent, and real-world noise, demonstrate the effectiveness of our approach.

**Code** — <https://github.com/Virusdoll/OGC>

## 1 Introduction

The effectiveness of supervised deep learning relies heavily on the availability of large-scale, qualifiedly-annotated data. Research has shown that an over-parameterized Deep Neural Network (DNN) can easily fit a dataset with randomly assigned labels (Zhang et al. 2017), underscoring the importance of high-quality annotations. However, in real-world scenarios, the annotation process will inevitably introduce noisy labels (incorrectly annotated labels), which is liable to compromise the performance of the model. Therefore, the study of noisy label learning has attracted significant attention. The goal of noisy label learning is to train a model on a given corrupted dataset, which contains mislabeled samples, while mitigating the adverse effects of noisy labels, and ultimately enabling the trained model to generalize well on a clean evaluation set.

\*Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Cross Entropy (CE), the most widely used loss function for classification, is nevertheless vulnerable to label noise. In contrast, symmetric loss functions like Mean Absolute Error (MAE) are noise tolerant but less effective for training classification models (Ghosh, Kumar, and Sastry 2017; Zhang and Sabuncu 2018). This discrepancy has driven the studies into robust loss functions, which primarily focus on how to make full use of both CE and MAE (Zhang and Sabuncu 2018; Feng et al. 2020; Engleson and Azizpour 2021; Wang et al. 2019; Ma et al. 2020; Ye et al. 2023) and have achieved notable success. Besides these, another promising approach, clipping-based techniques, also warrants greater attention. Among them, Menon et al. (2020) first introduced the PHuberCE method, which refines traditional gradient clipping by applying gradient clipping to model-predicted probabilities (see the definition in Section 3.1), offering an improved approach to handling noisy labels. More recently, Wei et al. (2023) proposed LogitClip, which involves clipping the logits to indirectly constrain the upper and lower bounds of the gradients.

Specifically, the two aforementioned clipping-based methods rely on the validation set to determine a fixed optimal clipping threshold, overlooking the dynamic nature of gradients during the training process. However, as demonstrated by the simple binary classification experiment in Figure 1, we observe distinct shifts in the gradient distribution throughout training, which significantly affect the suitability of a fixed threshold. Considering the fact that employing a higher threshold can enhance the model’s fitting ability, yet fails to bound the decreasing noisy gradients; conversely, employing a strict threshold ensures the noisy gradients remain bounded but compromises the model’s fitting ability, we expect a dynamically adjustable threshold for gradient clipping that continuously adapts in response to the shifts in the gradient distribution.

Motivated by this, in this paper, we propose an optimization-based strategy, Optimized Gradient Clipping (OGC), which ensures that the threshold adjusts itself as the distribution of gradients evolves during training, maintaining an optimal balance between fitting ability and noise suppression. Specifically, OGC 1) utilizes a 2-component Gaussian Mixture Model (2-GMM) to model the distribution of cross-entropy losses for clean and noisy samples, 2) and estimates the ratio of the noisy gradients to the clean gradi-

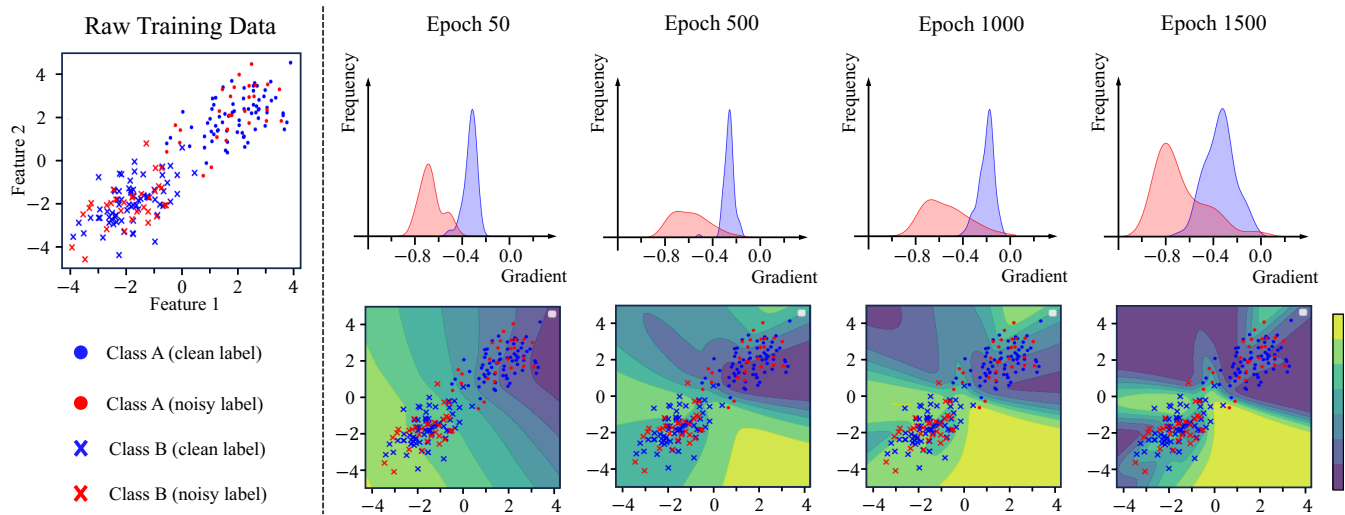


Figure 1: KDE visualizations of gradient distributions for clean and noisy labels, along with decision boundary visualizations for a simple binary classification task that utilizes gradient clipping with a fixed threshold. The leftmost plot in the first row shows the raw training data. The subsequent plots illustrate shifts in gradient distributions (before clipping) at various training epochs (50, 500, 1000, and 1500). The second row displays the corresponding decision boundaries for each epoch.

ents after clipping. This allows us to determine a clipping threshold at each training step that restricts the gradient ratio to a predefined limit, effectively bounding the influence of noisy gradients. Theoretically, we demonstrate that the dynamic threshold obtained by OGC can effectively enhance the noise-tolerant ability of CE under either symmetric, asymmetric, or instance-dependent noise, which is originally a non-robust loss function. Furthermore, we conduct extensive experiments on various types of label noise, including, symmetric, asymmetric, instance-dependent, and real-world noise, to verify the effectiveness of our proposed method. Our key contributions are summarized as follows:

- We first recognize the crucial role that gradient distribution shifts play in noisy label learning and propose a simple yet effective strategy to obtain a dynamic threshold that adapts to these continual shifts.
- We provide a theoretical analysis demonstrating that the threshold obtained by OGC can make non-robust loss functions robust against various types of label noise.
- We conduct extensive experiments to evaluate the effectiveness of our proposed method across multiple types of label noise, including symmetric, asymmetric, instance-dependent, and real-world noise.

## 2 Related Works

**Robust loss functions.** In recent years, several robust-loss-based methods have been developed for learning with noisy labels. Ghosh, Kumar, and Sastry (2017) theoretically demonstrated the robustness of symmetric loss functions, such as Mean Absolute Error (MAE), to label noise. Building on this, Zhang and Sabuncu (2018) introduced the Generalized Cross Entropy (GCE), which is a generalization of Cross Entropy (CE) and MAE. Further, Wang et al.

(2019) combined CE with scaled MAE to propose Symmetric Cross Entropy (SCE). Additionally, Menon et al. (2020) introduced a composite loss-based gradient clipping applied to CE, resulting in PHuber-CE. Feng et al. (2020) utilized the Taylor series to derive an alternative representation of CE, leading to the development of Taylor-CE. Moreover, Ma et al. (2020) proposed the Active Passive Loss (APL), aiming to create fully robust loss functions. Zhou et al. (2021) developed Asymmetric Loss Functions (ALFs) to address limitations posed by symmetric conditions. Wei et al. (2023) introduced logit clipping (LogitClip), a technique that clamps the norm of the logit vector to a constant upper bound. Recently, Ye et al. (2023) proposed Normalized Negative Loss Functions (NNLFs), which prioritize learning from well-understood samples.

**Other methods for noisy label learning.** In addition to robust loss functions, various other approaches are utilized for learning with noisy labels. A sample selection-based approach aims to identify and select clean samples for training, often leveraging the disagreement between two models (Yu et al. 2019; Wei et al. 2020). Some research further combines these methods with techniques from unsupervised or semi-supervised learning to make full use of the unselected corrupted samples (Nguyen et al. 2020; Li, Socher, and Hoi 2020; Wei et al. 2022b). A loss correction-based approach modifies the loss of each sample based on label-dependent weights (Natarajan et al. 2013), an estimated noise transition matrix (Han et al. 2018; Patrini et al. 2017), or model predictions (Arazo et al. 2019). Additionally, a regularization-based approach is also employed to combat label noise, including methods such as MixUp (Zhang et al. 2018), label smoothing (Lukasik et al. 2020), or leveraging previous prediction results (Liu et al. 2020).

### 3 Preliminaries

#### 3.1 Risk Minimization and Noisy Labels

In this work, we consider a typical classification problem with  $K$ -categories. Let  $\mathcal{X} \subset \mathbb{R}^d$  be the sample space and  $\mathcal{Y} = [K] := \{1, \dots, K\}$  be the label space. Given a dataset  $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$ , where each sample pair  $(\mathbf{x}_n, y_n)$  is drawn *i.i.d.* from an underlying “clean” distribution  $\mathcal{D}_c$ , over  $\mathcal{X} \times \mathcal{Y}$ . We define  $\mathbf{q}(k|\mathbf{x}_n)$  as the ground truth probabilistic distribution across various labels for a sample  $\mathbf{x}_n$ , along with the condition that  $\sum_{k=1}^K \mathbf{q}(k|\mathbf{x}_n) = 1$ . Specifically,  $\mathbf{q}(k|\mathbf{x}_n)$  is a degenerate distribution where  $\mathbf{q}(y_n|\mathbf{x}_n) = 1$  and  $\mathbf{q}(k \neq y_n|\mathbf{x}_n) = 0$ .

We define a classifier  $f: \mathcal{X} \rightarrow \mathbb{R}^K$  as a function that maps from sample space to logit space, parameterized by trainable parameters  $\theta$ . Specifically, we focus on the case that  $f$  is a deep neural network (DNN). Given a sample  $\mathbf{x}_n$ , the function  $f(\mathbf{x}_n)$  computes the logits. By applying a softmax output layer, we obtain the **model-predicted probability**  $\mathbf{p}(k|\mathbf{x}_n)$  for each category  $k \in \{1, \dots, K\}$ , with  $\sum_{k=1}^K \mathbf{p}(k|\mathbf{x}_n) = 1$ . Training a classifier  $f$  is to find an optimal classifier  $f^*$  that minimize the empirical risk:  $\sum_{n=1}^N \ell(f(\mathbf{x}_n), y_n)$ , where  $\ell: \mathbb{R}^K \times \mathcal{Y} \rightarrow \mathbb{R}^+$  denotes the loss function.

When label noise exists, we work with a corrupted dataset  $\{(\mathbf{x}_n, \tilde{y}_n)\}_{n=1}^N$ , where each sample is drawn *i.i.d.* from an unknown distribution  $\mathcal{D}_\eta$  and  $\tilde{y}$  denotes the possibly incorrect label. There are two ways to model label noise: instance-independent or instance-dependent.

In the **instance-independent** approach, it is assumed that, given the true label  $y$ , the corruption process is conditionally independent of input sample  $\mathbf{x}$  (Ghosh, Kumar, and Sastry 2017). So the stochastic corruption that the true label  $y$  is corrupted into label  $j$  is specified by a transition matrix,

$$[\eta_{yj}] = q_\eta(j|y), \quad \forall y, j \in [K], \quad (1)$$

where  $\eta_{yj}$  denotes the probability that true label  $y$  is corrupted into label  $j$ , and  $\eta_y = \sum_{j \neq y} \eta_{yj}$  denotes the noise rate for label  $y$ . Under this model assumption, label noise can be either symmetric or asymmetric. For symmetric noise,  $\eta_{ij} = \frac{\eta_i}{K-1}, \forall j \neq i$  and  $\eta_i = \eta, \forall i \in [k]$ , where  $\eta$  is a constant. For asymmetric noise,  $\eta_{ij}$  varies depending on both the true label  $i$  and corrupted label  $j$ .

The **instance-dependent** approach extends the setting above and assumes that the corruption process is conditionally dependent on both the true label  $y$  and input sample  $\mathbf{x}$  (Xia et al. 2020). Similarly, the probability of the event that true label  $y$  is corrupted into label  $j$  is given by:

$$\eta_{yj}(\mathbf{x}) = q_\eta(j|y, \mathbf{x}), \quad (2)$$

where  $\eta_{yj}(\mathbf{x})$  now depends on  $\mathbf{x}$ , and  $\eta_y(\mathbf{x}) = \sum_{j \neq y} \eta_{yj}(\mathbf{x})$  denotes the noise rate for sample  $\mathbf{x}$ .

As a closing remark, the unknown “noisy” distribution  $\mathcal{D}_\eta$  can be taken as a mixture of the “clean” distribution  $\mathcal{D}_c$  and the “noise” distribution  $\mathcal{D}_n$ . Specifically,  $q_\eta(\mathbf{x}, \tilde{y}) = q_\eta(\tilde{y} = y) \cdot q_\eta(\mathbf{x}, \tilde{y} | \tilde{y} = y) + q_\eta(\tilde{y} \neq y) \cdot q_\eta(\mathbf{x}, \tilde{y} | \tilde{y} \neq y)$ ;  $q_\eta(\mathbf{x}, \tilde{y} | \tilde{y} = y)$  is exactly the previously introduced “clean” distribution  $\mathcal{D}_c$ , and we further denote  $\mathcal{D}_n := q_\eta(\mathbf{x}, \tilde{y} | \tilde{y} \neq y)$ .

#### 3.2 Gradient Clipping

We first introduce the vector clipping function. Given a user-specified threshold  $\tau > 0$ , the clipping function can be formally defined by:

$$\text{clip}(w, \tau) = \begin{cases} \tau \cdot \frac{w}{\|w\|_2}, & \text{if } \|w\|_2 \geq \tau, \\ w, & \text{otherwise,} \end{cases} \quad (3)$$

where  $w$  denotes the vector to be clipped, and the output is guaranteed to have an  $\ell_2$  norm upper bounded by  $\tau$ .

In practice, the clipping function is usually applied to a mini-batch  $\{(x_m, y_m)\}_{m=1}^M$ . To clip the gradients for the model parameters  $\theta$ , the operation is defined as follows:

$$\begin{aligned} \bar{g}_\theta &= \text{clip}(g_\theta, \tau), \\ \text{where } g_\theta &= \frac{1}{M} \sum_{m=1}^M \nabla_\theta \ell(f(x_m), y_m). \end{aligned} \quad (4)$$

Standard gradient clipping is typically employed to avoid gradient explosion (Pascanu, Mikolov, and Bengio 2013) and accelerate the convergence of the model (Zhang et al. 2020).

As an implicit regularization approach, gradient clipping is adopted by practitioners to alleviate the issues caused by noisy labels. However, Menon et al. (2020) suggested that merely clipping the gradient of the model parameter  $\theta$  does not effectively counteract the detrimental impacts of noisy labels. They discovered that, on a linear model, applying gradient clipping to parameters is equivalent to applying it to logits. However, clipping gradients on logits modifies the loss function in a way that allows outliers to still influence the gradients, ultimately compromising the robustness. Instead, they recommend applying gradient clipping directly to the gradients of the model predicted probabilities  $\mathbf{p}$  in the following manner:

$$\bar{g}_\mathbf{p} = \text{clip}(g_\mathbf{p}, \tau), \quad \text{where } g_\mathbf{p} = \nabla_\mathbf{p} \ell(f(\mathbf{x}), y). \quad (5)$$

Applying gradient clipping to probabilities results in a modified loss function where the loss value remains nearly constant when logits exceed a certain threshold. This property helps the loss function assign small gradients to outlier samples (which are likely to be corrupted). Consequently, it bounds the gradient caused by noisy samples, enhancing the robustness of the loss function against noisy labels.

### 4 Optimized Gradient Clipping

We present the methodology of optimized gradient clipping in this section.

#### 4.1 Motivation

Given the observation that a fixed clipping threshold may be ineffective in addressing the inherent variability of gradient distributions, we reconsider the problem of selecting the clipping threshold from a new perspective.

Specifically, we are interested in understanding how the clipping threshold  $\tau^{(t)}$  affects  $r^{(t)}$ , which represents the ratio of expected gradients between clean and noisy samples:

$$r^{(t)} = \frac{\mathbb{E}_{(\mathbf{x}, \tilde{y}) \sim \mathcal{D}_n} [\|\text{clip}(\nabla_\mathbf{p} \ell(f(\mathbf{x}), \tilde{y}), \tau^{(t)})\|_2]}{\mathbb{E}_{(\mathbf{x}, \tilde{y}) \sim \mathcal{D}_c} [\|\text{clip}(\nabla_\mathbf{p} \ell(f(\mathbf{x}), \tilde{y}), \tau^{(t)})\|_2]}. \quad (6)$$

This formula calculates the ratio of expected gradient (with respect to the predicted probabilities) between noise and clean samples after clipping at time step  $t$ . Intuitively, this ratio reflects the extent to which noisy samples influence the training process. As the ratio increases, the influence of noisy samples begins to dominate, leading to model overfitting. Conversely, by controlling this ratio through the optimized threshold  $\tau^{(t)}$ , we can effectively manage the impact of noisy samples.

## 4.2 Method

Here, we formally introduce our method, named as Optimized Gradient Clipping (OGC). Our approach consists of the following three components: 1) Modeling clean and noisy distributions, 2) Determining the clipping threshold, and 3) Applying clipping to the loss functions. Our algorithm is outlined in Appendix.

To simplify our discussion, we focus on a specific type of loss function under the following assumption: the loss value  $\ell$  depends exclusively on  $p(y|x)$ , the predicted probability of the given label. It should be noted that most widely used loss functions follow this assumption. This includes both non-robust loss functions, such as Cross Entropy (CE) and Focal Loss (FL), and robust loss functions, such as Generalized Cross Entropy (GCE).

This assumption enables us to define two key mapping functions: 1)  $\phi_{H \rightarrow \ell}$ , which directly maps the cross entropy value  $H$  to the loss value  $\ell$ ; and 2)  $\phi_{\|g_p\|_2 \rightarrow \ell}$ , which directly maps  $\|g_p\|_2$ , the gradient norm of the predicted probability, to the loss value  $\ell$ .

**Modeling of clean and noisy distributions.** In the real world, we only have access to a corrupted dataset sampled from an unknown distribution  $D_\eta$ , making the direct calculation of  $r^{(t)}$  unfeasible. To address this, we propose a proxy ratio  $\tilde{r}^{(t)}$  to approximate the ground-truth  $r^{(t)}$ .

Given a mini-batch  $\{(\mathbf{x}_m, \tilde{y}_m)\}_{m=1}^M$  sampled at time  $t$ , we recall the cross entropy for a sample pair  $(\mathbf{x}, \tilde{y})$  is given as follows:

$$\begin{aligned} H(f(\mathbf{x}), \tilde{y}) &= - \sum_{k=1}^K q(k|\mathbf{x}) \log p(k|\mathbf{x}) \\ &= - \log p(\tilde{y}|\mathbf{x}). \end{aligned} \quad (7)$$

By respectively assuming  $(\mathbf{x}, \tilde{y}) \sim \mathcal{D}_c$  and  $(\mathbf{x}, \tilde{y}) \sim \mathcal{D}_n$ ,  $H(f(\mathbf{x}), \tilde{y})$  will exhibit two distinct distributions, which accordingly induce two random variables  $H_c, H_n$ .

The two-component Gaussian Mixture Model (2-GMM), which can be considered a non-parametric method, is commonly utilized to model the Cross Entropy value  $H$  in noisy label learning scenarios (Li, Socher, and Hoi 2020). This model enabling us to effectively characterize the clean and noisy distributions. Here, we fit a 2-GMM on the empirical distribution of  $H(f(\mathbf{x}), \tilde{y})$  for  $(\mathbf{x}, \tilde{y}) \sim \mathcal{D}_\eta$ , resulting in the following two Gaussian approximation distributions:

$$H_c \sim \mathcal{N}_c^{(t)}, \quad \text{where } \mathcal{N}_c^{(t)} = \mathcal{N}(\mu_c^{(t)}, \sigma_c^{(t)2}) \quad (8)$$

$$H_n \sim \mathcal{N}_n^{(t)}, \quad \text{where } \mathcal{N}_n^{(t)} = \mathcal{N}(\mu_n^{(t)}, \sigma_n^{(t)2}), \quad (9)$$

where  $\mu_c^{(t)} < \mu_n^{(t)}$ ,  $\mathcal{N}_c^{(t)}$  and  $\mathcal{N}_n^{(t)}$  denotes the clean and noise gaussian distribution at time  $t$ , respectively. It should be noted that the assumption that clean samples exhibit lower cross entropy values than noisy samples is widely accepted in the field of noisy label learning (Gui, Wang, and Tian 2021).

**Determine clipping threshold  $\tau^{(t)}$ .** Given the clean and noise distributions, we can calculate  $\tilde{r}^{(t)}$  in the following manner ( $\phi_{H \rightarrow \ell}(\cdot)$  is defined in the third paragraph of this subsection):

$$\tilde{r}^{(t)} = \frac{\int_0^{+\infty} \|\text{clip}_{\mathbb{P}}(\nabla_p \phi_{H \rightarrow \ell}(H), \tau^{(t)})\|_2 d\mathbb{P}_n(H)}{\int_0^{+\infty} \|\text{clip}_{\mathbb{P}}(\nabla_p \phi_{H \rightarrow \ell}(H), \tau^{(t)})\|_2 d\mathbb{P}_c(H)}, \quad (10)$$

where  $\mathbb{P}_c$  and  $\mathbb{P}_n$  are the distribution functions of the clean and noise Gaussian distributions  $\mathcal{N}_c^{(t)}$  and  $\mathcal{N}_n^{(t)}$ , respectively. And the clipping threshold  $\tau^{(t)}$  can be obtained by solving the following optimizing problem:

$$\min_{\tau^{(t)}} \tilde{r}^{(t)}(\tau^{(t)}). \quad (11)$$

However, directly optimizing Eq.11 may result in excessive clipping. Specifically, one could clip as many gradients as possible to minimize the ratio. In such cases, although most of the noisy gradients are effectively removed, the clean gradients are also impacted, leading to model underfitting. To address this, we instead employ a conditioned version of Eq.11 to balance the clipping of noisy gradients while preserving the clean gradients:

$$\min_{\tau^{(t)}} \tilde{r}^{(t)}(\tau^{(t)}) \quad \text{s.t.} \quad \tilde{r}^{(t)} \geq 1 + \epsilon \quad (12)$$

where  $\epsilon > 0$  is a hyper-parameter, and  $1 + \epsilon$  is the pre-defined gradient ratio. This allows us to consistently determine a clipping threshold  $\tau^{(t)}$  that preserves the desired ratio, regardless of variations in the gradient distributions of clean and noisy samples, while preventing excessive clipping of the clean gradients. In practice, a straightforward binary search algorithm is applied to solve Eq.12.

**Implement optimized clipping on loss functions.** According to Menon et al. (2020, Lemma 5), a loss function  $\ell$  equipped with gradient clipping is equivalent to a Huberised-like loss function as follows: if  $\|\nabla_p(\ell(f(\mathbf{x}), y))\|_2 \geq \tau^{(t)}$ , then  $\bar{\ell}(f(\mathbf{x}), y, \tau^{(t)}) = 1 - \tau^{(t)} \cdot p(y|\mathbf{x}) + \phi_{\|g_p\|_2 \rightarrow \ell}(\tau^{(t)})$ ; otherwise,  $\bar{\ell}(f(\mathbf{x}), y, \tau^{(t)}) = \ell(f(\mathbf{x}), y)$ .

Take Cross Entropy (CE) loss  $\ell_{\text{CE}}$  as an example. By applying Optimized Gradient Clipping (OGC), we obtain CE+OGC, which is induced by a new loss function  $\bar{\ell}_{\text{CE}}(f(\mathbf{x}), y, \tau^{(t)})$  as follows: if  $\frac{1}{p(y|\mathbf{x})} \geq \tau^{(t)}$ , then  $\bar{\ell}_{\text{CE}}(f(\mathbf{x}), y, \tau^{(t)}) = 1 - \tau^{(t)} \cdot p(y|\mathbf{x}) + \log \tau^{(t)}$ ; otherwise,  $\bar{\ell}_{\text{CE}}(f(\mathbf{x}), y, \tau^{(t)}) = -\log p(y|\mathbf{x})$ .

## 4.3 Robustness Analysis

In the following, we provide a formal analysis of the robustness exhibited by our proposed OGC. Specifically, we will first introduce an indicator of robustness, ‘‘excess risk’’,

Methods	Sym-50%	Sym-80%	Asymmetric	Dependent	Real	Average
Fixed	84.56	34.21	58.95	44.24	79.84	60.36
Linear	81.01	42.01	77.90	65.68	75.63	68.44
EMA	84.98	21.31	81.30	73.06	<b>82.23</b>	68.57
Optimized (ours)	<b>85.16</b>	<b>43.77</b>	<b>81.51</b>	<b>78.28</b>	81.45	<b>74.03</b>

Table 1: Test accuracies (%) of  $\tau^{(t)}$  adjusting methods on CIFAR-10 dataset with different label noise, **bold** for best results.

which is the focus of our analysis. To simplify our discussion, we specifically analyze the new loss function  $\bar{\ell}_{CE}$  induced by CE+OGC, which is Cross Entropy equipped with OGC. (Our results can be easily generalized to other loss functions, such as Focal Loss (FL) and Generalized Cross Entropy (GCE).)

Given any classifier  $f$  and loss function  $\bar{\ell}_{CE}$  with  $\tau^{(t)}$ , the population risks of  $f$  under the clean distribution  $\mathcal{D}_c$  and the unknown distribution  $\mathcal{D}_\eta$  are respectively defined as

$$\begin{aligned}\mathcal{R}_{\bar{\ell}_{CE}}(f, \tau^{(t)}) &= \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_c}[\bar{\ell}_{CE}(f(\mathbf{x}), y, \tau^{(t)})], \\ \mathcal{R}_{\bar{\ell}_{CE}}^\eta(f, \tau^{(t)}) &= \mathbb{E}_{(\mathbf{x}, \tilde{y}) \sim \mathcal{D}_\eta}[\bar{\ell}_{CE}(f(\mathbf{x}), \tilde{y}, \tau^{(t)})].\end{aligned}\quad (13)$$

Let  $f^*$  be the global minimizer of  $\mathcal{R}_{\bar{\ell}_{CE}}(f, \tau^{(t)})$ , and let  $\tilde{f}^*$  be the global minimizer of  $\mathcal{R}_{\bar{\ell}_{CE}}^\eta(f, \tau^{(t)})$ . With the notations, we can then define the ‘‘excess risk’’ as  $\mathcal{R}_{\bar{\ell}_{CE}}(\tilde{f}^*, \tau^{(t)}) - \mathcal{R}_{\bar{\ell}_{CE}}(f^*, \tau^{(t)})$  or  $\mathcal{R}_{\bar{\ell}_{CE}}^\eta(f^*, \tau^{(t)}) - \mathcal{R}_{\bar{\ell}_{CE}}^\eta(\tilde{f}^*, \tau^{(t)})$  (which indicates the performance gap of  $\tilde{f}^*$  and  $f^*$ ).

We first show the boundedness of  $\bar{\ell}_{CE}$ , which is important to the control of ‘‘excess risk’’.

**Proposition 1** *Given any classifier  $f$ , for any input sample pair  $(\mathbf{x}, y)$  and any  $\tau^{(t)} \geq 1$ , the CE+OGC loss  $\bar{\ell}_{CE}$  is both lower and upper bounded:*

$$\begin{aligned}1 - p(j|x) &\leq \bar{\ell}_{CE}(f(\mathbf{x}), j, \tau^{(t)}) \\ &\leq (1 - p(j|x))(1 + \log \tau^{(t)}).\end{aligned}\quad (14)$$

Moreover, the sum of  $\bar{\ell}_{CE}$  w.r.t all classes is thus also lower bound and upper bounded:

$$\begin{aligned}K - 1 &\leq \sum_{j=1}^K \bar{\ell}_{CE}(f(\mathbf{x}), j, \tau^{(t)}) \\ &\leq (K - 1)(1 + \log \tau^{(t)}).\end{aligned}\quad (15)$$

The proof is deferred to Appendix. Proposition 1 demonstrates that Cross Entropy equipped with our proposed OGC is always bounded. Moreover, as  $\tau^{(t)}$  approaches 1, the bound becomes tighter. When  $\tau^{(t)} = 1$ , our loss function is equivalent to the symmetric Mean Absolute Error (MAE). Notably, this symmetry is a key property that makes a loss function robust to label noise (Ghosh, Kumar, and Sastry 2017).

Proposition 1 paves the way for further noise robustness analysis of CE+OGC loss  $\bar{\ell}_{CE}$ . Firstly, under the common instance-independent setting, we can depict the bound of excess risk for both symmetric and asymmetric label noise (see Equation (1) for the definitions).

**Theorem 1** (Excess risk under instance-independent symmetric label noise) *Under symmetric label noise with  $\eta \leq 1 - \frac{1}{K}$ ,*

$$\begin{aligned}0 &\leq \mathcal{R}_{\bar{\ell}_{CE}}(\tilde{f}^*, \tau^{(t)}) - \mathcal{R}_{\bar{\ell}_{CE}}(f^*, \tau^{(t)}) \\ &\leq \log \tau^{(t)} / \left(1 - \frac{\eta K}{K - 1}\right).\end{aligned}\quad (16)$$

**Theorem 2** (Excess risk under instance-independent asymmetric label noise) *Under asymmetric label noise with  $\eta_{ij} < 1 - \eta_i, \forall j \neq i, \forall i, j \in [k]$ , where  $1 - \eta_i = \eta_{ii}$ , if  $\mathcal{R}_{\bar{\ell}_{CE}}(f^*, \tau^{(t)}) = 0$ , then*

$$\begin{aligned}0 &\leq \mathcal{R}_{\bar{\ell}_{CE}}^\eta(f^*, \tau^{(t)}) - \mathcal{R}_{\bar{\ell}_{CE}}^\eta(\tilde{f}^*, \tau^{(t)}) \\ &\leq (K - 1)(\log \tau^{(t)})\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_c}[1 - \eta_y].\end{aligned}\quad (17)$$

The proofs for Theorem 1 and Theorem 2 can be found in Appendix. These theorems illustrate that, under both symmetric and asymmetric label noise, with our proposed OGC method, the excess risk between  $\tilde{f}^*$  and  $f^*$  is consistently bounded. Meanwhile, these bounds get tighter as  $\tau^{(t)}$  decreases.

Next, we further extend our analysis to scenarios where the label noise is instance-dependent (see Equation 2). With a regular mild assumption on signal-to-noise ratio that  $\eta_{yj}(\mathbf{x}) < \eta_{yy}(\mathbf{x}), \forall \mathbf{x}, j \neq y$ , we have the following result:

**Theorem 3** (Excess risk under instance-dependent label noise) *Under instance-dependent label noise with  $\eta_{yj}(\mathbf{x}) < \eta_{yy}(\mathbf{x}), \forall \mathbf{x}, j \neq y$ , if  $\mathcal{R}_{\bar{\ell}_{CE}}(f^*, \tau^{(t)}) = 0$ , then*

$$\begin{aligned}0 &\leq \mathcal{R}_{\bar{\ell}_{CE}}^\eta(f^*, \tau^{(t)}) - \mathcal{R}_{\bar{\ell}_{CE}}^\eta(\tilde{f}^*, \tau^{(t)}) \\ &\leq (K - 1)(\log \tau^{(t)})\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_c}[1 - \eta_y(\mathbf{x})].\end{aligned}\quad (18)$$

The proof for Theorem 3 is provided in Appendix. This theorem demonstrates that, even under instance-dependent label noise, our proposed OGC method ensures that the risk differences between  $\tilde{f}^*$  and  $f^*$  remain consistently bounded.

## 5 Experiments

### 5.1 Empirical Understanding

In this subsection, to enhance our understanding of the proposed OGC, we conduct a series of experiments. Unless otherwise stated, our experimental settings follow those described in Section 5.2. For additional experimental results on parameter analysis, including predefined gradient ratio  $1 + \epsilon$ , queue  $Q$ , time frame  $s$ , and time consumption, please refer to Appendix.

Methods	Sym-50%	Sym-80%	Asymmetric	Dependent	Real	Average
GCE	85.36	34.72	59.87	66.76	81.33	65.61
GCE+OGC	86.23	38.68	80.88	77.22	81.75	72.95
Improvement	+ 0.87	+ 3.96	+ 21.01	+ 10.46	+ 0.42	+ 7.34

Table 2: Test accuracies (%) of GCE and GCE+OGC on CIFAR-10 dataset with different label noise.

Methods	Sym-20%	Sym-50%	Sym-80%	Asymmetric	Dependent	Real	Average
CE	81.25±0.35	51.58±0.41	35.97±5.46	76.20±0.38	60.37±0.68	64.43±0.24	61.63
FL	82.32±0.39	52.65±0.35	37.30±1.13	76.69±0.15	61.12±0.31	65.57±0.39	62.61
MAE	89.99±0.14	75.80±3.72	18.81±2.01	56.09±0.29	15.70±1.79	55.93±3.67	52.05
GCE	91.24±0.13	<u>85.36±0.20</u>	34.72±3.65	59.87±0.10	66.76±0.38	81.33±0.06	69.88
SCE	91.34±0.05	84.91±0.33	40.59±0.48	79.52±0.41	77.75±1.08	81.20±0.15	75.89
PHuber-CE	90.75±0.13	84.56±0.14	34.21±3.91	58.95±0.23	44.24±10.21	79.84±0.17	65.43
Taylor-CE	91.24±0.08	84.60±0.10	<u>43.69±5.26</u>	59.19±0.23	57.26±6.73	80.08±0.14	69.34
NCE+RCE	91.21±0.09	84.95±0.27	26.38±3.05	78.58±0.32	<u>78.75±0.37</u>	80.78±0.21	73.44
JS	91.65±0.05	83.79±0.21	40.74±5.98	75.45±0.35	<u>68.24±1.10</u>	78.55±0.31	71.40
LC-CE	91.69±0.08	79.29±0.35	35.77±3.27	75.48±0.06	61.19±0.40	74.67±0.43	69.68
NCE+NNCE	<u>90.96±0.08</u>	85.19±0.28	17.56±0.36	77.63±0.26	<b>80.92±0.66</b>	<u>82.39±0.08</u>	72.44
CE+OGC	<b>91.80±0.13</b>	85.16±0.26	<b>43.77±2.27</b>	81.51±0.31	78.28±0.12	81.45±0.68	<u>76.99</u>
FL+OGC	88.51±0.41	<b>85.43±0.13</b>	41.43±1.02	<b>86.34±0.49</b>	78.21±0.86	<b>82.69±0.30</b>	<b>77.10</b>

Table 3: Test accuracies (%) of different methods on CIFAR-10 datasets with different label noise. The results (mean±std) are reported over 3 random runs. **Bold** denotes the best results and underline denotes the second-best results.

**Comparison with manually designed methods.** To highlight the advantages of our optimize-based strategy, we compare it on CIFAR-10 against manually designed approaches for tuning  $\tau^{(t)}$ : 1) Linear Decrease Strategy: This approach is defined by  $\tau^{(t)} = \beta \cdot (1 - \frac{t}{T})$ , where  $T$  represents the total number of time steps and  $\beta$  is a parameter. 2) Exponential Moving Average (EMA) Strategy: This approach uses the formula  $\frac{1}{\tau^{(t)}} = \alpha \cdot \frac{1}{\tau^{(t-1)}} + (1 - \alpha) \cdot 1$ , where  $\alpha$  is a parameter. We also compare our optimize-based strategy with a fixed strategy proposed by Menon et al. (2020), where  $\tau^{(t)}$  remains constant.

The experimental results are presented in Table 1, and the experimental details can be found in Appendix.

As observed, our optimize-based strategy consistently outperforms other approaches in most label noise scenarios, particularly in cases with high noise rates (Sym-80%) and instance-dependent noise. Even in scenarios where it does not achieve the highest accuracy, it remains highly competitive, with only a 0.8% difference compared to the EMA method. Furthermore, our proposed optimize-based strategy achieves the highest average performance, outperforming the second-best method by 5%, highlighting its superior overall effectiveness. These experimental results underscore the robustness and efficacy of our optimize-based strategy across various types of label noise.

**Can OGC improves robust loss functions?** Our proposed OGC can be applied not only to non-robust loss functions, such as CE and FL, but also to robust loss functions. This naturally raises the question of whether an existing robust loss function can improve its performance by incorporating our proposed OGC. Here, we use Generalized Cross Entropy (GCE) (Zhang and Sabuncu 2018) as an example

and demonstrate that equipping with OGC can enhance the performance of a robust loss function.

Specifically, we conducted a series of experiments on the CIFAR-10 dataset under various types of label noise. For our GCE+OGC, GCE equipped with OGC, we set the parameters  $(q, \epsilon_0)$  to  $(0.7, 20.0)$ . The results, which reflect the average accuracies over the last 10 epochs, are presented in Table 2. As observed, equipping GCE with our OGC results in consistent performance improvements. Notably, under asymmetric label noise, we achieved a 21.01% improvement. These results demonstrate that our proposed OGC can enhance the effectiveness of existing robust loss functions.

## 5.2 Evaluation on Benchmark Datasets

In this subsection, we evaluate our proposed method using two well-known datasets: CIFAR-10 and CIFAR-100 (Krizhevsky, Hinton et al. 2009). To demonstrate the robustness of our proposed method against different types of label noise, we compare it with other state-of-the-art methods under various label noise conditions, including symmetric, asymmetric, instance-dependent (Xia et al. 2020), and real-world (Wei et al. 2022a) label noise settings.

**Baselines.** We consider several state-of-the-art methods: 1) Generalized Cross Entropy (GCE) (Zhang and Sabuncu 2018); 2) Symmetric Cross Entropy (SCE) (Wang et al. 2019); 3) Partially Huberised Cross Entropy (PHuber-CE) (Menon et al. 2020); 4) Taylor Cross Entropy (Taylor-CE) (Feng et al. 2020); 5) Normalized Loss Functions (Ma et al. 2020), specifically NCE+RCE; 6) Jensen-Shannon Divergence Loss (JS) (Engleson and Azizpour 2021); 7) Logit Clipping on Cross Entropy (LC-CE) (Wei et al. 2023); 8) Normalized Negative Loss Functions (Ye et al. 2023),

Methods	Sym-20%	Sym-50%	Sym-80%	Asymmetric	Dependent	Real	Average
CE	64.75±0.22	49.56±0.63	8.94±0.53	45.32±0.23	48.78±0.33	54.64±0.09	45.33
FL	64.58±0.77	50.27±0.28	9.52±0.55	46.55±0.14	49.67±0.35	53.80±0.11	45.73
MAE	5.46±0.69	3.63±0.12	0.99±0.01	2.81±0.45	1.36±0.27	2.79±0.80	14.86
GCE	71.14±0.23	65.18±0.52	11.99±0.37	42.59±0.56	51.43±0.69	56.58±0.21	49.65
SCE	66.01±0.23	57.00±0.37	6.44±0.17	39.93±0.41	45.53±0.46	53.45±0.13	45.23
PHuber-CE	54.08±1.67	48.94±0.34	15.79±1.07	31.97±0.41	20.22±0.91	34.71±0.66	37.33
Taylor-CE	69.87±0.36	63.84±0.12	23.10±0.59	41.94±0.29	43.44±0.75	49.76±0.48	50.20
NCE+RCE	69.95±0.19	61.31±0.10	14.87±0.39	42.65±0.25	51.24±0.22	56.40±0.30	48.74
JS	71.15±0.34	65.60±0.11	17.15±0.37	44.82±0.60	51.37±0.87	55.17±1.20	51.07
LC-CE	71.40±0.16	62.61±0.11	13.16±0.54	42.84±0.25	47.01±0.52	55.87±0.35	48.65
NCE+NNCE	69.07±0.24	62.09±0.12	7.93±0.88	<b>50.86±0.25</b>	55.31±2.13	57.55±0.19	50.81
CE+OGC	<b>71.41±0.42</b>	<b>67.61±0.46</b>	<u>28.70±1.54</u>	45.37±0.25	<u>57.28±1.40</u>	<u>59.93±0.15</u>	<u>54.78</u>
FL+OGC	71.13±0.34	<u>66.65±0.09</u>	<b>31.33±0.42</b>	<u>47.84±1.06</u>	<b>60.83±0.85</b>	<b>60.30±0.28</b>	<b>56.36</b>

Table 4: Test accuracies (%) of different methods on CIFAR-100 datasets with different label noise. The results (mean±std) are reported over 3 random runs. **Bold** denotes the best results and underline denotes the second-best results.

Method	CE	GCE	SCE	PHuber-CE	NCE+RCE	LC-CE	CE+OGC	FL+OGC
ILSVRC2012 Val	64.96	56.60	63.47	63.75	65.08	66.47	<u>67.40</u>	<b>68.28</b>
WebVision Val	68.51	58.20	65.63	65.43	66.75	69.56	<u>70.28</u>	<b>70.60</b>

Table 5: Top-1 validation accuracies (%) of different methods on the ILSVRC12 and WebVision validation sets, under the “Mini” setting (Jiang et al. 2018). **Bold** denotes the best results and underline denotes the second-best results.

specifically NCE+NNCE. Additionally, we employ Cross Entropy (CE), Focal Loss (FL) (Lin et al. 2017), and Mean Absolute Error (MAE) (Ghosh, Kumar, and Sastry 2017) for network training. For our proposed Optimized Gradient Clipping (OGC), we consider two loss functions: 1) CE+OGC and 2) FL+OGC.

**Experimental details.** Noise generation, training, and parameter settings are in Appendix.

**Results.** The results for CIFAR-10 and CIFAR-100 are presented in Table 3 and Table 4, respectively. The reported values reflect the average accuracies over the last 10 epochs. For both datasets, the incorporation of OGC enables CE and FL to consistently achieve the top-2 average accuracy across various scenarios. Notably, CE and FL are among the least robust loss functions, emphasizing OGC’s effectiveness in enhancing the performance of non-robust loss functions. Notably, CE and FL, typically among the least robust loss functions, show significant improvement with OGC, highlighting its effectiveness in enhancing the performance of non-robust loss functions. For instance, in Table 3, on CIFAR-10 with real-world label noise, integrating OGC with CE achieves a direct improvement (+17.02%). Moreover, in Table 4, with the integration of OGC, CE and FL consistently achieve the top two accuracies in all scenarios except for CIFAR-10 with asymmetric label noise. Additionally, due to our optimize-based strategy, our loss function requires only a single hyper-parameter, significantly reducing the time needed for hyper-parameter tuning compared to the previous SOTA method, NCE+NNCE, which involves three hyper-parameters. These observations underscore the effectiveness of our proposed OGC.

### 5.3 Evaluation on Real-world Noisy Dataset

To evaluate our proposed method on large-scale real-world noisy dataset, we conduct experiments on WebVision 1.0 dataset. WebVision 1.0 (Li et al. 2017) contains more than 2.4 million web images crawled from the internet by using queries generated from the 1,000 class labels of the ILSVRC 2012 (Deng et al. 2009) benchmark. Following the “Mini” setting in previous works (Jiang et al. 2018), we only take the first 50 classes of the Google resized image subset. We evaluate the trained networks on the same 50 classes of both the ILSVRC 2012 validation set and WebVision 1.0 validation set, these can be considered as clean validation sets.

The experimental details can be found in Appendix. The results of last epoch validation accuracies are reported in Table 5. As observed, integrating our proposed OGC with CE and FL results in significant improvements, outperforming existing robust loss functions on both validation sets. These results verify that our method is effective in enhancing noise-robustness in large-scale real-world scenarios.

## 6 Conclusions

In this paper, we introduce Optimized Gradient Clipping (OGC), a novel method for noisy label learning, which employs a 2-component Gaussian Mixture Model (2-GMM) to distinguish between clean and noisy samples, estimating the noise-to-clean gradient ratio after clipping. This allows us to set a clipping threshold at each training step, limiting the influence of noise gradients. OGC enhances the robustness of both non-robust and robust loss functions in noisy label settings. Our theoretical analysis highlights its noise tolerance, and extensive experiments demonstrate its effectiveness across various noise types.

## Acknowledgments

This work is supported by the Shanghai Engineering Research Center of Intelligent Computing System (grant No. 19DZ2252600) and the Research Grants Council (RGC) under grant ECS-22303424.

## References

- Arazo, E.; Ortego, D.; Albert, P.; O'Connor, N. E.; and McGuinness, K. 2019. Unsupervised Label Noise Modeling and Loss Correction. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, 312–321. PMLR.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, 248–255. IEEE Computer Society.
- Engleson, E.; and Azizpour, H. 2021. Generalized Jensen-Shannon Divergence Loss for Learning with Noisy Labels. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 30284–30297.
- Feng, L.; Shu, S.; Lin, Z.; Lv, F.; Li, L.; and An, B. 2020. Can Cross Entropy Loss Be Robust to Label Noise? In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020, 2206–2212*. ijcai.org.
- Ghosh, A.; Kumar, H.; and Sastry, P. S. 2017. Robust Loss Functions under Label Noise for Deep Neural Networks. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, 1919–1925. AAAI Press.
- Gui, X.; Wang, W.; and Tian, Z. 2021. Towards Understanding Deep Learning from Noisy Labels with Small-Loss Criterion. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, 2469–2475. ijcai.org.
- Han, B.; Yao, J.; Niu, G.; Zhou, M.; Tsang, I. W.; Zhang, Y.; and Sugiyama, M. 2018. Masking: A New Perspective of Noisy Supervision. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 5841–5851.
- Jiang, L.; Zhou, Z.; Leung, T.; Li, L.; and Fei-Fei, L. 2018. MentorNet: Learning Data-Driven Curriculum for Very Deep Neural Networks on Corrupted Labels. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, 2309–2318. PMLR.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Li, J.; Socher, R.; and Hoi, S. C. H. 2020. DivideMix: Learning with Noisy Labels as Semi-supervised Learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Li, W.; Wang, L.; Li, W.; Agustsson, E.; and Gool, L. V. 2017. WebVision Database: Visual Learning and Understanding from Web Data. *CoRR*, abs/1708.02862.
- Lin, T.; Goyal, P.; Girshick, R. B.; He, K.; and Dollár, P. 2017. Focal Loss for Dense Object Detection. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, 2999–3007. IEEE Computer Society.
- Liu, S.; Niles-Weed, J.; Razavian, N.; and Fernandez-Granda, C. 2020. Early-Learning Regularization Prevents Memorization of Noisy Labels. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Lukasik, M.; Bhojanapalli, S.; Menon, A. K.; and Kumar, S. 2020. Does label smoothing mitigate label noise? In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 6448–6458. PMLR.
- Ma, X.; Huang, H.; Wang, Y.; Romano, S.; Erfani, S. M.; and Bailey, J. 2020. Normalized Loss Functions for Deep Learning with Noisy Labels. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 6543–6553. PMLR.
- Menon, A. K.; Rawat, A. S.; Reddi, S. J.; and Kumar, S. 2020. Can gradient clipping mitigate label noise? In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Natarajan, N.; Dhillon, I. S.; Ravikumar, P.; and Tewari, A. 2013. Learning with Noisy Labels. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, 1196–1204.
- Nguyen, D. T.; Mummadi, C. K.; Ngo, T.; Nguyen, T. H. P.; Beggel, L.; and Brox, T. 2020. SELF: Learning to Filter Noisy Labels with Self-Ensembling. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Pascanu, R.; Mikolov, T.; and Bengio, Y. 2013. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, volume 28 of *JMLR Workshop and Conference Proceedings*, 1310–1318. JMLR.org.
- Patrini, G.; Rozza, A.; Menon, A. K.; Nock, R.; and Qu, L. 2017. Making Deep Neural Networks Robust to Label Noise: A Loss Correction Approach. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR*

- 2017, Honolulu, HI, USA, July 21-26, 2017, 2233–2241. IEEE Computer Society.
- Wang, Y.; Ma, X.; Chen, Z.; Luo, Y.; Yi, J.; and Bailey, J. 2019. Symmetric Cross Entropy for Robust Learning With Noisy Labels. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, 322–330. IEEE.
- Wei, H.; Feng, L.; Chen, X.; and An, B. 2020. Combating Noisy Labels by Agreement: A Joint Training Method with Co-Regularization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, 13723–13732. Computer Vision Foundation / IEEE.
- Wei, H.; Zhuang, H.; Xie, R.; Feng, L.; Niu, G.; An, B.; and Li, Y. 2023. Mitigating Memorization of Noisy Labels by Clipping the Model Prediction. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, 36868–36886. PMLR.
- Wei, J.; Zhu, Z.; Cheng, H.; Liu, T.; Niu, G.; and Liu, Y. 2022a. Learning with Noisy Labels Revisited: A Study Using Real-World Human Annotations. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Wei, Q.; Sun, H.; Lu, X.; and Yin, Y. 2022b. Self-Filtering: A Noise-Aware Sample Selection for Label Noise with Confidence Penalization. In *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXX*, volume 13690 of *Lecture Notes in Computer Science*, 516–532. Springer.
- Xia, X.; Liu, T.; Han, B.; Wang, N.; Gong, M.; Liu, H.; Niu, G.; Tao, D.; and Sugiyama, M. 2020. Part-dependent Label Noise: Towards Instance-dependent Label Noise. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Ye, X.; Li, X.; Dai, S.; Liu, T.; Sun, Y.; and Tong, W. 2023. Active Negative Loss Functions for Learning with Noisy Labels. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Yu, X.; Han, B.; Yao, J.; Niu, G.; Tsang, I. W.; and Sugiyama, M. 2019. How does Disagreement Help Generalization against Label Corruption? In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, 7164–7173. PMLR.
- Zhang, C.; Bengio, S.; Hardt, M.; Recht, B.; and Vinyals, O. 2017. Understanding deep learning requires rethinking generalization. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Zhang, H.; Cissé, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond Empirical Risk Minimization. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Zhang, J.; He, T.; Sra, S.; and Jadbabaie, A. 2020. Why Gradient Clipping Accelerates Training: A Theoretical Justification for Adaptivity. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Zhang, Z.; and Sabuncu, M. R. 2018. Generalized Cross Entropy Loss for Training Deep Neural Networks with Noisy Labels. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 8792–8802.
- Zhou, X.; Liu, X.; Jiang, J.; Gao, X.; and Ji, X. 2021. Asymmetric Loss Functions for Learning with Noisy Labels. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, 12846–12856. PMLR.