

# Learning from Noisy Labels with Decoupled Meta Label Purifier

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>  
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 Univeristy, Shanghai

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

{2030809, zhaocairong}@tongji.edu.cn

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

## Abstract

Training deep neural networks (DNN) with noisy labels is challenging since DNN can easily memorize inaccurate labels, leading to poor generalization ability. Recently, the meta-learning based label correction strategy is widely adopted to tackle this problem via identifying and correcting potential noisy labels with the help of a small set of clean validation data. Although training with purified labels can effectively improve performance, solving the meta-learning problem inevitably involves a nested loop of bi-level optimization between model weights and hyper-parameters (i.e., label distribution). As compromise, previous methods resort to a coupled learning process with alternating update. In this paper, we empirically find such simultaneous optimization over both model weights and label distribution can not achieve an optimal routine, consequently limiting the representation ability of backbone and accuracy of corrected labels. From this observation, a novel multi-stage label purifier named DMLP is proposed. DMLP decouples the label correction process into label-free representation learning and a simple meta label purifier. In this way, DMLP can focus on extracting discriminative feature and label correction in two distinctive stages. DMLP is a plug-and-play label purifier, the purified labels can be directly reused in naive end-to-end network retraining or other robust learning methods, where state-of-the-art results are obtained on several synthetic and real-world noisy datasets, especially under high noise levels. Code is available at <https://github.com/yuanpengtu/DMLP>.

## 1. Introduction

Deep learning has achieved significant progress on various recognition tasks. The key to its success is the availabil-

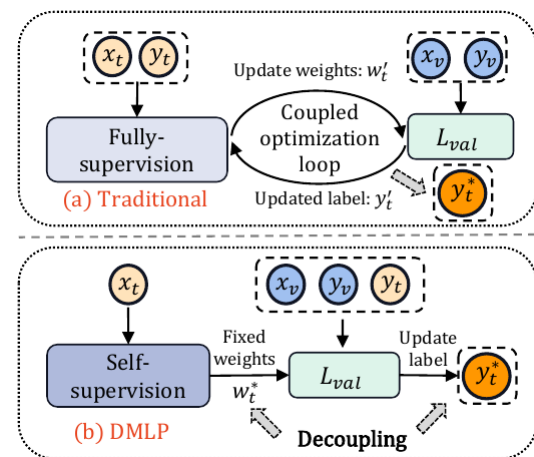


Figure 1. (a) Traditional coupled alternating update to solve meta label purification problem, and (b) the proposed DMLP method that decouples the label purification process into representation learning and a simple non-nested meta label purifier.

ity of large-scale datasets with reliable annotations. Collecting such datasets, however, is time-consuming and expensive. Easy ways to obtain labeled data, such as web crawling [31], inevitably yield samples with noisy labels, which is not appropriate to be directly utilized to train DNN since these complex models are vulnerable to memorize noisy labels [2].

Towards this problem, numerous Learning with Noisy Label (LNL) approaches were proposed. Classical LNL methods focus on identifying the noisy samples and reducing their effect on parameter updates by abandoning [12] or assigning smaller importance. However, when it comes to extremely noisy and complex scenarios, such scheme struggles since there is no sufficient clean data to train a discriminative classifier. Therefore, label correction approaches are proposed to augment clean training samples by revising noisy labels to underlying correct ones. Among them, meta-learning based approaches [9, 16, 25] achieve state-of-the-art performance via resorting to a small clean validation set and taking noisy labels as hyper-parameters, which provides sound guidance

\* Equal contribution.

† Corresponding author.

toward underlying label distribution of clean samples. However, such meta purification inevitably involves a nested bi-level optimization problem on both model weight and hyper-parameters (shown as Fig. 1 (a)), which is computationally infeasible. As a compromise, the alternating update between model weights and hyper-parameters is adopted to optimize the objective [9, 16, 25], leading to a coupled solution for representation learning and label purification.

**Empirical observation.** Intuitively, alternate optimization over a large search space (model weight and hyper-parameters) may lead to sub-optimal solutions. To investigate how such approximation affects results in robust learning, we conduct empirical analysis on CIFAR-10 [14] with recent label purification methods MLC [41] and MSLC [9], which consist of a deep model and a meta label correction network, and make observation as Fig. 2.

• **Coupled optimization hinders quality of corrected labels.** We first compare the *Coupled* meta corrector MLC with its extremely *Decoupled* variant where the model weights are first optimized for 70 epochs with noisy labels and get fixed, then labels are purified with the guidance of validation set. We adopt the accuracy of corrected label to measure the performance of purification. From Fig. 2 (a), we can clearly observe that compared with *Decoupled* counterpart, joint optimization yields inferior correction performance, and these miscorrection will reversely affect the representation learning in coupled optimization.

• **Coupled optimization hinders representation ability.** We investigate the representation quality by evaluating the linear prob accuracy [6] of extracted feature in Fig. 2 (b). We find the representation quality of *Coupled* training is much worse at the beginning, which leads to slow and unstable representation learning in the later stage. To further investigate the effect on representation learning, we also resort to a well pretrained backbone with self-supervised learning [5] as initialization, recent research [40] shows pretrained representation is substantially helpful for LNL framework. However, we find this conclusion does not strictly hold for coupled meta label correctors. As shown in Fig. 2 (c), by comparing the classification accuracy from classifier of MLC/MSLC, we observe the pretrained model only brings marginal improvement if model weights is still coupled with hyper-parameters. In contrast, when the weight of backbone is fixed and decoupled from the label purification and classifier, the improvement becomes more significant.

**Decoupled Meta Purification.** From the observation above, we find the decoupling between model weights and hyperparameters of meta correctors is essential to label accuracy and final results. Therefore, in this paper, we aim at detaching the meta label purification from representation learning and designing a simple meta label purifier which is more friendly to optimization of label distribution problem than existing complex meta networks [9, 41]. Hence

we propose a general multi-stage label correction strategy, named Decoupled Meta Label Purifier (DMLP). The core of DMLP is a meta-learning based label purifier, however, to avoid solving the bi-level optimization with a coupled solution, DMLP decouples this process into self-supervised representation learning and a linear meta-learner to fit underlying correct label distribution (illustrated as Fig. 1 (b)), thus simplifies the label purification stage as a single-level optimization problem. The simple meta-learner is carefully designed with two mutually reinforcing correcting processes, named intrinsic primary correction (IPC) and extrinsic auxiliary correction (EAC) respectively. IPC plays the role of purifying labels in a global sense at a steady pace, while EAC targets at accelerating the purification process via looking ahead (i.e., training with) the updated labels from IPC. The two processes can enhance the ability of each other and form a positive loop of label correction. Our DMLP framework is flexible for application, the purified labels can either be directly applied for naive end-to-end network retraining, or exploited to boost the performance of existing LNL frameworks. Extensive experiments conducted on mainstream benchmarks, including synthetic (noisy versions of CIFAR) and real-world (Clothing1M) datasets, demonstrate the superiority of DMLP. In a nutshell, the key contributions of this paper include:

- We analyze the necessity of decoupled optimization for label correction in robust learning, based on which we propose DMLP, a flexible and novel multi-stage label purifier that solves bi-level meta-learning problem with a decoupled manner, which consists of representation learning and non-nested meta label purification;
- In DMLP, a novel non-nested meta label purifier equipped with two correctors, IPC and EAC is proposed. IPC is a global and steady corrector, while EAC accelerates the correction process via training with the updated labels from IPC. The two processes form a positive training loop to learn more accurate label distribution;
- Deep models trained with purified labels from DMLP achieve state-of-the-art results on several synthetic and real-world noisy datasets across various types and levels of label noise, especially under high noise levels. Extensive ablation studies are provided to verify the effectiveness.

## 2. Related Works

The existing LNL approaches that are related to our work can be coarsely categorized into two groups: noisy sample detection and label correction.

**Noisy sample detection** methods aim to identify and reduce the importance of suspicious false-labeled samples during training. The detected noisy samples are abandoned [12, 23], assigned with smaller weights via a sample re-weight training scheme [20, 37], or used to formulate a semi-supervised learning problem by throwing away the




Figure 2. The corrected label accuracy (a) and linear probe accuracy of representations (b) between coupled [41] and decoupled label correction schemes during training under 50% and 90% symmetric label noise on CIFAR-10. (c) investigates the effect of representation learning on ‘Ori’-original coupled network training from scratch, ‘SimCLR-Coupled’-initialization with stronger self-supervised pretrained weights and ‘SimCLR-Decoupled’-further fix the pretrained backbone during label purification.

labels while keeping the unlabeled data [15, 38]. These methods show robustness under certain noise levels, but struggle when it comes to extremely noisy and complex scenarios since there is no sufficient clean data to train a classifier.

**Label correction** approaches attempt to augment the training set by finding and correcting noisy labels to their underlying true ones. To do so, some works [22] try to estimate the noise transition matrix. However, these methods usually assume that the noise type is class-dependent, which may be inappropriate for more complex noise settings, such as real-world noisy datasets [31]. Some other works resort to exploiting the prediction of the network, both soft [1, 13, 24, 35] and hard [27, 28] label correction schemes are designed. However, the predictions of over-parameterized backbone network can be unreliable since it tends to fluctuate during training in the presence of false-labeled data ([38]). Another line of works utilize robust representations learned via unsupervised contrastive learning methods [11, 17, 38, 40] to eliminate the interference of noisy labels, which provides a reliable initialization of deep models. Recently, meta-learning based methods [9, 16, 25, 41] show great potential towards LNL problems with the help of a small clean validation set to provide sound guidance toward underlying label distribution of clean sample. However, these approaches involves a bi-level optimization problem on model weights and hyper-parameters, which is too computationally expensive to optimize. As a compromise, the one-step approximation is commonly adopted [9, 41] to convert the nested objective into a coupled update procedure between model weights and hyper-parameters, leading to sub-optimal performance.

Accordingly, DMLP belongs to the label correction group via meta-learning strategy, but unlike previous meta-learning based methods [9, 41], the learning process on model weights and labels are decoupled into individual stages within DMLP. Next, we will introduce the proposed DMLP in detail.

### 3. Method

#### 3.1. Decoupled Solution to Meta Label Purification

For convenience, notations within DMLP are clarified first. The noisy training dataset is denoted as  $D_t = \{(x_i, y_i) | 1 \leq i \leq N_t\}$ , where  $x_i \in \mathbb{R}^{H \times W \times 3}$ ,  $y_i \in \{0, 1\}^c$  are the image and corrupted label of the  $i$ -th instance,  $c$  is the class number. Similarly, a small clean validation dataset is denoted as  $D_v = \{(x_i, y_i) | 1 \leq i \leq N_v\}$ .  $N$  denotes the dataset size, and  $N_v \ll N_t$ . The subscripts  $t, v$  represent the data is from training set or validation set.

Typical meta-learning based purification requires solving a bi-level optimization problem on both model weights and hyper-parameters with objective:

$$\begin{aligned} & \min_{\theta_\alpha} E_{(x_v, y_v) \in D_v} \mathcal{L}_{val}(x_v, y_v; w^*(\theta_\alpha)), \\ \text{s.t. } & w^*(\theta_\alpha) = \arg \min_w E_{(x_t, y_t) \in D_t} \mathcal{L}_{train}(x_t, y_t; w, \theta_\alpha), \end{aligned} \quad (1)$$

where  $w$  and  $\theta_\alpha$  denote the model weights and the meta hyper-parameters respectively. In the classical label purification pipeline,  $\theta_\alpha$  is reparameterized as a function of noisy label distribution, i.e.  $\theta_\alpha = g(y_t)$ .  $\mathcal{L}_{val}$  and  $\mathcal{L}_{train}$  are the loss function on different datasets. Since both loss terms are *not analytical* and involve complicated forward pass of DNN, solving the nested optimization objective is computationally expensive. Meanwhile, alternating one-step approximation [9, 25] cannot guarantee that the optimization direction is optimal due to coupled update.

In contrast, in order to avoid coupled optimization over the large searching space of network parameters  $w$ , we reformulate the meta-learning objective in DMLP as:

$$\begin{aligned} & \min_{y_t} E_{(x_v, y_v) \in D_v} \mathcal{L}_{val}(\mathbf{f}_v, y_v; w^*(y_t)), \\ \text{s.t. } & w^*(y_t) = \arg \min_w E_{(x_t, y_t) \in D_t} \mathcal{L}_{train}(\mathbf{f}_t, y_t; w), \quad (2) \\ & \mathbf{f}_t = G(x_t; \theta_G^*), \quad \mathbf{f}_v = G(x_v; \theta_G^*), \end{aligned}$$

where a pretrained feature extractor  $G: \mathbb{R}^{H \times W \times 3} \rightarrow \mathbb{R}^d$  is designed to extract  $d$ -dimensional image representation  $\mathbf{f}$ . In

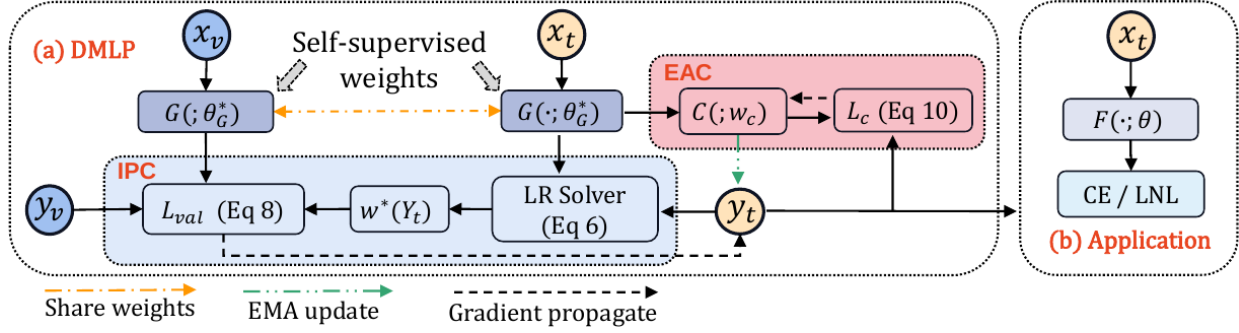


Figure 3. The overall framework of DMLP. (a) Decoupled meta label purification (Sec. 3.2). (b) The purified labels can be applied in normal network retraining (CE) or other LNL loss (Sec. 3.3).

order to obtain representation of high quality, these extracted features are utilized in a contrastive self-supervised learning framework [5, 6] to update the parameters  $\theta_G$  as pretraining without noisy labels. Subsequently, the established feature extractor  $\mathbf{f} = G(x; \theta_G^*)$  can learn a noise-agnostic descriptor of image data, which is also highly separable in high dimensional feature space [38]. In this manner, we detach the representation learning from noisy label correction, while keeping strong separability of features.

Further, since the feature is representative and separable, the loss term  $\mathcal{L}_{train}$  and  $\mathcal{L}_{val}$  can be formulated with simple risk estimation functions (e.g. linear discrimination) instead of complex DNN forward pass, making it possible to solve the problem of Eq. (2) in a non-nested manner with an *analytical solution*, which will be introduced in Sec. 3.2.

### 3.2. Non-nested Meta Label Purifier

To solve the purification problem of Eq. (2), we propose two mutually reinforced solutions to seek purified training labels as shown in Fig. 3, the **intrinsic primary correction** (IPC) and **extrinsic auxiliary correction** (EAC) processes.

**Intrinsic Primary Correction.** IPC aims at performing global-wise label purification at a slow and steady pace. Specifically, as shown in Fig. 3 (a), the features and labels of a batch of  $b$  training data are gathered into matrix  $F_t \in \mathbb{R}^{b \times d}$  and  $Y_t \in \mathbb{R}^{b \times c}$  as:

$$F_t = [\mathbf{f}_{t,1}, \mathbf{f}_{t,2}, \dots, \mathbf{f}_{t,b}]^T, \quad Y_t = [y_{t,1}, y_{t,2}, \dots, y_{t,b}]^T. \quad (3)$$

Since the feature descriptors are representative, we assume there exists a simple linear estimation transform  $w \in \mathbb{R}^{d \times c}$ , which accurately predict the categorical distribution with ridge regression:

$$\min_w \mathcal{L}_{train}(F_t, Y_t; w) = \|\sigma(\alpha Y_t) - F_t w\|^2 + \lambda \|w\|^2, \quad (4)$$

where  $\sigma(\cdot)$  is the softmax function along the categorical dimension to meet the normalization constraint, and  $\alpha$  is a scaling factor. By solving the linear regression problem of Eq. (4) through least square method (see supplementary material for details), we can obtain its closed-form solution

$w^*(Y_t)$  on the training batch  $(F_t, Y_t)$  and derive its optimal prediction via linear regression on samples  $\mathbf{f}_{v,i}$  from validation set  $D_v$ :

$$\begin{aligned} y'_{v,i}(Y_t) &= w^*(Y_t)^T \mathbf{f}_{v,i} \\ &= \sigma(\alpha Y_t)^T F_t (F_t^T F_t + \lambda I)^{-1} \mathbf{f}_{v,i}. \end{aligned} \quad (5)$$

Intuitively, the discrepancy between the predicted results  $y'_{v,i}$  and the ground truth labels  $y_{v,i}$  of  $D_v$  is due to the potential noise from  $Y_t$ , therefore we take the prediction discrepancy as objective for label purification:

$$\mathcal{L}_{val}(Y_t) = \frac{1}{N_v} \sum_{i=1}^{N_v} \|y'_{v,i}(Y_t) - y_{v,i}\|^2 + \mathcal{H}(y'_{v,i}(Y_t)), \quad (6)$$

where  $\mathcal{H}(\cdot)$  represents the entropy of input distribution as a regularization term to sharpen the predicted label distribution as [35]. With Eq. (5) and Eq. (6), the validation loss can be expressed analytically by training labels  $Y_t$  in a batch, thus the noisy labels can be corrected steadily with correction rate  $\eta_I$  by the gradient back-propagated from Eq. (6) as

$$Y_t^{p+1} := Y_t^p - \eta_I \nabla(\mathcal{L}_{val}(Y_t^p)). \quad (7)$$

**Extrinsic Auxiliary Correction.** To accelerate the label correction process, an external correction process is further proposed. Specifically, an accompanied linear classifier  $C(\cdot; w_c) : \mathbb{R}^d \rightarrow \mathbb{R}^c$  with learnable parameter  $w_c$  is trained along with the updated labels from IPC:

$$\mathcal{L}_c(w_c) = \mathcal{L}_{ce}(C(\mathbf{f}_t; w_c), y'_t) + \mathcal{H}(C(\mathbf{f}_t; w_c)), \quad (8)$$

where  $y'_t$  is the updated training label from IPC, and  $\mathcal{L}_{ce}$  denotes the cross entropy loss function. Since the accompanied linear classifier is intrinsically robust to noisy labels in  $y'_t$  [24], it can quickly achieve high correction accuracy. With this intuition, the predicted results of the classifier  $C(\mathbf{f}_t; w_c)$  are used to correct labels periodically, specifically, the update rule of  $Y_t$  switches to momentum update (i.e., EMA) every  $T$  iterations:

$$Y_t^{p+1} := (1 - \eta_E) Y_t^p + \eta_E C(F_t; w_c) \quad \text{if } p = nT, \quad (9)$$

Table 1. Comparison with state-of-the-art methods on CIFAR-10/100 datasets with symmetric noise. ‘‘CE’’ is the standard ConvNet trained with Cross-Entropy loss in an end-to-end manner. ‘‘Classifier’’ means adopts the pre-trained SimCLR features to re-train a linear classifier. ‘‘Val’’ denotes using a small clean validation set. DivideMix\* denotes training DivideMix with the same validation set as additional data.

Dataset Method	Val	Noise ratio	CIFAR-10				CIFAR-100			
			20%	50%	80%	90%	20%	50%	80%	90%
Cross-Entropy (CE)	✗	Best	86.8	79.4	62.9	42.7	62.0	46.7	19.9	10.1
		Last	82.7	57.9	26.1	16.8	61.8	37.3	8.8	3.5
Co-teaching+ [36]	✗	Best	89.5	85.7	67.4	47.9	65.6	51.8	27.9	13.7
		Last	88.2	84.1	45.5	30.1	64.1	45.3	15.5	8.8
PENCIL [35]	✗	Best	92.4	89.1	77.5	58.9	69.4	57.5	31.1	15.3
		Last	92.0	88.7	76.5	58.2	68.1	56.4	20.7	8.8
REED [38]	✗	Best	95.8	95.6	94.3	93.6	76.7	73.0	66.9	59.6
		Last	95.7	95.4	94.1	93.5	76.5	72.2	66.5	59.4
Sel-CL+ [18]	✗	Best	95.5	93.9	89.2	81.9	76.5	72.4	59.6	48.8
		Last	95.1	93.3	88.7	81.6	76.1	72.0	59.2	48.6
MOIT+ [21]	✗	Best	94.1	91.8	81.1	74.7	75.9	70.6	47.6	41.8
		Last	93.8	91.3	80.6	74.0	75.2	70.1	46.9	41.2
C2D-DivideMix [40]	✗	Best	<b>96.3</b>	95.2	94.4	93.5	78.6	76.4	67.7	58.7
		Last	<b>96.2</b>	95.1	94.1	93.4	78.3	76.0	67.4	58.4
DivideMix [15]	✗	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
Meta-Learning [16]	✓	Best	92.9	89.3	77.4	58.7	68.5	59.2	42.4	19.5
		Last	92.0	88.8	76.1	58.3	67.7	58.0	40.1	14.3
MLC [41]	✓	Best	92.6	88.1	77.4	67.9	66.8	52.7	21.8	15.0
		Last	91.8	87.5	77.1	67.0	66.5	52.4	18.9	14.2
MSLC [9]	✓	Best	93.4	89.9	69.8	56.1	72.5	65.4	24.3	16.7
		Last	93.3	89.4	68.8	55.2	72.0	64.9	20.5	14.6
DivideMix* [15]	✓	Best	96.1	94.9	93.6	77.3	77.7	74.8	60.7	32.5
		Last	95.9	94.6	93.0	76.5	77.1	74.3	60.5	32.2
DMLP-Naive	✓	Best	94.7	94.2	93.5	92.8	72.7	68.0	63.5	61.3
		Last	94.2	94.0	93.2	92.0	72.3	67.4	63.2	60.9
DMLP-DivideMix	✓	Best	<b>96.3</b>	<b>95.8</b>	<b>94.5</b>	<b>94.3</b>	<b>79.9</b>	<b>76.8</b>	<b>68.6</b>	<b>65.8</b>
		Last	<b>96.2</b>	<b>95.6</b>	<b>94.3</b>	<b>94.0</b>	<b>79.4</b>	<b>76.1</b>	<b>68.5</b>	<b>65.4</b>

where  $T$  and  $\eta_E$  are the period and momentum for update, and  $n$  is an arbitrary positive integer to denote the  $n$ -th time of EAC update. In a global sense, after  $T$  iterations of training, EAC can quickly achieve locally optimal label estimation by mimic of gradually updated labels from IPC, which reversely facilitates the label correction of IPC by providing cleaner training labels. Subsequently, IPC and EAC form a positive loop and mutually improve the quality of label correction of each other.

### 3.3. Application of DMLP

The DMLP is a flexible label purifier, the corrected training labels  $y_t^*$  can be applied in different ways in robust learning scenarios, as shown in Fig. 3.

**Naive classification network with DMLP.** In a simple and direct manner, we can take the purified labels to retrain a neural network with simple cross-entropy loss (CE), here we term this simple application as DMLP-Naive.

**LNL framework boosted by DMLP.** Considering that there may still be a small number of noisy or miscorrected labels after purification, another effective way to apply DMLP is to take the purified labels as new training samples for the existing LNL framework. In this work, we extend some classic frameworks [12, 15, 19, 30] with DMLP, and the boosted LNL methods are denoted with prefix of ‘‘DMLP-’’ (e.g. DMLP-DivideMix, DMLP-ELR+, etc.)

## 4. Experiments

### 4.1. Experimental Settings

**CIFAR-10/100.** For the self-supervised pre-training stage, we adopt the popular SimCLR algorithm [5] with ResNet as the backbone network. Classifiers in meta-learner are trained for 100 epochs with the Adam optimizer. For the final DivideMix algorithm, ResNet18 is adopted for a fair comparison.  $\eta_I$  and  $\eta_E$  are set as 0.01 and 1.0 respec-

Table 2. Evaluation results with asymmetric noise of different noisy ratio on CIFAR-10. “Validation” denotes the method exploits a small clean validation set.

Method	Validation	Noisy ratio	
		20%	40%
Joint-Optim [28]	✗	92.8	91.7
PENCIL [35]	✗	92.4	91.2
M-correction [1]	✗	-	86.3
Iterative-CV [4]	✗	-	88.0
DivideMix [15]	✗	93.4	93.4
REED [38]	✗	95.0	92.3
C2D-DivideMix [40]	✗	93.8	93.4
Sel-CL+ [18]	✗	<b>95.2</b>	93.4
GCE [11]	✗	87.3	78.1
RRL [17]	✗	-	92.4
Zhang, et al. [39]	✓	92.7	90.2
Meta-Learning [16]	✓	-	88.6
MSLC [9]	✓	94.4	91.6
DMLP-Naive	✓	94.6	93.9
DMLP-DivideMix	✓	<b>95.2</b>	<b>95.0</b>

Table 3. Top-1 testing accuracy on Clothing-1M testset. “Validation” denotes using the validation provided by [31].

Method	Validation	Top-1 Accuracy
PENCIL [35]	✗	73.49
DivideMix [15]	✗	74.76
RRL [17]	✗	74.90
GCE [11]	✗	73.30
C2D-DivideMix [40]	✗	74.30
REED [38]	✗	75.81
Meta-Learning [16]	✓	73.47
Self-Learning [13]	✓	76.44
MLC [41]	✓	75.78
MSLC [9]	✓	74.02
Meta-Cleaner [10]	✓	72.50
Meta-Weight [8]	✓	73.72
FaMUS [34]	✓	74.40
MSLG [7]	✓	76.02
DMLP-Naive	✓	77.77
DMLP-DivideMix	✓	<b>78.23</b>

tively. The scaling factor  $\alpha$  is set as 1.0. To ensure fair evaluation, we follow previous meta-learning based LNL methods [8, 9] to randomly separate 1,000 images as the clean validation set for CIFAR-10/100 [14], leaving the rest as training samples. We strictly follow the protocol in [12] to generate label noise. Specifically, symmetric noise is generated by replacing labels with one of the other classes uniformly, while the labels in asymmetric noise are disturbed to their similar classes to simulate label noise in real-world scenarios. Our experiments are conducted under different noisy rates:  $\pi \in \{20\%, 50\%, 80\%, 90\%\}$  for symmetric and  $\pi \in \{20\%, 40\%\}$  for asymmetric noises.

**Clothing1M.** For the first stage on the Clothing1M [32],

the ResNet50 is trained with the official MoCo-v2 [6] to fully leverage its advantages on large-scale datasets. Afterward, the meta-learner is trained for 50 epochs. For DMLP-DivideMix, ResNet50 is adopted and initialized with weights from previous stages and trained for 80 epochs. Due to the space constraint, we put more detailed experimental settings and descriptions of each compared LNL method in our supplementary materials.

## 4.2. Experimental Results

**Comparison with state-of-the-art methods.** We compare our method with multiple recent competitive methods on CIFAR-10/100 under various noisy settings (detailed descriptions of these methods are provided in the supplementary materials). Both test accuracy of the best and last epoch are reported. As shown in Table 1, the simple DMLP-Naive can already achieve competitive results to most methods, the superiority is more obvious in extremely noisy cases, further, the DMLP-DivideMix achieves state-of-the-art performance across all the settings. It is worth noting that directly utilizing the validation data to train DivideMix (i.e., DivideMix\*) only brings marginal improvement, while when equipped with the purified labels by DMLP, significant improvements are obtained, indicating that DMLP is effective in terms of utilizing validation set towards LNL problem. On the other hand, though there exist other meta-learning methods utilizing validation set [9, 16, 41], DMLP shows great advantages over them. It is also noticeable that compared with the original DivideMix, the purified version of DMLP-DivideMix achieves better results by a large margin, indicating that the purified label from our approach is more friendly to boost LNL frameworks. Table 2 shows comparison with the recent methods on asymmetric noisy CIFAR-10 dataset. DMLP-DivideMix outperforms REED by 0.2% and 2.7% under different noisy ratios and obtains greater improvements over the rest methods, demonstrating the ability of DMLP in handling harder semantic-related noise. Finally, DMLP-based methods suffer less from increasing noisy ratio than other competitors, indicating *its robustness to variant noisy levels*.

Furthermore, we also evaluate DMLP on the large-scale real-world noisy dataset Clothing1M. As shown in Table 3, simple DMLP-Naive can outperform all other methods by a large margin, and DMLP-DivideMix further improves the accuracy by about 0.46%. The results indicate that *DMLP is more suitable for noise from real-world situations*.

**Label correction accuracy.** Fig. 4 compares the label accuracy after correction in our meta-learner against coupled label purifiers MLC and MSLC on CIFAR-10. Specifically, the one-hot form of the corrected pseudo-labels is compared with the ground truth for evaluation. As Fig. 4 shows, labels can be rapidly corrected to accuracy over 92% in low noise cases. For severe label noise, DMLP can still improve label accuracy similar to low-noise settings. *The overall corrected*

Table 4. Comparison between the LNL methods and their DMLP applications with symmetric noise on CIFAR-10/100. Specifically, the 9-layer CNN is adopted as the backbone network of Co-teaching.

Dataset Method/Noise ratio		CIFAR-10				CIFAR-100			
		20%	50%	80%	90%	20%	50%	80%	90%
Co-teaching [12]	Best	82.6	73.0	24.0	14.6	50.5	38.2	11.8	4.9
	Last	81.9	72.6	23.5	11.7	50.3	38.0	11.3	4.3
DMLP-Co-teaching	Best	<b>85.8</b>	<b>85.8</b>	<b>85.4</b>	<b>84.6</b>	<b>51.2</b>	<b>49.8</b>	<b>48.1</b>	<b>45.3</b>
	Last	<b>85.6</b>	<b>85.6</b>	<b>85.3</b>	<b>84.5</b>	<b>51.0</b>	<b>49.3</b>	<b>47.8</b>	<b>45.1</b>
CDR [30]	Best	90.4	85.0	47.2	12.3	63.3	39.5	29.2	8.0
	Last	82.7	49.4	16.6	10.1	62.9	39.5	9.7	4.5
DMLP-CDR	Best	<b>91.4</b>	<b>91.2</b>	<b>91.2</b>	<b>90.2</b>	<b>69.2</b>	<b>64.8</b>	<b>61.4</b>	<b>58.5</b>
	Last	<b>91.2</b>	<b>90.8</b>	<b>90.6</b>	<b>89.3</b>	<b>68.3</b>	<b>64.3</b>	<b>61.1</b>	<b>57.9</b>
ELR+ [19]	Best	94.6	93.8	91.1	75.2	77.5	72.4	58.2	30.8
	Last	94.4	93.7	90.5	73.5	76.2	72.2	56.8	30.6
DMLP-ELR+	Best	<b>94.9</b>	<b>94.1</b>	<b>93.0</b>	<b>92.5</b>	<b>77.8</b>	<b>73.6</b>	<b>63.9</b>	<b>60.5</b>
	Last	<b>94.6</b>	<b>94.0</b>	<b>92.7</b>	<b>92.1</b>	<b>77.1</b>	<b>73.4</b>	<b>63.6</b>	<b>60.5</b>




Figure 4. Comparison of corrected label accuracy under symmetric-20% (left), symmetric-80% (right) noise settings on CIFAR-10.

*label accuracy within DMLP is superior against other competitors across all noise settings* (More detailed experimental results can be found in the supplementary material).

**Generality of DMLP.** To validate the generalization ability of DMLP, other than DivideMix, another three popular LNL methods, ELR+ [19], Co-teaching [12], and CDR [30] are further adopted to work collaboratively with purified labels of DMLP. As shown in Table 4, all the applications of DMLP perform consistently better over their corresponding baselines, and the improvement becomes more significant when it comes to high-level noise cases. It is worth noting that since CDR highly relies on the early stopping technique, it suffers from a severe memorization effect in the training process, leading to a discrepancy between best results and last performance. In contrast, when training CDR with our purified labels, this discrepancy almost disappears, demonstrating the labels output by DMLP have a better quality to suppress the memorization effect, thus alleviating the reliance on early stopping. Therefore, the results indicate that *the purified labels of DMLP are friendly to boost LNL frameworks*. More detailed experimental results can be found in the supplementary materials.

### 4.3. Ablation Studies

**Analysis on different label correction process.** We explore the influence of IPC and EAC on the performance

Table 5. Ablation study for the effectiveness of IPC and EAC in DMLP-Naive on CIFAR-10.

Component		CIFAR-10				Clothing 1M	
IPC	EAC	20%	50%	80%	90%		
✗	✓	Best	93.7	93.3	91.1	67.4	76.5
		Last	93.0	92.9	90.6	66.5	76.1
✓	✗	Best	87.8	85.7	79.9	76.0	76.8
		Last	87.2	85.5	79.4	75.4	76.5
✓	✓	Best	<b>94.7</b>	<b>94.2</b>	<b>93.5</b>	<b>92.8</b>	<b>77.7</b>
		Last	<b>94.2</b>	<b>94.0</b>	<b>93.2</b>	<b>92.0</b>	<b>77.6</b>

of DMLP-Naive in Table 5, when one of the process is excluded, the updated labels from the other are applied to retrain a new model. It is observed that EAC performs well in low noise cases due to its intrinsic robustness, as the noise ratio increases, the performance drops rapidly. On the other hand, IPC is more robust to high-level noise, but there exists a large gap compared with the full DMLP pipeline due to its slow optimization process. In contrast, when IPC and EAC work collaboratively, DMLP can achieve optimal results.

**Comparison against other coupled purifiers with pre-training.** To solely evaluate the influence of decoupled purification, we train two coupled meta label correction methods MLC and MSLC with the same self-supervised pretrained weights and apply their corrected labels to naive training or DivideMix for fair comparisons. As shown in Table 6, though self-supervised weights can marginally boost the performance of coupled label correctors, there still exists a large gap between their performance and DMLP, especially for high noisy cases. Moreover, when further applying the corrected labels to mainstream LNL framework DivideMix, our method can also consistently outperform the coupled counterparts across all the noisy settings, demonstrating our corrected labels are of better quality. Therefore, these results verify that *the superiority mainly attributes to decoupled*

Table 6. Comparison with coupled meta label correction methods MLC [41] and MSLC [9] on CIFAR-10. "\*" denotes training with SimCLR pretrained ResNet-18.

Method		Noisy ratio			
		20%	50%	80%	90%
MLC*	Best	91.8	86.2	77.6	72.9
	Last	91.6	85.9	77.5	72.6
MSLC*	Best	92.0	87.7	78.0	67.8
	Last	92.0	87.5	77.9	67.3
DMLP-Naive*	Best	<b>94.0</b>	<b>93.7</b>	<b>93.1</b>	<b>92.3</b>
	Last	<b>93.9</b>	<b>93.4</b>	<b>92.9</b>	<b>91.9</b>
MLC*-DivideMix	Best	95.3	94.0	93.0	86.6
	Last	95.0	93.6	92.7	86.5
MSLC*-DivideMix	Best	95.7	94.9	93.8	83.0
	Last	95.5	94.8	93.1	82.8
DMLP*-DivideMix	Best	<b>96.3</b>	<b>95.6</b>	<b>94.1</b>	<b>93.8</b>
	Last	<b>96.0</b>	<b>95.2</b>	<b>94.0</b>	<b>93.6</b>

Table 7. Investigation of the validation set size  $\tau$  on Clothing1M.

$\tau$	10%	20%	30%	40%	50%	100%
Accuracy (%)	75.50	76.40	76.61	77.00	77.30	<b>77.31</b>

label correction instead of self-supervised pretraining.

**Effect of different feature representation for purification.** The quality of features plays a crucial role in the label correction process of DMLP since the distribution of learned features is closely related to the rationality behind the linear estimation assumption in high-dimensional space. Therefore, we study the influence of different features on performance. Specifically, two types of features are investigated, including features from the ResNet-18/50 which load the self-supervised pre-trained weights. As the results in Table 8 show, the features from the ResNet-18 lead to slightly poor performance, while it brings performance improvements when using features from self-supervised ResNet-50. This observation indicates that although feature representation of higher quality benefits the purification results, DMLP is not very sensitive to the representation ability of input feature.

**Effect of validation size.** We examine how the number of validation set affect performance. Specifically, validation sizes from 10% to 100% of the whole validation set are evaluated on Clothing1M for DMLP-Naive. As shown in Table 7, DMLP-Naive achieves similar performance regardless of the validation size  $N_v$ , demonstrating DMLP is not sensitive to the number of images. It is worth noting that even using only 10% of the validation set (around 0.1% of training data), DMLP-Naive still achieves high accuracy and outperforms most methods in Table.3, indicating that the effectiveness of DMLP is not heavily relied on validation size.

**Performance under extremely noisy setting.** In an extremely noisy scenario where all labels in the training set are unreliable except the given clean validation set, the LNL problem is converted into a semi-supervised learning prob-

Table 8. Ablation study for adopting different features in DMLP-Naive on CIFAR-10, where "R18/50" denote "ResNet-18/50" and "M/S" represent "MoCo/SimCLR".

Feature Source		Noisy ratio			
		20%	50%	80%	90%
R18 (M)	Best	93.8	93.3	92.2	90.4
	Last	93.7	92.7	92.1	90.0
R18 (S)	Best	94.0	93.7	93.1	92.3
	Last	93.9	93.4	92.9	91.9
R50 (S)	Best	<b>94.7</b>	<b>94.2</b>	<b>93.5</b>	<b>92.8</b>
	Last	<b>94.2</b>	<b>94.0</b>	<b>93.2</b>	<b>92.0</b>

Table 9. Comparison between recent semi-supervised methods and DMLP-DivideMix on CIFAR-10/100 with 100% noisy ratio.

Method	CIFAR-10	CIFAR-100
MeanTeacher	83.0	31.0
MixMatch	87.9	57.7
FixMatch	88.1	56.3
UDA	88.2	56.1
Ours	<b>91.7</b>	<b>60.1</b>

lem, therefore we further compare DMLP-DivideMix with some state-of-the-art semi-supervised learning algorithms, including MeanTeacher [29], MixMatch [3], FixMatch [26] and UDA [33], under the 100% symmetric noise case on CIFAR-10 and CIFAR-100. From the results in Table 9, DMLP-DivideMix performs optimally among all methods when using the validation set as labeled samples, indicating the proposed method is potential for broader applications.

## 5. Conclusion

In this paper, we propose a flexible and novel multi-stage robust learning approach termed as DMLP. The core of DMLP is a carefully-designed meta-learning based label purifier, which decouples the complex bi-level optimization problem into representation and label distribution learning, thus helping the meta-learner focus on correcting noisy labels in a faster and more precise manner even under extremely noisy scenarios. Further, DMLP can be applied either for naive retraining on noisy data or assistance of existing LNL methods to boost performance. Extensive experiments conducted on several synthetic and real-world noisy datasets verify the superiority of the proposed method.

## 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.
- [2] Devansh Arpit, Stanislaw Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron C. Courville, Yoshua Bengio, and Simon Lacoste-Julien. A closer look at memorization in deep networks. In *Proc. International Conference on Machine Learning (ICML)*, 2017.
- [3] David Berthelot, Nicholas Carlini, Ian J. Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [4] 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.
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proc. International Conference on Machine Learning (ICML)*, pages 1597–1607. PMLR, 2020.
- [6] Xinlei Chen, Haoqi Fan, Ross B. Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*.
- [7] Algan et al. Meta soft label generation for noisy labels. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7142–7148, 2021.
- [8] Shu et al. Meta-weight-net: Learning an explicit mapping for sample weighting. In *NeurIPS*, 2019.
- [9] Wu et al. Learning to purify noisy labels via meta soft label corrector. In *AAAI*, 2021.
- [10] Zhang et al. Metacleaner: Learning to hallucinate clean representations for noisy-labeled visual recognition. In *CVPR*, 2019.
- [11] 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.
- [12] 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.
- [13] Jiangfan Han, Ping Luo, and Xiaogang Wang. Deep self-learning from noisy labels. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [14] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [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.
- [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.
- [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.
- [18] Shikun Li, Xiaobo Xia, Shiming Ge, and Tongliang Liu. Selective-supervised contrastive learning with noisy labels. *arXiv preprint arXiv:2203.04181*, 2022.
- [19] 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.
- [20] Tongliang Liu and Dacheng Tao. Classification with noisy labels by importance reweighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2015.
- [21] 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.
- [22] 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.
- [23] Geoff Pleiss, Tianyi Zhang, Ethan R. Elenberg, and Kilian Q. Weinberger. Identifying mislabeled data using the area under the margin ranking. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [24] Scott E. Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. Training deep neural networks on noisy labels with bootstrapping. In *Proc. International Conference on Learning Representations (ICLR)*, 2015.
- [25] 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)*, 2018.
- [26] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [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.
- [28] 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.
- [29] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [30] 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.
- [31] 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.
- [32] 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)*, pages 2691–2699, 2015.
- [33] Qizhe Xie, Zihang Dai, Eduard H. Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [34] Youjiang Xu, Linchao Zhu, Lu Jiang, and Yi Yang. Faster meta update strategy for noise-robust deep learning. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 144–153, 2021.
- [35] 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)*, 2019.
- [36] 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.
- [37] 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.
- [38] Hui Zhang and Quanming Yao. Decoupling representation and classifier for noisy label learning. *arXiv preprint arXiv:2011.08145*, 2020.
- [39] 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.
- [40] 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.
- [41] 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.