

# Learning with Noisy labels via Self-supervised Adversarial Noisy Masking

Yuanpeng Tu<sup>1\*</sup> Boshen Zhang<sup>2\*</sup> Yuxi Li<sup>2</sup> Liang Liu<sup>2</sup> Jian Li<sup>2</sup>  
Jiangning Zhang<sup>2</sup> Yabiao Wang<sup>2†</sup> Chengjie Wang<sup>2,3</sup> Cai Rong Zhao<sup>1†</sup>

<sup>1</sup>Dept. of Electronic and Information Engineering, Tongji University, Shanghai

<sup>2</sup>YouTu Lab, Tencent, Shanghai, <sup>3</sup>Shanghai Jiao Tong University

{2030809, zhaocairong}@tongji.edu.cn

{boshenzhang, yukiyxli, leoneliu, swordli, vtzhang, caseywang, jasoncjwang}@tencent.com

## Abstract

Collecting large-scale datasets is crucial for training deep models, annotating the data, however, inevitably yields noisy labels, which poses challenges to deep learning algorithms. Previous efforts tend to mitigate this problem via identifying and removing noisy samples or correcting their labels according to the statistical properties (e.g., loss values) among training samples. In this paper, we aim to tackle this problem from a new perspective, delving into the deep feature maps, we empirically find that models trained with clean and mislabeled samples manifest distinguishable activation feature distributions. From this observation, a novel robust training approach termed adversarial noisy masking is proposed. The idea is to regularize deep features with a label quality guided masking scheme, which adaptively modulates the input data and label simultaneously, preventing the model to overfit noisy samples. Further, an auxiliary task is designed to reconstruct input data, it naturally provides noise-free self-supervised signals to reinforce the generalization ability of models. The proposed method is simple yet effective, it is tested on synthetic and real-world noisy datasets, where significant improvements are obtained over previous methods. Code is available at <https://github.com/yuanpengtu/SANM>.

## 1. Introduction

Deep learning has achieved remarkable success, relying on large-scale datasets with human-annotated accurate labels. However, collecting such high-quality labels is time-consuming and expensive. As an alternative, inexpensive strategies are usually used for generating labels for large-scale samples, such as web crawling, leveraging search engines, or using machine-generated annotations. All these

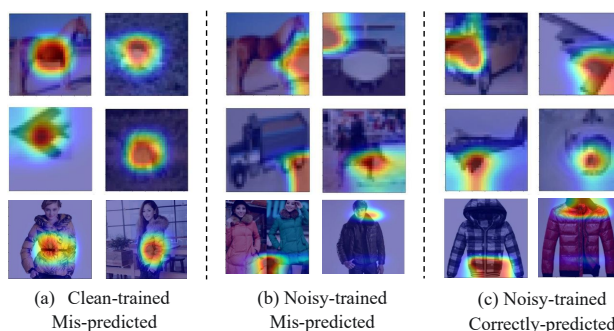


Figure 1. Activation maps of mis-predicted (a-b) and correctly-predicted (c) samples when training PreAct ResNet-18 with clean (i.e., clean-trained) and noisy (i.e., noisy-trained) data on CIFAR-10 (1st and 2nd row) and Clothing1M [38] (3rd row).

alternative methods inevitably yield numerous noisy samples. However, previous research [2] has revealed that deep networks can easily overfit to noisy labels and suffer from dramatic degradation in the generalization performance.

Towards this problem, numerous Learning with Noisy labels (LNL) approaches have been proposed. Sample selection methods [3, 8, 34, 41] are the most straightforward methods, which attempt to select clean samples based on certain criterion (e.g., loss value) and then reweight or discard the noisy instances to reduce the interference of mislabeled training samples. However, these methods fail to leverage the potential information of the discarded samples. Similarly, label correction based methods [32, 40, 48] try to correct labels, which often impose assumptions on the existence of a small correctly-labeled validation subset or directly utilize predictions of deep models. However, such assumptions can not be always full-filled in real-world noisy datasets [27], leading to limited application scenarios. Besides, the predictions of deep model tend to fluctuate when training with mislabeled samples, making the label correction process unstable and sub-optimal [36]. On the contrary, regularization based methods [2, 6, 15, 20, 26, 29, 43, 47, 50] aim to alleviate the side effect of label noise by preventing the deep models from overfitting to all training sam-

\* Equal contribution.

† Corresponding author.

ples, which is flexible and can work collaboratively with other LNL methods. Both explicit and implicit regularization techniques were proposed, the former generally calibrates the parameter update of networks by modifying the expected training loss, i.e., dropout. The latter focuses on improving the generalization by utilizing the stochasticity, i.e., data augmentation strategy and stochastic gradient descent [29]. However, most of these methods are designed for general fully-supervised learning tasks with correctly-labeled samples, and poor generalization ability could be obtained when the label noise is severe [29].

In this paper, we present a novel regularization based method specialized for LNL problem. Our method is inspired by an observation that some differences exist in activation maps between the predictions of deep models trained with clean and noisy labels. As shown in Fig. 1 (a), activation maps of mis-predicted samples by the model trained with clean labels are always focused on their foreground areas. By contrast, when the model is trained with noisy labels, it tends to generate results focused on the meaningless background area for the mis-predicted samples (Fig. 1 (b)). And even for cases that the prediction is correct, i.e., Fig. 1 (c), the high-activated region drifts to irrelevant (i.e., edge) areas of the object, rather than the main regions of the object. This indicates that *model trained with mislabeled samples is likely to overfit to some corner parts of the object from noisy data or remember the less-informative regions (i.e., background)*. Therefore, from this observation, we propose a novel Self-supervised Adversarial Noisy Masking (SANM) method for LNL. The idea is to explicitly regularize the model to generate more diverse activation maps through masking the certain region (e.g., max-activated) of the input image and thus alleviates the confirmation bias [8].

Moreover, to investigate how the masking strategies affect the performance of training deep models with noisy labels, we conduct a simple experiment, where different adversarial masking strategies (i.e., masking the max-activated region of input images with fixed, random, and noise-aware mask ratio) were applied to the popular LNL framework DivideMix [15]. As shown in Fig. 2, we find that the performance gain is much more sensitive to the masking strategy. A fixed or random mask ratio obtain only marginal gain, but when dealing with the correctly-labeled and mislabeled samples with different mask ratios, a significant performance gain is achieved. This indicates that how to design an effective adversarial masking method in the case of LNL is non-trivial and not well addressed.

Towards the problem above, a label quality guided adaptive masking strategy is proposed, which modulates the image label and the mask ratio of image masking simultaneously. The label is updated via first leveraging the soft-distributed model prediction and then reducing the proba-

bility of max-activated class with the noise-aware masking ratio, while at the same time lifting the probability of other classes. Intuitively, a sample with a high probability of being mislabeled will possess a larger masking ratio, leading to a more strictly regularized input image and modulated label. Therefore, the negative impact of noisy labels can be greatly reduced. As for the correctly-labeled samples, the masking ratio is relatively small, which plays the role of general regularization strategy and improves the model generalization ability by preventing the model from overfitting training samples. Further, we specially customized an auxiliary decode branch to reconstruct the original input image, which provides noise-free self-supervised information for learning a robust deep model. The proposed SANM is flexible and can further boost existing LNL methods. It is tested on synthetic and real-world large-scale noisy datasets, and elaborately ablative experiments were conducted to verify each design component of the proposed method. In a nutshell, the key contributions of this work are:

- We propose a novel self-supervised adversarial noisy masking method named SANM to explicitly impose regularization for LNL problem, preventing the model from overfitting to some corner parts of the object or less-informative regions from noisy data;
- A label quality guided masking strategy is proposed to differently adjust the process for clean and noisy samples according to the label quality estimation. This strategy modulates the image label and the ratio of image masking simultaneously;
- A self-supervised mask reconstruction auxiliary task is designed to reconstruct the original images based on the features of masked ones, which aims at enhancing generalization by providing noise-free supervision signals.

## 2. Related Work

### 2.1. Learning with Noisy Labels

A variety of methods have been proposed to improve the robustness of DNNs on noisy datasets, which consist of regularization, label correction, and noisy sample selection.

**Regularization** based methods help resist memorizing noisy data since previous works [2] have revealed that networks always firstly learn simple patterns before memorizing noisy hard data. Previous research [29] claims that these methods can be roughly divided into two categories: explicit/implicit regularization. Explicit regularization based methods generally modify the losses for the training samples. Among them, ELR [20] tries to alleviate the influence of noisy data by regularizing the loss gradient. SR [50] proposes a sparse regularization method to restrict network output to the permutation set over a one-hot vector. Meta learning based methods [26] impose regularization on the sample gradient directions by bi-level optimization with a small un-

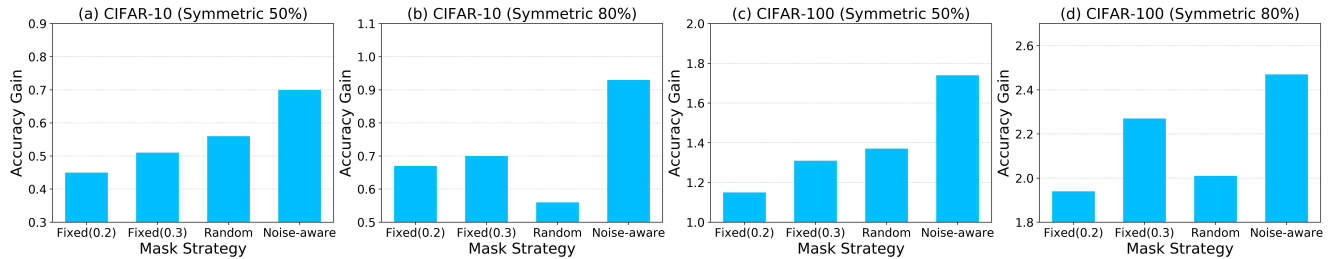


Figure 2. A experiment for masking the max-activated region with different mask ratios. The performance gains of different mask strategies under 50% and 80% symmetric noise of CIFAR-10/100 [14] are reported, where DivideMix [15] is adopted as the baseline. "Fixed(0.2/0.3)" denotes masking all the images with the same mask ratio of 0.2/0.3. "Random" represents masking images with a random mask ratio between 0.2 and 0.4. "Noise-aware" is masking noisy samples with a mask ratio of 0.3 while the ratio for clean ones is 0.2.

biased validation. In addition to these regularizers, example re-weighting strategies are also popular explicit regularization solutions to the over-fitting problem. However, these methods often need sensitive hyper-parameters or deeper networks to improve the model capacity. For implicit regularization methods, they generally utilize the stochasticity effect, i.e., augmentation and stochastic gradient descent. Among them, semi-supervised based LNL methods [15, 47] generally use mixup [43] as an effective regularization strategy to hinder over-fitting, which makes the network favor simple linear behavior in-between training examples. Adversarial samples based methods [6] improve the robustness by helping the model classify original and perturbed inputs. Since the model is unable to distinguish right from wrong by itself, regularization can generally work surprisingly well in certain synthetic noisy cases.

**Label correction** methods seek to transform noisy labels to correct ones directly. The pseudo labeling strategy is commonly adopted by these methods, which utilizes the confident output of the network. Joint [32] optimizes labels and network parameters in an alternating strategy. PENCIL [40] corrects and updates labels probabilistically through the back-propagation mechanism. MLC [48] trains a label correction network through meta learning to generate pseudo labels for the training of main models. However, these approaches either need extra clean validation data or suffer from significant performance drop under high noise settings due to the low quality of generated pseudo labels.

**Noisy sample selection** approaches mostly exploit the small-loss criterion, which is based on the observation that samples with noisy labels tend to have much higher losses than the clean ones. Among these methods, MentorNet uses a pre-trained teacher model to select clean samples for the training of a student model. Co-teaching based methods [8, 34, 41, 42] update the parameters of two parallel networks with clean instances chosen from their peer networks, which target avoiding the confirmation bias problem. INCV [3] attempts to separate clean data from the noisy training set by using cross-validation and removing high-loss samples. Though these methods may reduce the risk of error correction, they fail to use the potential infor-

mation of the abandoned samples. Moreover, most of them need prior knowledge, i.e., the noise rate.

## 2.2. Masking Regularization Techniques

As a commonly adopted strategy in weakly-supervised semantic segmentation, adversarial masking regularization aims at helping models discover less apparent semantic objects. Specifically, a Class Activation feature Map (CAM) [49] branch is first used to generate initial attention maps and selectively mask the discovered areas in the images through a threshold. Then the masked images are used for the training of another network. AE-PSL [35] iteratively erases the images and expands target object regions with the erased areas to produce final contiguous segmentation masks. Self-erasing [12] utilizes the background prior to make models generate more refined and focused attention on the foreground area. GAIN [18] tries to impose constraints on the activation maps in a regularized style to generate more complete attention regions. PSA [30] trains two networks with opposing losses to eliminate the drawbacks of certain mask strategies. In addition to weakly-supervised semantic segmentation, adversarial masking is also used in fine-grained visual classification. DB [31] designs a diversification block to mask out the most salient features, aiming at forcing networks to mine subtle differences between images. WS-DAN [13] randomly crop and drop one of the discriminative regions in the attention maps to encourage models to mine local features. Several adversarial masking strategies are proposed in [7, 39] to preserve human privacy by masking the minimal number of pixels required for identification. [33] designs an adversarial discriminator for mask generation at inference to correct segmentation errors. However, the purpose of these methods is quite different from our approach. Based on the assumption that the samples are all correctly-labeled, previous methods generally retain the original most salient regions while helping models mine more subtle information through masking regularization. By contrast, our method aims to regularize the model through adversarial masking to make the activation map more dispersed rather than concentrated to specific regions, and a similar masking strategy is applied to gener-

ate more smooth pseudo labels, thus avoiding over-fitting to noisy labels and alleviating the confirmation bias.

### 3. Method

**Overview.** For convenience, we begin by introducing the notations. Consider LNL as a classification task, the class number is  $c$ . The noisily-labeled training dataset is denoted as  $D = \{(x_i, y_i)\}_{i=1}^N$ , where  $x_i \in \mathbb{R}^{H_i \times W_i \times 3}$  and  $y_i \in \{0, 1\}^c$  are the  $i$ -th sample and its noisy label,  $H_i, W_i$  are the height and width of samples.  $N$  represents the dataset size. The encoder and decoder models are denoted as  $M_E$  and  $M_D$  respectively. We firstly introduce the overall framework of the proposed SANM. As shown in Fig. 3, a label quality guided masking scheme is proposed to adaptively perform adversarial noisy masking for each sample where a contrastive regularization strategy is imposed on the noisy predictions as well. Afterward, the representations of the masked images are used for the self-supervised reconstruction of original images, which aims at preventing models from overfitting to the noisy labels by providing noise-free signals. For application, SANM can be easily integrated with existing LNL frameworks to further boost their performance.

#### 3.1. Adversarial Noisy Masking

**Adversarial Masking Generation.** To highlight the semantic regions of correctly labeled samples and keep the captured attention areas of the noisy samples from expanding to meaningless background regions, we propose a novel Adversarial noisy Masking Generation (AMG) strategy during the training process. Specifically, we begin by extracting the feature of image  $x_i$  through  $M_E$  and the obtained feature map and prediction are denoted as  $F_i \in \mathbb{R}^{H_F \times W_F \times N_C}$  and  $\tilde{y}_i$  respectively, where  $H_F, W_F, N_C$  are the height, width and channel number of the feature layer. For each sample, a two-component Gaussian Mixture Model (GMM) [25] is fit to the cross-entropy loss with the Expectation-Maximization algorithm, where the generated weight for  $x_i$  is denoted as  $G_i$ . Afterward,  $F_i$  is used to generate the activation maps [49]  $A_i \in \mathbb{R}^{H_i \times W_i}$  of the category with the largest prediction, which is derived as:

$$A_i = \text{CAM}(F_i, \text{argmax}(\tilde{y}_i)) \quad (1)$$

where  $\text{CAM}(\cdot)$  is the activation map generation function. Then the coordinates with the maximum and minimum activation values for  $x_i$  are recorded as  $\{h_i^{\max}, w_i^{\max}\}$  and  $\{h_i^{\min}, w_i^{\min}\}$  respectively. For the  $\{h_i^{\max}, w_i^{\max}\}$ , assuming that the GMM division is overall correct, it is believed that the activation maps of correct-labeled samples are mainly focused on the target area of the foreground, while due to the interference of noisy labels, the max-activated region in the activation maps of mislabeled sam-

ples are mostly concentrated in meaningless background areas. To breakdown the over-fitting to noisy labels, we intend to adversarially augment training samples via a novel making strategy. In particular, denote the loss function as  $\mathcal{L}(x_i, y_i)$ , we try to maximize  $\mathcal{L}$  by adaptively generating region perturbations for noisy and clean samples across the training process. The objective can be derived as  $\max \mathcal{L}(x_i + \sigma, y_i)$ , where  $\sigma$  is the generated perturbations. Specifically, we generate masks with larger ratios for the max-activated region of noisy samples to force the model to reduce its response to these areas and pay more attention to the rest area, where the target is more likely to be found. For the correct-labeled ones, high response areas in their activation maps are likely to exist in the edge area of the object since they are affected by the gradient of noisy samples. Thus smaller mask ratios are used for the max-activated regions of them, which keep the model from ignoring the masked area and force it to mine discriminative information from the rest areas of the target since the correct labels help limit the network to pay more attention to the foreground area. However, it will make the activation map of the well-trained model scattered to the meaningless area when filtering out nearly all the peak locations for each sample through a threshold like existing methods, which violates our purpose. Therefore, we choose to mask random rectangular areas based on recorded coordinates. Specifically, the mask ratio  $r_i$  for  $x_i$  is calculated as follows:

$$r_i = \mu \times (1 - G_i) \quad (2)$$

where  $\mu$  is the basic mask ratio. The more likely a sample is to be noisy, the smaller the  $G_i$  is. Based on Eq. (2), the greater mask ratio will be adopted for this sample, thus stronger regularization will be imposed [15]. Subsequently, the aspect of the masked area  $\delta_i$  is generated as follows:

$$\delta_i \sim \text{Uniform}(\delta, \frac{1}{\delta}) \quad (3)$$

where  $\delta$  is a hyper-parameter for generating the aspect. Based on  $r_i$  and  $\delta_i$ , the height and width of the masked area are given by:

$$\begin{cases} h_i^{up} = \max \left( h_i^{\max} - \sqrt{\frac{(H_x \times W_x) \times r_i \times \delta_i}{4}}, 0 \right) \\ h_i^{dn} = \min \left( h_i^{\max} + \sqrt{\frac{(H_x \times W_x) \times r_i \times \delta_i}{4}}, H_x \right) \\ w_i^{lt} = \max \left( w_i^{\max} - \sqrt{\frac{(H_x \times W_x) \times r_i}{4\delta_i}}, 0 \right) \\ w_i^{rt} = \min \left( w_i^{\max} + \sqrt{\frac{(H_x \times W_x) \times r_i}{4\delta_i}}, W_x \right) \end{cases} \quad (4)$$

The selected locations in the masked area of  $x_i$  are assigned random values generated from a uniform distribution. Therefore, the adversarial noisy masked image  $x_i^m$

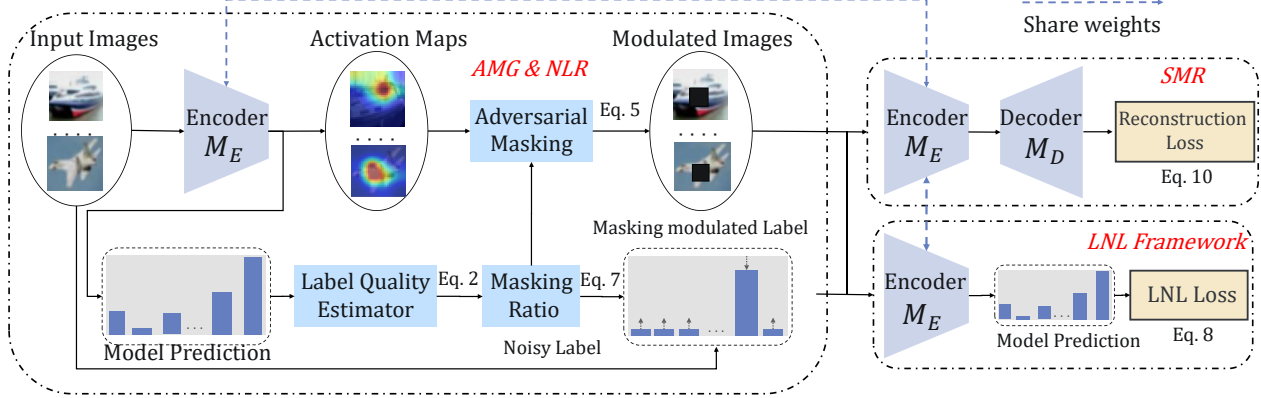


Figure 3. The framework of SANM. Three components are included: AMG (Adversarial Masking Generation), NLR (Noisy Label Regularization), SMR (Self-supervised Masking Reconstruction). In AMG and NLR, firstly feed the images to an encoder and generate activation maps. And a label quality guided adversarial masking strategy is proposed to modulate the images and noisy labels simultaneously. Further, an auxiliary decode branch is designed in SMR to reconstruct input images from the features of masked images. Finally, the generated modulated images and labels of SANM together with the reconstruction loss can be directly adopted for existing LNL frameworks.

can be derived as follows:

$$x_i^m(m, n) = \begin{cases} U(0, 1), & \text{if } m \in [h_i^{up}, h_i^{dn}], n \in [w_i^{lt}, w_i^{rt}], \\ x_i(m, n), & \text{otherwise.} \end{cases} \quad (5)$$

where  $U$  denotes the Uniform distribution. Similar operation is performed on the regions of minimum activation value with the coordinates  $\{h_i^{min}, w_i^{min}\}$  as well. Different from Eq. (2), since the regions of minimum activation value for correct-labeled samples are mainly focused on the background, a smaller mask ratio is adopted for them, helping models narrow the search scope for targets. Therefore, we set the mask ratio of the regions with minimum activation values as  $\mu \times G_i$ .

**Noisy Label Regularization.** To penalize the overconfident predictions of the network, a contrastive Noisy Label Regularization (NLR) strategy is designed. Specifically, the generated  $x_i^m$  is fed into  $M_E$  to obtain the prediction  $\tilde{y}_i^m$  and feature  $\mathbf{f}_i^m$  as follows:

$$\mathbf{f}_i^m = M_E(x_i^m; \theta_E) \quad (6)$$

where  $\theta_E$  denotes the parameters of the encoder  $M_E$ . Specifically, we give the pseudo labels of  $x_i^m$  with the predictions of  $x_i$ . Since the mask generation is based on the activation maps of the class with the largest predictions, the corresponding prediction is penalized with the mask ratio  $r_i$ . Besides, there exists an equal possibility for the masked region to be predicted as each class. Thus the prediction values for all the classes are increased uniformly. To sum up, the regularized pseudo label  $y_i^r$  for  $x_i^m$  can be derived as follows:

$$y_i^r(j) = \begin{cases} y_i(j) - r_i + r_i/c, & \text{if } j = \text{argmax}(y_i(j)) \\ y_i(j) + r_i/c, & \text{otherwise} \end{cases}, \quad (7)$$

where  $0 \leq j < c$ . Afterward,  $y_i^r$  and  $\tilde{y}_i^m$  are used for supervised loss calculation as:

$$\mathcal{L}_c = \mathcal{L}_{ce}(\tilde{y}_i^m, y_i^r) \quad (8)$$

where  $\mathcal{L}_{ce}$  denotes the cross-entropy loss.

### 3.2. Self-supervised Masking Reconstruction

Inspired by [11], a Self-supervised Masking Reconstruction branch (SMR) is designed to impose implicit regularization by providing extra noise-free supervised signal. Specifically, the decoder  $M_D$  reconstructs the original image from the latent representation of the masked image  $x_i^m$  by predicting the pixel values for the masked patch, which is defined as:

$$x_i^r = M_D(\mathbf{f}_i^m; \theta_D) \quad (9)$$

where  $\theta_D$  is the parameters of the decoder model  $M_D$ . Then the mean squared error (MSE) loss is computed between the reconstructed image  $x_i^r$  and the original unmasked image  $x_i$  in the pixel space, which is shown as follows:

$$\mathcal{L}_r = \|x_i^r - x_i\|^2 \quad (10)$$

In this way, a challenging self-supervised pretext task that requires holistic understanding beyond low-level image statistics is created for the model. This supervision signal is not only free of interference from noisy labels but can also enhance the model generalization ability. Therefore the overall loss function for SANM is constructed as follows:

$$\mathcal{L}_{\text{train}} = \mathcal{L}_c + \beta \mathcal{L}_r \quad (11)$$

where  $\beta$  is a parameter to balance the weight of two loss terms. Putting this all together, Algorithm 1 delineates the proposed SANM. Firstly, the original images are masked

Table 1. Comparison with state-of-the-art methods on CIFAR-10/100 datasets with symmetric noise.

Dataset Method/Noise ratio	CIFAR-10				CIFAR-100			
	20%	50%	80%	90%	20%	50%	80%	90%
Cross-Entropy (CE)	86.8	79.4	62.9	42.7	62.0	46.7	19.9	10.1
Co-teaching+ [41]	89.5	85.7	67.4	47.9	65.6	51.8	27.9	13.7
Mixup [43]	95.6	87.1	71.6	52.2	67.8	57.3	30.8	14.6
PENCIL [40]	92.4	89.1	77.5	58.9	69.4	57.5	31.1	15.3
Meta-Learning [16]	92.9	89.3	77.4	58.7	68.5	59.2	42.4	19.5
M-correction [1]	94.0	92.0	86.8	69.1	73.9	66.1	48.2	24.3
DivideMix [15]	96.1	94.6	93.2	76.0	77.3	74.6	60.2	31.5
C2D [47]	96.3	95.2	94.4	93.5	78.6	76.4	67.7	58.7
AugDesc [22]	96.3	95.4	93.8	91.9	79.5	77.2	66.4	41.2
GCE [5]	90.0	89.3	73.9	36.5	68.1	53.3	22.1	8.9
Self-CL+ [19]	95.5	93.9	89.2	81.9	76.5	72.4	59.6	48.8
MOIT+ [23]	94.1	91.8	81.1	74.7	75.9	70.6	47.6	41.8
SANM(DivideMix)	96.4	95.8	94.6	92.3	81.2	78.2	68.7	43.5
SANM(C2D)	<b>96.6</b>	<b>96.4</b>	<b>95.7</b>	<b>95.1</b>	<b>81.9</b>	<b>79.3</b>	<b>71.6</b>	<b>61.9</b>

with the label quality guided masking scheme, where mis-labeled and correct-labeled samples are modulated adaptively. Meanwhile, the predictions of the original images are used for regularized label generation. Then the regularized features of the masked images are utilized for self-supervised masking reconstruction, which provides noise-free supervisions. Finally, SANM can be plugged into existing LNL frameworks to further boost performance.

## 4. Experiments

### 4.1. Datasets and Implementation Details

**Simulated noisy datasets.** We evaluate on two simulated noisy datasets: CIFAR-10/100 [14], which consists of 50k training samples and 10k test samples. Following previous work [15], two simulated noisy settings are considered: symmetric and asymmetric noise. Following noise rates are considered:  $\tau \in \{20\%, 50\%, 80\%, 90\%\}$  for symmetric noise and  $\tau \in \{20\%, 40\%\}$  for asymmetric noise. An 18-layer PreAct ResNet is used and we train for 300 epochs with a batch size of 64 and SGD optimizer. The initial learning rate is set as 0.02 and is reduced by a factor of 10 after 150 epochs.  $\beta$  is empirically set as 1.

**Real-world noisy datasets.** We also validate our method on real-world noisy datasets Clothing-1M [38] and Animal-10N [28], the former contains 1 million training samples with nearly 40% mislabeled samples, the latter contains 55,000 human-labeled online images for 10 confusing animals. It includes approximately 8% noisy-labeled samples. The original protocol is used to split the training and test images. We strictly follow the previous experimental settings, more details can be found in the supplementary materials.

---

### Algorithm 1 The proposed SANM framework

---

**Input:** Noisy training set  $D$ , encoder model  $M_E(\cdot; \theta_E)$ , decoder model  $M_D(\cdot; \theta_D)$ , batch size  $b$ , max iterations  $m$ , basic mask ratio  $\mu$ .

**Procedure:**

- 1: **for**  $i = 1$  to  $m$  **do**
- 2:    $\{x_i, y_i\}_{i=1}^b \leftarrow \text{SampleMiniBatch}(D, b)$ .
- 3:   Feed  $\{x_i\}_{i=1}^b$  into  $M_E$  and generate feature maps  $\{F_i\}_{i=1}^b$  and predictions  $\{\tilde{y}_i\}_{i=1}^b$ .
- 4:   Generate activation maps  $\{A_i\}_{i=1}^b$  by Eq. (1).
- 5:   Calculate mask ratios  $\{r_i\}_{i=1}^b$  and adversarial masked images  $\{x_i^m\}_{i=1}^b$  by Eq. (2-5).
- 6:   Feed  $\{x_i^m\}_{i=1}^b$  into  $M_E$  and generate predictions  $\{\tilde{y}_i^m\}_{i=1}^b$  and features  $\{f_i^m\}_{i=1}^b$  by Eq. (6).
- 7:   Calculate the regularized labels  $\{y_i^r\}_{i=1}^b$  by Eq. (7).
- 8:   Calculate cross-entropy loss  $\mathcal{L}_c$  by Eq. (8).
- 9:   Feed  $\{f_i^m\}_{i=1}^b$  into  $M_D$  and generate reconstructed images  $\{x_i^r\}_{i=1}^b$  by Eq. (9).
- 10:   Calculate self-supervised reconstruction loss  $\mathcal{L}_r$  by Eq. (10).
- 11:   Update parameter  $\theta_E, \theta_D$  in backward process.
- 12: **end for**

**Output:** The final encoder  $M_E(\cdot; \theta_E)$ .

---

### 4.2. Experimental Results

• **Results on Simulated Noisy Datasets.** Here we compare the proposed method with multiple baselines on CIFAR-10/100 using different symmetric noise levels. As shown in Tab. 1, by integrating with C2D [47], our SANM outperforms the naive C2D significantly and achieves state-of-the-art performance among these methods. Specifically, an increase of around 3 ~ 4% across all the symmetric noise cases on CIFAR-100 can be observed. Besides, when combined with DivideMix [15], SANM can bring greater improvement to its performance, especially for severe label noise cases. Under 90% symmetric noise setting, our SANM (DivideMix) outperforms DivideMix by a large

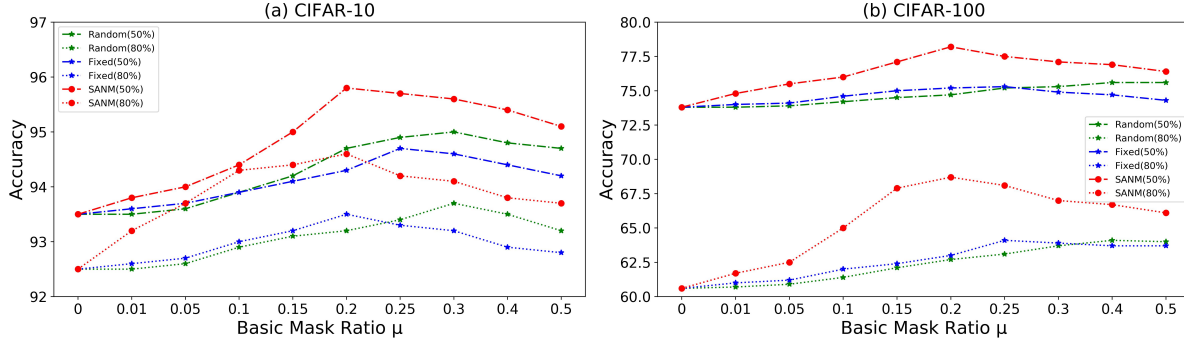


Figure 4. Investigation of the basic mask ratio  $\mu$  under 50% and 80% symmetric noise of CIFAR-10/100.

Table 2. Asymmetric noise on CIFAR-10.

Method	Noisy ratio	
	20%	40%
Joint-Optim [32]	92.8	91.7
PENCIL [40]	92.4	91.2
F-correction [24]	89.9	-
Distilling [46]	92.7	90.2
Meta-Learning [16]	-	88.6
M-correction [1]	-	86.3
Iterative-CV [3]	-	88.0
DivideMix [15]	93.4	93.4
REED [44]	95.0	92.3
C2D [47]	93.8	93.4
Sel-CL+ [19]	95.2	93.4
GCE [5]	87.3	78.1
RRL [17]	-	92.4
SANM(DivideMix)	<b>95.4</b>	<b>94.8</b>

Table 3. Testing accuracy on Clothing-1M.

Method	Acc
CrossEntropy	69.21
F-correction [24]	69.84
M-correction [1]	71.00
Joint-Optim [32]	72.16
Meta-Cleaner [45]	72.50
Meta-Learning [16]	73.47
PENCIL [40]	73.49
Self-Learning [10]	74.45
DivideMix [15]	74.76
Nested [4]	74.90
AugDesc [22]	75.11
RRL [17]	74.90
GCE [5]	73.30
C2D [47]	74.30
SANM(DivideMix)	<b>75.63</b>

Table 4. Comparison with masks generated from pre-trained backbones on CIFAR-10/100. M: Method. P: Pretrained Backbone. S: SANM. D: DivideMix. C: C2D.

Dataset			CIFAR-10			
M	P	S	20%	50%	80%	90%
D	✓	✗	96.0	95.1	93.7	81.5
	✗	✓	<b>96.4</b>	<b>95.8</b>	<b>94.6</b>	<b>92.3</b>
C	✓	✗	96.3	95.4	95.0	93.9
	✗	✓	<b>96.6</b>	<b>96.4</b>	<b>95.7</b>	<b>95.1</b>
Dataset			CIFAR-100			
M	P	S	20%	50%	80%	90%
D	✓	✗	78.8	76.6	64.9	36.4
	✗	✓	<b>81.2</b>	<b>78.2</b>	<b>68.7</b>	<b>43.5</b>
C	✓	✗	80.0	77.1	69.2	58.4
	✗	✓	<b>81.9</b>	<b>79.3</b>	<b>71.6</b>	<b>61.9</b>

margin of 12%. A similar performance boost can be observed on the CIFAR-100 as well, especially in the high noise case. When compared with the implicit regularization method AugDesc [22], our SANM (DivideMix) can also achieve much better accuracy across all the settings. Tab. 2 shows comparisons of the test accuracy on CIFAR-10 with varied asymmetric noise levels. SANM shows a consistent improvement with respect to the state-of-the-art methods, demonstrating the ability of SANM in enhancing the robustness of the model to semantic noise.

• **Results on Real-world Noisy Datasets.** The results on Clothing1M dataset are shown in Tab. 3, one can see that SANM can bring around 0.87% improvement in accuracy to the baseline method DivideMix, and get much better result than other state-of-the-art methods, verifying that SANM is also effective in real-world noise scenarios. Moreover, SANM obtains 89.3% accuracy on Animal-10N (detailed comparison on Animal-10N is in supplementary materials), which sets the new state-of-the-art, indicating that SANM can handle fine-grained real-world noisy datasets as well.

**Generality of SANM.** To further verify the effectiveness and generalization of SANM, we adopt vanilla CE, Co-teaching [9], CDR [37] and ELR+ [21], other than DivideMix and C2D, to work collaboratively with SANM. As shown in Tab. 6, SANM brings consistent performance

boost to all the corresponding baselines across all the noisy cases. For example, when seamlessly combined with SANM, an accuracy increase of around 14.9% on the symmetric 80% noise setting (CIFAR-100) can be achieved by vanilla CE, demonstrating that SANM is capable of effectively imposing adaptive regularization on noisy and clean samples respectively and thus alleviates over-fitting to noisy labels. Therefore, all the results indicate the generalization ability of SANM to boost existing LNL approaches.

### 4.3. Ablation Study

**Component Analysis.** To explore the impact of each component in SANM, i.e., AMG-adversarial noisy masking generation, NLR-noisy label regularization, SMR-self-supervised masking reconstruction, we conduct an ablation study on CIFAR-10 and CIFAR-100. As shown in Tab. 5, we report test accuracy for all the noisy settings for SANM where the contribution of each component is clearly shown. Specifically, it can be observed that AMG is crucial to performance, especially under severe label noisy cases. Besides, the accuracy is further improved when further combined with NLR. By introducing noise-free self-supervised signals, SMR can also boost the performance of the baseline across all the settings. Best accuracy can be achieved by SANM when all the components are included. This ver-

Table 5. Ablation study for the effectiveness of each key component. AMG: adversarial noisy masking generation, NLR: noisy label regularization, SMR: self-supervised masking reconstruction.

Component				CIFAR-10				CIFAR-100			
AMG	NLR	SMR		20%	50%	80%	90%	20%	50%	80%	90%
✗	✗	✗	Best	96.1	94.6	93.2	76.0	77.3	74.6	60.2	31.5
			Last	95.7	94.4	92.9	75.4	76.9	74.2	59.6	31.0
✓	✗	✗	Best	96.3	95.3	94.0	91.4	80.2	77.3	68.0	42.7
			Last	96.2	95.1	93.6	90.8	79.7	77.0	67.5	42.5
✓	✓	✗	Best	96.4	95.5	94.2	91.6	80.5	77.5	68.3	43.0
			Last	96.3	95.4	94.0	91.5	80.2	77.1	68.2	42.7
✓	✓	✓	Best	<b>96.4</b>	<b>95.8</b>	<b>94.6</b>	<b>92.3</b>	<b>81.2</b>	<b>78.2</b>	<b>68.7</b>	<b>43.5</b>
			Last	<b>96.3</b>	<b>95.6</b>	<b>94.3</b>	<b>92.1</b>	<b>80.4</b>	<b>78.0</b>	<b>68.3</b>	<b>43.0</b>

Table 6. Comparison between the LNL methods and their SANM applications with symmetric noise on CIFAR-10/100. Specifically, the 9-layer CNN is adopted as the backbone network of Co-teaching.

Dataset		CIFAR-10				CIFAR-100			
Method/Noise ratio		20%	50%	80%	90%	20%	50%	80%	90%
CE	Best	86.8	79.4	62.9	42.7	62.0	46.7	19.9	10.1
SANM(CE)	Best	<b>92.4</b>	<b>89.7</b>	<b>72.1</b>	<b>51.5</b>	<b>70.9</b>	<b>53.1</b>	<b>34.8</b>	<b>18.6</b>
Co-teaching [9]	Best	82.6	73.0	24.0	14.6	50.5	38.2	11.8	4.9
SANM(Co-teaching)	Best	<b>89.2</b>	<b>78.2</b>	<b>36.4</b>	<b>20.7</b>	<b>58.2</b>	<b>51.3</b>	<b>19.4</b>	<b>13.4</b>
CDR [37]	Best	90.4	85.0	47.2	12.3	63.3	39.5	29.2	8.0
SANM(CDR)	Best	<b>92.6</b>	<b>91.6</b>	<b>55.3</b>	<b>16.7</b>	<b>72.7</b>	<b>56.4</b>	<b>36.6</b>	<b>20.8</b>
ELR+ [21]	Best	94.6	93.8	91.1	75.2	77.5	72.4	58.2	30.8
SANM(ELR+)	Best	<b>96.3</b>	<b>95.7</b>	<b>94.1</b>	<b>82.9</b>	<b>79.8</b>	<b>77.3</b>	<b>65.0</b>	<b>38.7</b>

ifies the effectiveness of all three key designs in SANM.

**Parameter Sensitivity.** The basic mask ratio  $\mu$  in Eq. (2) is a important parameter within the proposed method. To investigate how  $\mu$  affects the performance, we conduct parameter analysis experiments under 50%/80% symmetric noise levels on CIFAR-10/100. Best test accuracy with respect to different  $\mu$  is shown in Fig. 4. Compared with the results of baseline methods, i.e., masking the images with fixed ratios and random ratios (sampled from  $Uniform(0, \mu)$ ) for each sample, SANM achieves superior performance. Moreover, the results show that keeping  $\mu$  between 0.1 and 0.3 leads to better performance than without adversarial noisy masking ( $\mu = 0$ ). This is mainly because large  $\mu$  makes the masked training images and the test images differ significantly in image style, and thus the model is unable to learn class-specific discriminative information when trained with such strong regularization. However, an extremely small  $\mu$  results in encoder overfit to noisy labels more severely. Therefore, we set  $\mu$  as 0.2 to obtain the best performance.

**Comparison with Masks from Pre-trained Backbones.** The effectiveness of the proposed adversarial masking generation is evaluated in Tab. 5. And further we conduct experiments that utilizing a strong backbone (i.e., pre-trained ResNet-18 with DivideMix and C2D) to generate more object-centric masks as comparison to our method. The results shown in Tab. 4 indicate that even compared with the mask generation strategy by well-trained strong backbones, our adversarial masking generation can consistently

performs better since the mask generation strategy of fixed well-trained backbones cannot adaptively adjust the mask strategy according to the training state of the base model, leading to sub-optimal results.

## 5. Conclusion

In this paper, we propose SANM, a novel self-supervised adversarial noisy masking LNL method. SANM prevents models from overfitting noisy instances by explicitly imposing regularization on the features. A label quality guided adversarial masking scheme is designed to adaptively modify the inputs and noisy labels. Moreover, an auxiliary decoder is introduced to reconstruct images based on the modulated features, which enhances the generalization by providing extra noise-free self-supervised signals. Besides, SANM is flexible and can be easily plugged into existing LNL frameworks. State-of-the-art performance is obtained by SANM on both synthetic and real-world noisy datasets.

## Acknowledgement

This work was supported by National Natural Science Fund of China (62076184, 61976158, 61976160, 62076182, 62276190), in part by Fundamental Research Funds for the Central Universities and State Key Laboratory of Integrated Services Networks (Xidian University), in part by Shanghai Innovation Action Project of Science and Technology (20511100700) and Shanghai Natural Science Foundation (22ZR1466700).



## References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. In *Proc. International Conference on Machine Learning (ICML)*, 2019. 6, 7
- [2] Devansh Arpit, Stanislaw Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *Proc. International Conference on Machine Learning (ICML)*, pages 233–242. PMLR, 2017. 1, 2
- [3] Pengfei Chen, Benben Liao, Guangyong Chen, and Shengyu Zhang. Understanding and utilizing deep neural networks trained with noisy labels. In *Proc. International Conference on Machine Learning (ICML)*, 2019. 1, 3, 7
- [4] Yingyi Chen, Xi Shen, Shell Xu Hu, and Johan A. K. Suykens. Boosting co-teaching with compression regularization for label noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 2688–2692, June 2021. 7
- [5] Aritra Ghosh and Andrew Lan. Contrastive learning improves model robustness under label noise. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2703–2708, 2021. 6, 7
- [6] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 1, 3
- [7] Aayush Gupta, Ayush Jaiswal, Yue Wu, Vivek Yadav, and Pradeep Natarajan. Adversarial mask generation for preserving visual privacy. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, pages 1–5. IEEE, 2021. 3
- [8] 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 *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, 2018. 1, 2, 3
- [9] 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 *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 7, 8
- [10] Jiangfan Han, Ping Luo, and Xiaogang Wang. Deep self-learning from noisy labels. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 7
- [11] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv:2111.06377*, 2021. 5
- [12] Qibin Hou, PengTao Jiang, Yunchao Wei, and Ming-Ming Cheng. Self-erasing network for integral object attention. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, 2018. 3
- [13] Tao Hu, Honggang Qi, Qingming Huang, and Yan Lu. See better before looking closer: Weakly supervised data augmentation network for fine-grained visual classification. *arXiv preprint arXiv:1901.09891*, 2019. 3
- [14] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 3, 6
- [15] Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *Proc. International Conference on Learning Representations (ICLR)*, 2020. 1, 2, 3, 4, 6, 7
- [16] Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S. Kankanhalli. Learning to learn from noisy labeled data. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 6, 7
- [17] Junnan Li, Caiming Xiong, and Steven CH Hoi. Learning from noisy data with robust representation learning. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, pages 9485–9494, 2021. 7
- [18] Kunpeng Li, Ziyang Wu, Kuan-Chuan Peng, Jan Ernst, and Yun Fu. Tell me where to look: Guided attention inference network. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9215–9223, 2018. 3
- [19] Shikun Li, Xiaobo Xia, Shiming Ge, and Tongliang Liu. Selective-supervised contrastive learning with noisy labels. *arXiv preprint arXiv:2203.04181*, 2022. 6, 7
- [20] Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 20331–20342, 2020. 1, 2
- [21] Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 7, 8
- [22] Kento Nishi, Yi Ding, Alex Rich, and Tobias Hollerer. Augmentation strategies for learning with noisy labels. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8022–8031, 2021. 6, 7
- [23] Diego Ortego, Eric Arazo, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Multi-objective interpolation training for robustness to label noise. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6606–6615, 2021. 6
- [24] 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 *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 7
- [25] Haim Permuter, Joseph Francos, and Ian Jermyn. A study of gaussian mixture models of color and texture features for image classification and segmentation. *Pattern recognition*, 39(4):695–706, 2006. 4
- [26] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *Proc. International Conference on Machine Learning (ICML)*, pages 4334–4343. PMLR, 2018. 1, 2
- [27] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. Selfie: Refurbishing unclean samples for robust deep learning. In *Proc. International Conference on Machine Learning (ICML)*, 2019. 1

- [28] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. SELFIE: Refurbishing unclean samples for robust deep learning. In *Proc. International Conference on Machine Learning (ICML)*, pages 5907–5915, 2019. 6
- [29] Hwanjun Song, Minseok Kim, Dongmin Park, Yooju Shin, and Jae-Gil Lee. Learning from noisy labels with deep neural networks: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 1, 2
- [30] Erik Stammes, Tom FH Runia, Michael Hofmann, and Mohsen Ghafoorian. Find it if you can: end-to-end adversarial erasing for weakly-supervised semantic segmentation. In *Thirteenth International Conference on Digital Image Processing (ICDIP 2021)*, volume 11878, page 1187825. International Society for Optics and Photonics, 2021. 3
- [31] Guolei Sun, Hisham Cholakkal, Salman Khan, Fahad Khan, and Ling Shao. Fine-grained recognition: Accounting for subtle differences between similar classes. In *Association for the Advancement of Artificial Intelligence (AAAI)*, volume 34, pages 12047–12054, 2020. 3
- [32] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 3, 7
- [33] Gabriele Valvano, Andrea Leo, and Sotirios A Tsaftaris. Re-using adversarial mask discriminators for test-time training under distribution shifts. *arXiv preprint arXiv:2108.11926*, 2021. 3
- [34] Hongxin Wei, Lei Feng, Xiangyu Chen, and Bo An. Combating noisy labels by agreement: A joint training method with co-regularization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13726–13735, 2020. 1, 3
- [35] Yunchao Wei, Jiashi Feng, Xiaodan Liang, Ming-Ming Cheng, Yao Zhao, and Shuicheng Yan. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1568–1576, 2017. 3
- [36] Yichen Wu, Jun Shu, Qi Xie, Qian Zhao, and Deyu Meng. Learning to purify noisy labels via meta soft label corrector. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2021. 1
- [37] Xiaobo Xia, Tongliang Liu, Bo Han, Chen Gong, Nannan Wang, Zongyuan Ge, and Yi Chang. Robust early-learning: Hindering the memorization of noisy labels. In *Proc. International Conference on Learning Representations (ICLR)*, 2021. 7, 8
- [38] Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 1, 6
- [39] Xiao Yang, Yinpeng Dong, Tianyu Pang, Hang Su, Jun Zhu, Yuefeng Chen, and Hui Xue. Towards face encryption by generating adversarial identity masks. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, pages 3897–3907, 2021. 3
- [40] Kun Yi and Jianxin Wu. Probabilistic end-to-end noise correction for learning with noisy labels. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7017–7025, 2019. 1, 3, 6, 7
- [41] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor W Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *Proc. International Conference on Machine Learning (ICML)*, 2019. 1, 3, 6
- [42] Boshen Zhang, Yuxi Li, Yuanpeng Tu, Jinlong Peng, Yabiao Wang, Cunlin Wu, Yang Xiao, and Cairong Zhao. Learning from noisy labels with coarse-to-fine sample credibility modeling. In *Computer Vision—ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 21–38. Springer, 2023. 3
- [43] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 1, 3, 6
- [44] Hui Zhang and Quanming Yao. Decoupling representation and classifier for noisy label learning. *arXiv preprint arXiv:2011.08145*, 2020. 7
- [45] Weihe Zhang, Yali Wang, and Yu Qiao. Metacleaner: Learning to hallucinate clean representations for noisy-labeled visual recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 7
- [46] Z. Zhang, H. Zhang, S. Arik, H. Lee, and T. Pfister. Distilling effective supervision from severe label noise. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 7
- [47] Evgenii Zheltonozhskii, Chaim Baskin, Avi Mendelson, Alex M Bronstein, and Or Litany. Contrast to divide: Self-supervised pre-training for learning with noisy labels. In *Proc. IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1657–1667, 2022. 1, 3, 6, 7
- [48] Guoqing Zheng, Ahmed Hassan Awadallah, and Susan Dumais. Meta label correction for noisy label learning. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2021. 1, 3
- [49] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2921–2929, 2016. 3, 4
- [50] Xiong Zhou, Xianming Liu, Chenyang Wang, Deming Zhai, Junjun Jiang, and Xiangyang Ji. Learning with noisy labels via sparse regularization. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, pages 72–81, 2021. 1, 2