ROBUST LEARNING WITH DECOUPLED META LABEL PURIFIER

Anonymous authors
Paper under double-blind review

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.

1 INTRODUCTION

Deep learning has achieved significant progress on various recognition tasks. The key to its success is the availability 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 (Xiao et al., 2015a), 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 (Arpit et al., 2017).

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 (Han et al., 2018) 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 (Ren et al., 2018; Li et al., 2019; Wu et al., 2021) 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 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 (Ren et al., 2018; Li et al., 2019; Wu et al., 2021), resulting in a coupled solution for both 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 with recent label

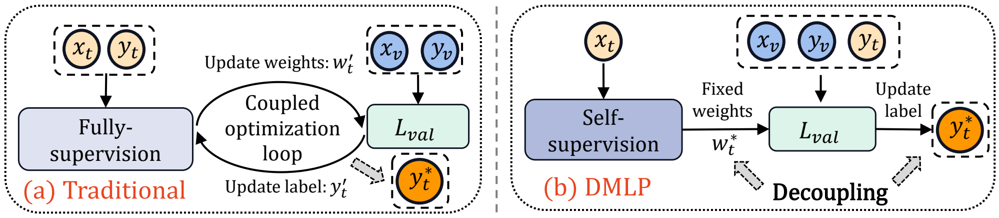

Figure 1: The comparison between (a) traditional coupled alternating update to solve meta label purification problem, and (b) the proposed DMLP method that decouples the label purification process into representation learning and a simple non-nested meta label purifier.

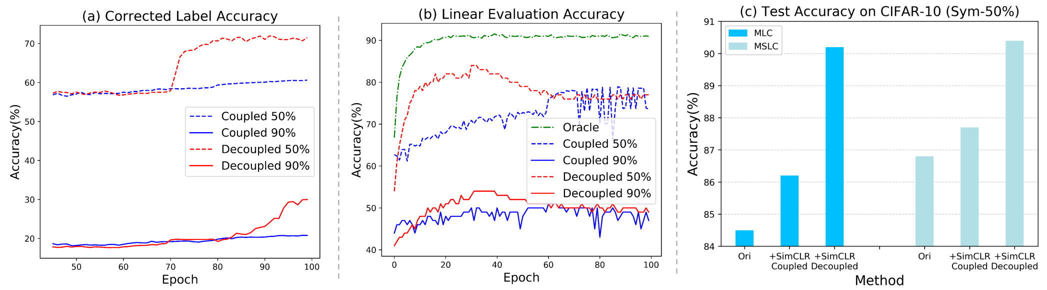

Figure 2: The corrected label accuracy (a) and linear probe accuracy of representations (b) between coupled (Zheng et al., 2021) and decoupled label correction schemes during training under 50% and 90% symmetric label noise on CIFAR-10. (c) investigates the effect of representation learning on 'Ori'-original 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.

purification methods MLC (Zheng et al., 2021) and MSLC (Wu et al., 2021), 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 (Chen et al.) 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 (Chen et al., 2020) as initialization, recent effort (Zheltonozhskii et al., 2022) 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). We observe the pretrained model only brings marginal improvement if model weights optimization is still coupled with hyper-parameters. In contrast, when the weight of backbone is fixed and decoupled from the purification, 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 purification 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 for decoupled optimization of label distribution than existing complex meta networks (Zheng et al., 2021; Wu et al., 2021). 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 (Han et al., 2018; Pleiss et al., 2020), assigned with smaller weights via a sample re-weight training scheme (Liu & Tao, 2015), or used to formulate a semi-supervised learning problem by throwing away the labels while keeping the unlabeled data (Li et al., 2020; Zhang & Yao, 2020). 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 discriminative 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 (Patrini et al., 2017) 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 (Xiao et al., 2015a). Some other works resort to exploiting the prediction of the network, both soft (Reed et al., 2015; Han et al., 2019; Yi & Wu, 2019; Arazo et al., 2019) and hard (Tanaka et al., 2018; Song et al., 2019) 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 (Zhang & Yao, 2020). Another line of works utilize robust representations learned via unsupervised contrastive learning methods (Zhang & Yao, 2020; Li et al., 2021; Ghosh & Lan, 2021; Zheltonozhskii et al., 2022) to eliminate the interference of noisy labels, which provides a reliable initialization of deep models. Recently, meta-learning based methods (Ren et al., 2018; Li et al., 2019; Wu et al., 2021; Zheng et al., 2021) 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 (Wu et al., 2021; Zheng et al., 2021) 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, the learning process on model weights and labels are decoupled into individual stages within DMLP. Together with the proposed non-nested meta-label purifier, DMLP yields more accurate labels than coupled label correction methods (Wu et al., 2021; Zheng et al., 2021) and further set the new state-of-the-art on CIFAR and Clothing1M.

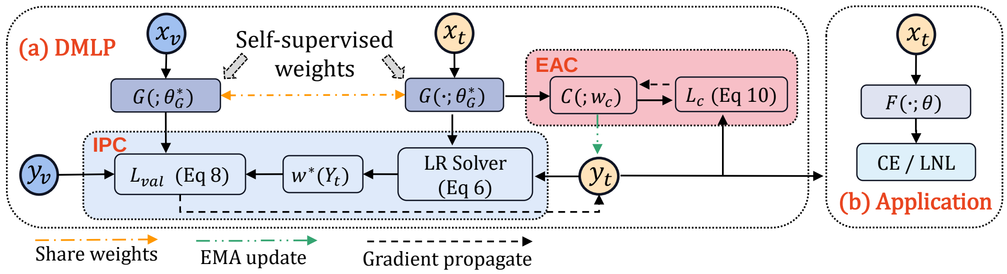

Figure 3: The overall framework of DMLP. (a) Decoupled meta label purification (Sec. 3.2). (b) The purified labels can be applied in normal network retraining (CE) or other LNL loss (Sec. 3.3).

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, the 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, validation set.

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

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

where w and θ_α denote the model weights and the meta hyper-parameters respectively. In the classical label purification pipeline, θ_α 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 (Ren et al., 2018; Wu et al., 2021) 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:

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

where a pretrained feature extractor $G : \mathbb{R}^{H\times W\times 3} \to \mathbb{R}^d$ is designed to extract d-dimensional representation $\mathbf{f}$. These extracted features are utilized in a contrastive learning framework (Chen et al., 2020; Chen et al.) to update the parameters θ_G. Subsequently, the established feature extractor $\mathbf{f} = G(x;\theta_G^*)$ can learn a noise-agnostic descriptor of images, which is also highly separable in high dimensional feature space (Zhang & Yao, 2020). In this manner, we detach representation learning from noisy label correction, while keeping strong separability of features.

Further, since the feature is representative and separable, the loss term can be formulated with simple estimation risk function (e.g. linear discrimination) instead of complex DNN forward pass, making it possible to solve the meta-purification 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 **i**ntrinsic **p**rimary **c**orrection (IPC) and **e**xtrinsic **a**uxiliary **c**orrection (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 (b), a batch of b training data is gathered into matrix:

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

Since the feature descriptors are representative, we assume there exists a simple linear estimation transform $w \in \mathbb{R}^{d \times c}$, which accurately regresses the feature descriptor of a sample to the one-hot form belonging to its semantic label:

$$\min_{w} \mathcal{L}_{train}(F_t, Y_t; w) = \|Y_t - F_t w\|^2. \tag{4}$$

By solving the linear regression problem of Eq. (4) through least square method, we can obtain its closed-form solution $w^*(Y_t)$ on the training batch and derive its optimal linear prediction on samples $\mathbf{f}_{v,i}$ from validation set D_v:

$$y'_{v,i}(Y_t) = w^*(Y_t)^T \mathbf{f}_{v,i} = Y_t^T F_t \left(F_t{}^T F_t\right)^{-1} \mathbf{f}_{v,i}. \tag{5}$$

To ensure inverse is feasible, we will add a small unit diagonal matrix to $F_t^T F_t$ if the inverse calculation fails at first. Intuitively, the discrepancy between the predicted results and the ground truth labels of D_v can be due to the potential noise in Y_t, therefore we take the prediction discrepancy as objective for label purification:

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

where $\mathcal{H}(\cdot)$ represents the entropy of input distribution as a regularization term to sharpen the predicted label distribution, similar to Yi & Wu (2019). With Eq(5,6), the validation loss can be expressed analytically by training labels Y_t in a batch, thus the noisy labels are corrected steadily with correction rate η_I by the gradient from Eq. (6), in implementation, a softmax function is applied to ensures the label vector sums to one:

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

• **Extrinsic Auxiliary Correction.** To accelerate the label correction process, an external correction process is further proposed. Specifically, an accompanied linear classifier $C(\cdot; w_c)$ is trained together with the updated labels from IPC:

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

where y'_t is the updated training labels, and $\mathcal{L}_{\text{ce}}$ denotes the cross entropy loss function. Since the accompanied classifier is intrinsically robust to noisy labels in y'_t, it can quickly achieve relatively high correction accuracy. With this intuition, the predicted results of the classifier are used to update noisy labels periodically in a momentum manner:

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

where T and η_E are the period and momentum for 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 performance.

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 the whole 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, four classic frameworks are adopted, including DivideMix (Li et al., 2020), ELR+ (Liu et al., 2020), CDR (Xia et al., 2021) and Co-teaching (Han et al., 2018) . The boosted LNL methods are denoted with prefix of "DMLP-" (DMLP-DivideMix, DMLP-ELR+, etc.)

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 20%	CIFAR-10 50%	CIFAR-10 80%	CIFAR-10 90%	CIFAR-100 20%	CIFAR-100 50%	CIFAR-100 80%	CIFAR-100 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+ (Yu et al., 2019)	✗	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 (Yi & Wu, 2019)	✗	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 (Zhang & Yao, 2020)	✗	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
RRL (Li et al., 2021)	✗	Best	95.9	94.5	-	-	79.4	75.0	-	-
		Last	95.6	94.1	-	-	79.0	74.5	-	-
Sel-CL+ (Li et al., 2022)	✗	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+ (Ortego et al., 2021)	✗	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 (Zheltonozhskii et al., 2022)	✗	Best	**96.3**	95.2	94.4	93.5	78.6	76.4	67.7	58.7
		Last	**96.2**	95.1	94.1	93.4	78.3	76.0	67.4	58.4
DivideMix (Li et al., 2020)	✗	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 (Li et al., 2019)	✔	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 (Zheng et al., 2021)	✔	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 (Wu et al., 2021)	✔	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* (Li et al., 2020)	✔	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	**96.3**	**95.8**	**94.5**	**94.3**	**79.9**	**76.8**	**68.6**	**65.8**
		Last	**96.2**	**95.6**	**94.3**	**94.0**	**79.4**	**76.1**	**68.5**	**65.4**

4 Experiments

4.1 Experimental Settings

CIFAR-10/100. For the self-supervised pre-training stage, we adopt the popular SimCLR algorithm (Chen et al., 2020) 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. η_I and η_E are set as 0.01 and 1.0 respectively. To ensure fair evaluation, we only randomly separate 1,000 images as the clean set for both CIFAR-10/100 (Krizhevsky et al., 2009), leaving the rest as training samples. We strictly follow the protocol in (Han et al., 2018) to generate 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 (Xiao et al., 2015b), the ResNet50 is trained with the official MoCo-v2 (Chen et al.) to fully leverage its advantages on large-scale datasets. Afterward, the meta-learner is trained for 50 epochs.For DMLP-Mix, ResNet50 is adopted and initialized with weights from previous stages and trained for 80 epochs. More details can be found in the supplementary materials.

4.2 Experimental Results

• **Comparison with state-of-the-art methods.** We compare our method with multiple recent competing 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

Table 2: Evaluation with asymmetric noise on CIFAR-10. "Val" denotes the method exploits a small clean validation set.

Method	Val	Noisy ratio 20%	Noisy ratio 40%
Joint-Optim(Tanaka et al., 2018)	✗	92.8	91.7
PENCIL (Yi & Wu, 2019)	✗	92.4	91.2
M-correction (Arazo et al., 2019)	✗	-	86.3
Iterative-CV (Chen et al., 2019)	✗	-	88.0
DivideMix (Li et al., 2020)	✗	93.4	93.4
REED (Zhang & Yao, 2020)	✗	95.0	92.3
C2D (Zheltonozhskii et al., 2022)	✗	93.8	93.4
Sel-CL+ (Li et al., 2022)	✗	**95.2**	93.4
GCE (Ghosh & Lan, 2021)	✗	87.3	78.1
RRL (Li et al., 2021)	✗	-	92.4
Zhang, et al. (Zhang et al., 2020)	✔	92.7	90.2
Meta-Learning (Li et al., 2019)	✔	-	88.6
MSLC (Wu et al., 2021)	✔	94.4	91.6
DMLP-Naive	✔	94.6	93.9
DMLP-DivideMix	✔	**95.2**	**95.0**

Table 3: Top-1 testing accuracy on Clothing-1M. "Val" denotes using the validation provided by (Xiao et al., 2015a).

Method	Val	Top-1
PENCIL (Yi & Wu, 2019)	✗	73.49
DivideMix (Li et al., 2020)	✗	74.76
RRL (Li et al., 2021)	✗	74.90
GCE (Ghosh & Lan, 2021)	✗	73.30
C2D (Zheltonozhskii et al., 2022)	✗	74.30
REED (Zhang & Yao, 2020)	✗	75.81
Meta-Learning (Li et al., 2019)	✔	73.47
Self-Learning (Han et al., 2019)	✔	76.44
MLC (Zheng et al., 2021)	✔	75.78
MSLC (Wu et al., 2021)	✔	74.02
Meta-Cleaner (Zhang et al., 2019)	✔	72.50
Meta-Weight (Shu et al., 2019)	✔	73.72
FaMUS (Xu et al., 2021)	✔	74.40
MSLG (Algan & Ulusoy, 2021)	✔	76.02
DMLP-Naive	✔	77.77
DMLP-DivideMix	✔	**78.23**

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		CIFAR-10				CIFAR-100			
Method/Noise ratio		20%	50%	80%	90%	20%	50%	80%	90%
Co-teaching (Han et al., 2018)	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	**85.8**	**85.8**	**85.4**	**84.6**	**51.2**	**49.8**	**48.1**	**45.3**
	Last	**85.6**	**85.6**	**85.3**	**84.5**	**51.0**	**49.3**	**47.8**	**45.1**
CDR (Xia et al., 2021)	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	**91.4**	**91.2**	**91.2**	**90.2**	**69.2**	**64.8**	**61.4**	**58.5**
	Last	**91.2**	**90.8**	**90.6**	**89.3**	**68.3**	**64.3**	**61.1**	**57.9**
ELR+ (Liu et al., 2020)	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	**94.9**	**94.1**	**93.0**	**92.5**	**77.8**	**73.6**	**63.9**	**60.5**
	Last	**94.6**	**94.0**	**92.7**	**92.1**	**77.1**	**73.4**	**63.6**	**60.5**

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 (Li et al., 2019; Wu et al., 2021; Zheng et al., 2021), 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.

In addition to artificial noise, 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 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

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

Component IPC	Component EAC		CIFAR-10 20%	CIFAR-10 50%	CIFAR-10 80%	CIFAR-10 90%	Clothing1M
✗	✓	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	**94.7**	**94.2**	**93.5**	**92.8**	**77.7**
		Last	**94.2**	**94.0**	**93.2**	**92.0**	**77.6**

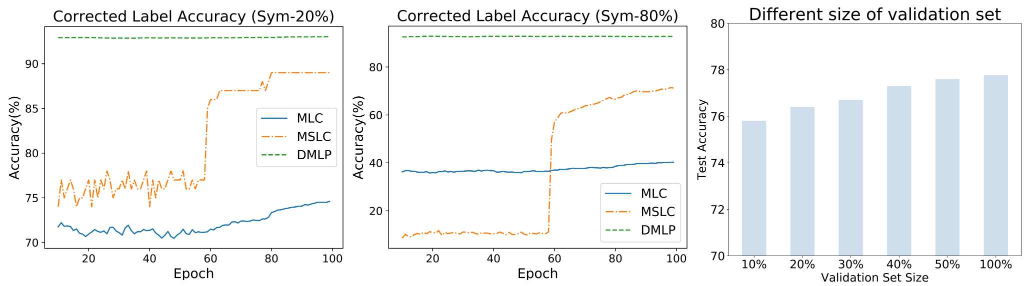

Figure 4: Comparison of corrected label accuracy curve under symmetric-20% (left), symmetric-80% (middle) noise settings on CIFAR-10. And investigation of the validation set size on Clothing1M (right).

can be rapidly corrected to accuracy over 92% in low noise cases. For severe label noise, DMLP can still steadily improve label accuracy similar to low-noise settings. And the overall corrected 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 3 popular LNL methods, ELR+ (Liu et al., 2020), Co-teaching (Han et al., 2018), and CDR (Xia et al., 2021) are further adopted to work collaboratively with purified labels of DMLP. As show in Table 4, all the applications of DMLP perform consistently better over their corresponding baselines, especially under 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. More detailed experimental results can be found in the supplementary materials. Therefore, the results indicate that the purified labels of DMLP are friendly to boost LNL frameworks.

4.3 Ablation Studies

• **Component analysis.** We explore the influence of IPC and EAC on the performance of DMLP-Naive in Table 5, when EAC is excluded, the optimization results $w^*(Y_t)$ in Eq. (5) are applied as a linear classifier. 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 pretraining.** 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 coupled counterparts across all the noisy cases, demonstrating our corrected labels are of higher quality. Therefore, all these comparison verify **that the improvement mainly attributes to our decoupled label correction strategy instead of self-supervised pretraining**.

Table 6: Comparison with coupled meta label correction methods MLC (Zheng et al., 2021) and MSLC (Wu et al., 2021). "*" denotes training with SimCLR pre-trained 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	**94.0**	**93.7**	**93.1**	**92.3**
	Last	**93.9**	**93.4**	**92.9**	**91.9**
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	**96.3**	**95.6**	**94.1**	**93.8**
	Last	**96.0**	**95.2**	**94.0**	**93.6**

Table 7: 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	**94.7**	**94.2**	**93.5**	**92.8**
	Last	**94.2**	**94.0**	**93.2**	**92.0**

Table 8: Results of recent semi-supervised methods and DMLP-DivideMix on CIFAR-10.

	Mean Teacher	Mix Match	Fix Match	UDA	Ours
Acc	83.0	87.9	88.1	88.2	**91.7**

• **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 7 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 the performance. Specifically, validation sizes from 10% to 100% of the whole validation set are evaluated on Clothing1M for DMLP-Naive. As shown in Fig. 4, DMLP-Naive achieves similar performance regardless of the validation size N_v, demonstrating the performance of 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 validation set, the LNL problem is converted into a partially-labeled semi-supervised learning problem, therefore we further compare DMLP-DivideMix with state-of-the-art semi-supervised learning methods in the 100% symmetric noise case, including MeanTeacher (Tarvainen & Valpola, 2017), MixMatch (Berthelot et al., 2019), FixMatch (Sohn et al., 2020), UDA (Xie et al., 2020). From the results in Table 8, DMLP-DivideMix performs optimally among all methods when using the validation set as labeled training samples, indicating that the proposed method is suitable 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 direct inference 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.

REFERENCES

Görkem Algan and Ilkay Ulusoy. Meta soft label generation for noisy labels. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pp. 7142–7148, 2021.

Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. In *ICML*, 2019.

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 *ICML*, 2017.

David Berthelot, Nicholas Carlini, Ian J. Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *NeurIPS*, 2019.

Pengfei Chen, Benben Liao, Guangyong Chen, and Shengyu Zhang. Understanding and utilizing deep neural networks trained with noisy labels. In *ICML*, 2019.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pp. 1597–1607. PMLR, 2020.

Xinlei Chen, Haoqi Fan, Ross B. Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*.

Aritra Ghosh and Andrew Lan. Contrastive learning improves model robustness under label noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2703–2708, 2021.

Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, 2018.

Jiangfan Han, Ping Luo, and Xiaogang Wang. Deep self-learning from noisy labels. In *ICCV*, 2019.

Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S. Kankanhalli. Learning to learn from noisy labeled data. In *CVPR*, 2019.

Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *ICLR*, 2020.

Junnan Li, Caiming Xiong, and Steven CH Hoi. Learning from noisy data with robust representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9485–9494, 2021.

Shikun Li, Xiaobo Xia, Shiming Ge, and Tongliang Liu. Selective-supervised contrastive learning with noisy labels. *arXiv preprint arXiv:2203.04181*, 2022.

Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. In *NeurIPS*, 2020.

Tongliang Liu and Dacheng Tao. Classification with noisy labels by importance reweighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2015.

Diego Ortego, Eric Arazo, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Multi-objective interpolation training for robustness to label noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6606–6615, 2021.

Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *CVPR*, 2017.

Geoff Pleiss, Tianyi Zhang, Ethan R. Elenberg, and Kilian Q. Weinberger. Identifying mislabeled data using the area under the margin ranking. In *NeurIPS*, 2020.

Scott E. Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. Training deep neural networks on noisy labels with bootstrapping. In *ICLR*, 2015.

Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *ICML*, 2018.

Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. *Advances in neural information processing systems*, 32, 2019.

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 *NeurIPS*, 2020.

Hwanjun Song, Minseok Kim, and Jae-Gil Lee. Selfie: Refurbishing unclean samples for robust deep learning. In *ICML*, 2019.

Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *CVPR*, 2018.

Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017.

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.

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 *9th International Conference on Learning Representations, ICLR 2021*, 2021.

Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *CVPR*, 2015a.

Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2691–2699, 2015b.

Qizhe Xie, Zihang Dai, Eduard H. Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. In *NeurIPS*, 2020.

Youjiang Xu, Linchao Zhu, Lu Jiang, and Yi Yang. Faster meta update strategy for noise-robust deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 144–153, 2021.

Kun Yi and Jianxin Wu. Probabilistic end-to-end noise correction for learning with noisy labels. In *CVPR*, 2019.

Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor W Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *ICML*, 2019.

Hui Zhang and Quanming Yao. Decoupling representation and classifier for noisy label learning. *arXiv preprint arXiv:2011.08145*, 2020.

Weihe Zhang, Yali Wang, and Yu Qiao. Metacleaner: Learning to hallucinate clean representations for noisy-labeled visual recognition. In *CVPR*, 2019.

Z. Zhang, H. Zhang, S. Arik, H. Lee, and T. Pfister. Distilling effective supervision from severe label noise. In *CVPR*, 2020.

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 *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1657–1667, 2022.

Guoqing Zheng, Ahmed Hassan Awadallah, and Susan Dumais. Meta label correction for noisy label learning. *AAAI 2021*, 2021.

A Appendix

A.1 Algorithm Details

In this appendix, we provide material that could not be included in the main manuscript due to space constraints. First, Section A.1 presents the detailed algorithm of the proposed method. Section A.2 provides insights about how IPC works from the perspective of label aggregation. Finally, Section A.3 presents additional experimental results of DMLP.

Algorithm 1 delineates the proposed DMLP in detail, where the equation labels are consistent with the main text of the paper. Firstly, the feature extractor $G(\cdot;\theta_G)$ is pre-trained via a self-supervised learning stage. IPC (Line 5-9) and EAC (Line 12-16) are two mutually reinforcing label correcting processes of DMLP. After m iterations of optimization, the purified labels $Y_t{}^*$ are obtained. Finally, the well-trained classifier $C(\cdot;w_c)$ can directly infer on test data or re-train a LNL framework with $Y_t{}^*$ to further boost performance.

A.2 Detailed Interpretation of IPC

In this section we further discuss how our IPC utilizes the decoupled feature representation to model the risk resulted from label noises from the perspective of label aggregation. In the main part of our paper, the IPC process seeks to predict the labels of validation samples via its optimally estimated linear estimator as

$$y'_{v,i}(Y_t) = Y_t^T F_t (F_t^T F_t)^{-1} \mathbf{f}_{v,i}. \tag{10}$$

Since the representation features are normalized after self-supervised training, thus $(F_t^T F_t)^{-1}$ can be interpreted as the inversed covariance matrix of distribution from training data, and the estimated covariance matrix maps the feature of validation data into the observation space of training data

$$\mathbf{f}'_{v,i} = (F_t^T F_t)^{-1} \mathbf{f}_{v,i}. \tag{11}$$

With the definition of Eq. (11), the term of matrix multiplication in Eq. (10) can be interpreted as a similarity matrix

$$\alpha = F_t \left[(F_t^T F_t)^{-1} \mathbf{f}_{v,i}\right] = F_t \mathbf{f}'_{v,i}, \tag{12}$$

each entry of vector $\alpha \in \mathbb{R}^b$ represents the similarity between $f'_{v,i}$ and each training sample of F_t. Finally, the output prediction on validation data can be regarded as the attentive aggregation over all labels in a training batch

$$y'_{v,i}(Y_t) = Y_t^T \alpha. \tag{13}$$

Ideally, since the self-supervised training process is trained without noisy labels in a contrastive manner, the feature distribution intrinsically forms a cluster-like manifold, i.e. samples with the same semantic label are closer in feature space

$$\mathbf{f}_{t,j}^T \mathbf{f}'_{v,i} > \mathbf{f}_{t,k}^T \mathbf{f}'_{v,i} \quad \textbf{when} \quad y_{v,i} = y_{t,j}, y_{v,i} \neq y_{t,k}. \tag{14}$$

Consequently, when a training sample in F_t is more similar to the validation sample, its semantic label will contribute more to prediction $y'_{v,i}(Y_t)$ and vice versa. Therefore, when penalizing on the discrepancy between $y'_{v,i}(Y_t)$ and $y_{v,i}$, we put more penalty on training samples with similar feature distribution but different labels from $(x_{v,i}, y_{v,i})$, which is more likely to be noisy samples.

A.3 Experimental Details

A.3.1 Experimental Settings

For the CIFAR10/100 dataset, most parameters of the SimCLR (Chen et al., 2020) algorithm are set as suggested in the original implementation. The classifier $C(\cdot;w_c)$ is trained with a learning rate of 0.005 and batch size of 500, while the batch size for the label purifier in IPC is set to 7,000. The classifier and the purifier are both trained with the Adam optimizer. For DMLP-Mix, all the hyper-parameters of DivideMix are set as the authors suggested. As for the real-world Clothing1M dataset, we follow the original protocol to split the training and testing sets. The given validation set

Algorithm 1 The workflow of DMLP.

Input: Noisy training set D_t, clean validation set D_v, feature extractor $G(\cdot;\theta_G)$, classifier $C(\cdot;w_c)$, batch size b, max iterations m, period for regular label substitution T.

Procedure:

Self-supervised training for $G(\cdot;\theta_G)$
Generate features **f** by Eq. (2)
for $i = 1$ to m **do**
    /***IPC** starts*/
    $\{F_t, Y_t\} \leftarrow SampleMiniBatch(\mathbf{f}, \mathrm{D}_t, b)$
    Calculate closed-form solution $w^*(Y_t)$ by Eq. (5)
    Predict validation set labels ${y'}_v$ by Eq. (5)
    Calculate label purification loss $\mathcal{L}_{\text{val}}(Y_t)$ by Eq. (6).
    Update training labels Y_t in backward process by Eq. (7).

    /***EAC** starts*/
    Calculate loss for the classifier $C(\cdot;w_c)$ by Eq. (8)
    Update classifier parameter w_c in backward process.
    if $i = nT$ **then**
        Update training labels Y_t by Eq. (9)
    end if
end for

Output: The purified labels ${Y_t}^*$.

is utilized for meta-learning. The batch sizes for the classifier and the label purifier in IPC are set to 500 and 10,000 respectively. The learning rate for the former is set to 0.03 while for the latter is 0.02. Similar to the setting on CIFAR-10/100, the Adam optimizer is also adopted.

A.3.2 Details of Compared Methods

As discussed in manuscript, DMLP is compared with most recent relevant works. Specifically, the competing works can be coarsely categorized into two group, noisy sample detection and label correction. The former usually identifying and reducing the importance of suspicious false-labeled samples during training, either by directly selecting the clean samples out of training set (Co-teaching (Han et al., 2018), Co-teaching+ (Yu et al., 2019), Iterative-CV (Chen et al., 2019), Sel-CL+ (Li et al., 2022), ELR+ (Liu et al., 2020), C2D (Zheltonozhskii et al., 2022), DivideMix (Li et al., 2020), REED (Zhang & Yao, 2020), MOIT+ (Ortego et al., 2021)) or adjusting the soft weight of each training sample (RRL (Li et al., 2021), M-correction (Arazo et al., 2019), GCE (Ghosh & Lan, 2021), CDR (Xia et al., 2021)). The latter aims to correct the corrupted labels and augment the training data. Typical paradigms are correction-by-prediction, i.e., utilizing the prediction of deep model to correct labels, including Joint-Optim(Tanaka et al., 2018), PENCIL (Yi & Wu, 2019), Self-Learning (Han et al., 2019). Others resort to a small set of clean validation set with meta-learning training strategies, i.e., Meta-Learning (Li et al., 2019), MLC (Zheng et al., 2021), MSLC (Wu et al., 2021), Meta-Cleaner (Zhang et al., 2019), Meta-Weight (Shu et al., 2019), FaMUS (Xu et al., 2021), MSLG (Algan & Ulusoy, 2021), Zhang, et al. (Zhang et al., 2020).

A.3.3 More Applications of DMLP

As mentioned in the manuscript, DMLP can be applied to work collaboratively with the existing LNL framework to boost performance. To further verify the effectiveness of DMLP, we also plot the accuracy curve of the proposed DMLP and its baseline methods in Fig. 8-13. It is observed that all the applications of DMLP perform consistently better over their corresponding baselines throughout the training process, especially under high-level noise cases.

A.3.4 Details of Experimental Result

• **Detailed Comparison with Coupled Methods.** In the manuscript we compare DMLP against coupled meta label correction methods MLC (Zheng et al., 2021) and MSLC (Wu et al., 2021)

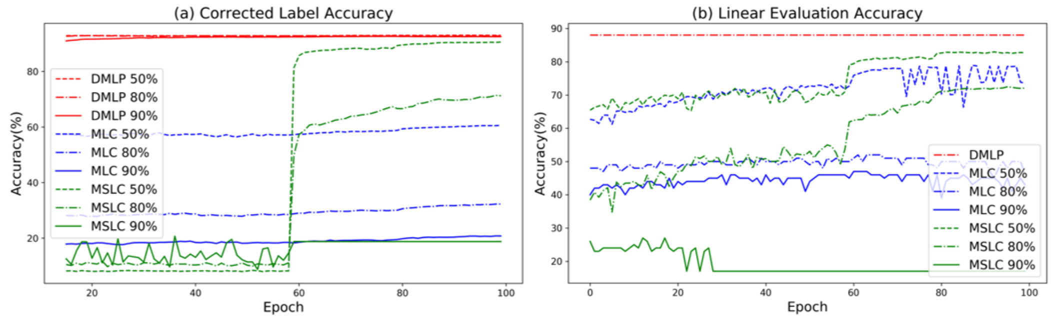

Figure 5: Comparison with two state-of-the-art coupled optimization based meta label correction methods MLC (Zheng et al., 2021) and MSLC (Wu et al., 2021) on corrected label accuracy and linear evaluation accuracy.

Table 9: Investigation on the influence of meta-purification on CIFAR-10.

Method		Noisy ratio 20%	50%	80%	90%
SimCLR-DivideMix	Best	92.2	91.2	92.1	85.7
	Last	82.8	81.3	77.0	10.9
DMLP-DivideMix	Best	**96.3**	**95.8**	**94.5**	**94.3**
	Last	**96.2**	**95.6**	**94.3**	**94.0**

with same self-supervised pretrained weights. Here we provide more detailed comparison results with the original implemented MLC and MSLC in terms of corrected label accuracy and backbone quality (revealed by linear evaluation accuracy). As shown in Fig. 5, by simplifying the complex coupled meta-learning process into individual representation learning and non-nested meta label purification, DMLP can achieve superior performance to these methods in the sense of purified label accuracy and meanwhile obtain representations of better quality, which further verifies our empirical findings of Fig. 1 in the manuscript.

• **Detailed Comparison with Decoupled Baselines.** The proposed non-nested meta label purifier plays a crucial role in DMLP. To fairly verify its superiority to existing label correction methods, we train MLC and MSLC in a decoupled way where their backbone is fixed with SimCLR self-supervised weights as in DMLP. As shown in Table 10, decoupled training scheme can largely boosts their performance, which is also in line with our empirical findings in Fig. 1 of the manuscript. Besides, it can be obviously observed that DMLP-Naive shows great advantage over the decoupled MLC and MSLC across all the settings especially under high noise, demonstrating the effectiveness of the non-nested meta label purifier.

• **Effect of the Meta-purification on DMLP-DivideMix.** We further investigate the effect of the non-nested meta label purifier on the afterward LNL framework in DMLP-DivideMix setting. To do this, we initialize a model with SimCLR and apply DivideMix to train the model as our baseline. According to the results in Table 9, the accuracy suffers from a significant performance drop when removing the non-nested meta label purifier from our pipeline (i.e, SimCLR-DivideMix), especially for severe label noise cases, this can be attributed to the DNN inevitably gradually memorizes the noisy labels when updating the backbone and classifier simultaneously. In contrast, in DMLP-DivideMix, labels purified by the meta-learner yield higher accuracy, guiding the afterward LNL framework to learn more robust and discriminative decision boundaries.

• **Detailed Comparison on the label accuracy.** In addition to the comparison of 50% and 90% noise between MLC, MSLC and DMLP on the CIFAR-10 in the manuscript, we also visualize the label accuracy curves for 20% and 90% noise. As shown in Fig. 6, DMLP consistently shows great superiority over MLC and MSLC throughout the training process. Moreover, Fig.7 shows corrected label accuracy curve of DMLP under symmetric and asymmetric noise settings on CIFAR-10, which demonstrates that high quality labels can be generated by DMLP across all noisy settings.

Table 10: Comparison with MLC and MSLC on CIFAR-10/100. "†" denotes training with fixed self-supervised pretrained ResNet-18. "*" denotes training with self-supervised pretrained ResNet-18.

Dataset / Method	Noise ratio	CIFAR-10 20%	50%	80%	90%	CIFAR-100 20%	50%	80%	90%
MLC* (Zhenget al., 2021)	Best	91.8	86.2	77.6	72.9	62.2	53.8	46.5	39.6
	Last	91.6	85.9	77.5	72.6	61.6	53.0	46.2	39.2
MSLC* (Wuet al., 2021)	Best	92.0	87.7	78.0	67.8	70.8	64.1	36.4	19.8
	Last	92.0	87.5	77.9	67.3	70.2	63.8	34.3	18.7
MLC† (Zhenget al., 2021)	Best	92.0	90.2	89.0	88.9	65.9	59.4	54.4	54.2
	Last	91.6	89.4	88.5	88.1	65.2	59.2	54.1	54.0
MSLC† (Wuet al., 2021)	Best	92.1	90.4	87.3	84.7	71.7	64.7	53.3	46.8
	Last	92.0	90.0	87.2	84.2	71.6	64.4	53.0	46.4
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**

Table 11: Comparison between DMLP-EAC, DMLP-Naive and DMLP-DivideMix on CIFAR-10/100 and Clothing1M datasets. "‡" denotes reproduced results.

Dataset / Method	Noise ratio	CIFAR-10 20%	50%	80%	90%	CIFAR-100 20%	50%	80%	90%	Clothing1M
REED(no stage-3)‡	Best	89.1	88.6	87.3	85.1	62.6	61.5	58.4	53.2	46.05
	Last	88.9	88.5	87.0	84.9	62.5	61.4	58.2	52.9	45.81
DMLP-EAC	Best	91.3	91.0	90.3	89.3	65.7	64.1	60.2	57.6	77.31
	Last	91.2	90.7	90.2	89.2	65.5	63.8	60.1	57.5	77.31
DMLP-Naive	Best	94.7	94.2	93.5	92.8	72.7	68.0	63.5	61.3	77.77
	Last	94.2	94.0	93.2	92.0	72.3	67.4	63.2	60.9	77.70
DMLP-DivideMix	Best	**96.3**	**95.8**	**94.5**	**94.3**	**79.9**	**76.8**	**68.6**	**65.8**	**78.23**
	Last	**96.2**	**95.6**	**94.3**	**94.0**	**79.4**	**76.1**	**68.5**	**65.4**	**78.23**

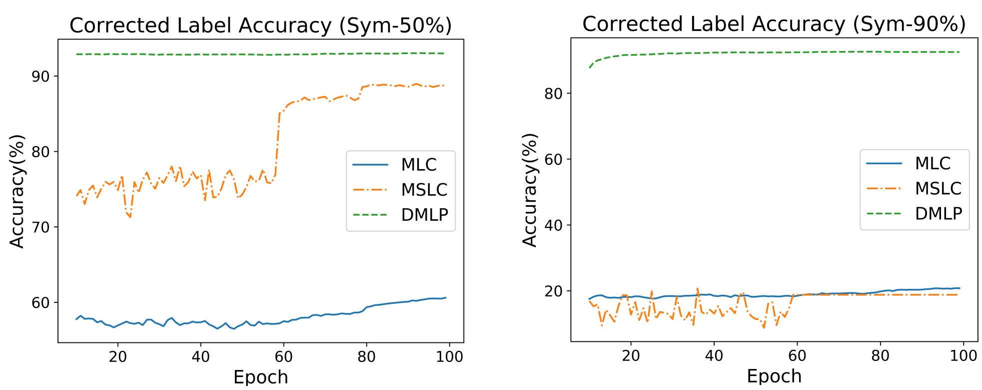

Figure 6: Comparison of corrected label accuracy curve under symmetric-50% (left), symmetric-90% (middle) noise settings on CIFAR-10.

A.3.5 EAC AS CLASSIFIER

Besides DMLP-Naive, we can also take the well-trained linear classifier $C(\cdot; w_c)$ in the non-nested meta label purifier for the test set prediction, this is termed as DMLP-EAC. As shown in Table. 11, though DMLP-EAC is only an individual linear classifier, it can also perform well especially under high noisy settings on CIFAR-10/100. Moreover, DMLP-EAC can already outperform most of state-of-the-art LNL methods by a considerable margin and achieve comparable performance to DMLP-Naive on the Clothing1M dataset, which further demonstrates that DMLP is more suitable to tackle with real-world noise. Finally, Table. 11 shows the comparison between DMLP-EAC and REED (no stage-3) (Zhang & Yao, 2020), which simply trains a linear classifier on well-established representations without extra operations. Though REED (no stage-3) achieves overall good results, DMLP-EAC can still obtain consistent performance gains over this baseline under all noise settings, especially on the high noise level.

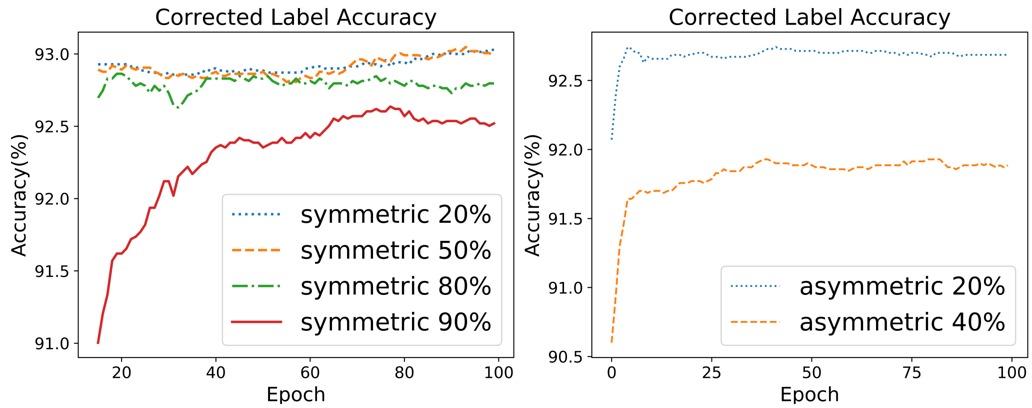

Figure 7: Corrected label accuracy curve of DMLP under symmetric (left), asymmetric (middle) noise settings on CIFAR-10.

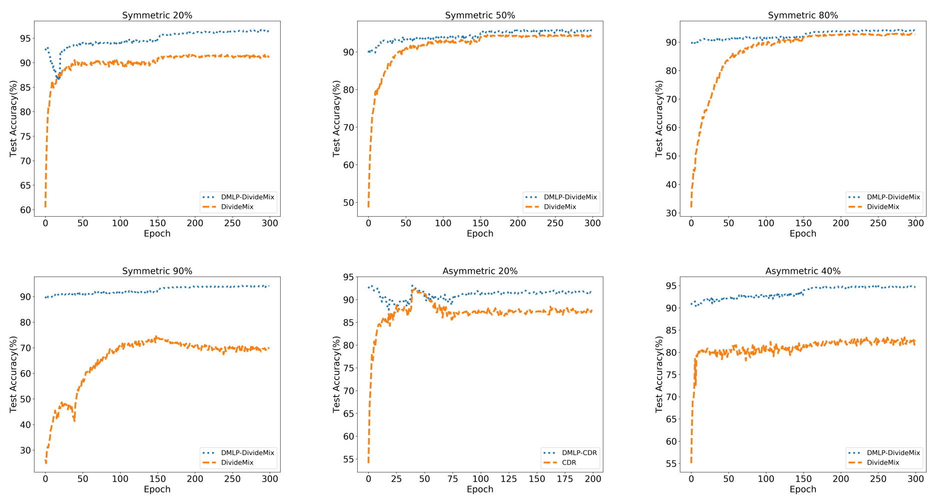

Figure 8: Accuracy curve of DMLP-DivideMix and DivideMix on CIFAR-10 under different noise settings.

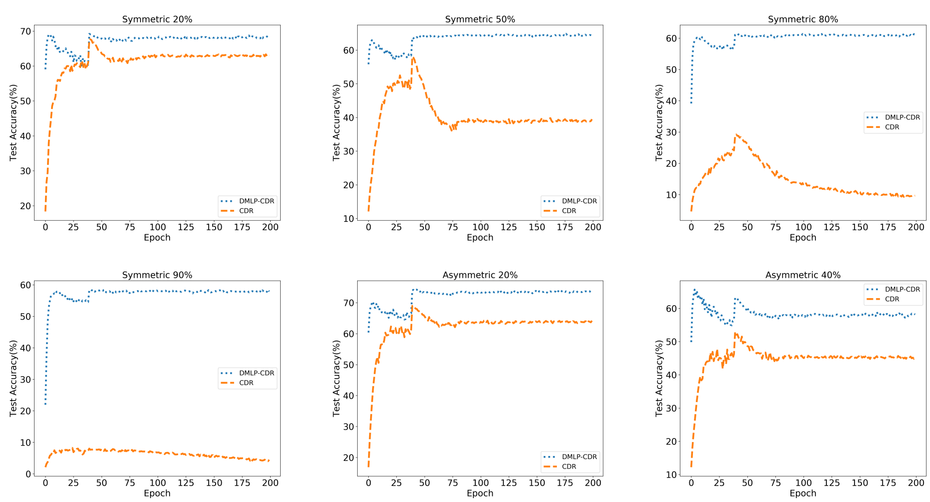

Figure 9: Accuracy curve of DMLP-DivideMix and DivideMix on CIFAR-100 under different noise settings.

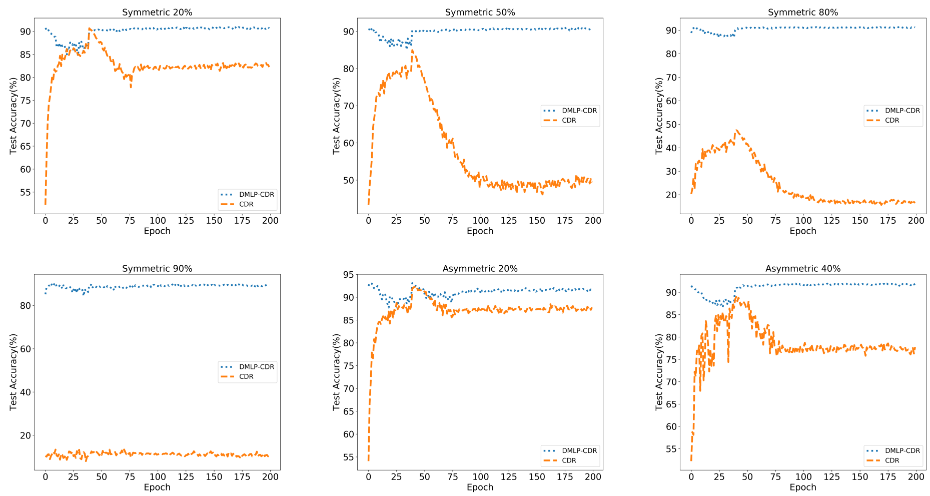

Figure 10: Accuracy curve of DMLP-CDR and CDR on CIFAR-10 under different noise settings.

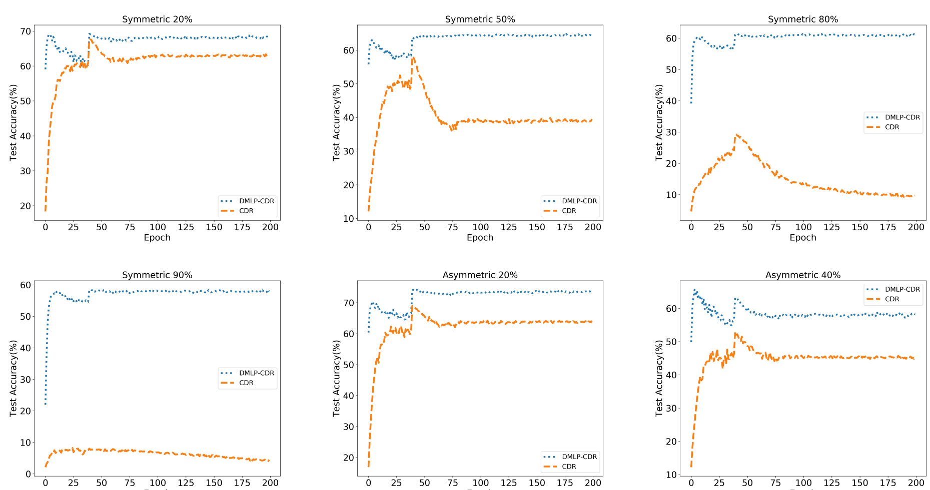

Figure 11: Accuracy curve of of DMLP-CDR and CDR on CIFAR-100 under different noise settings.

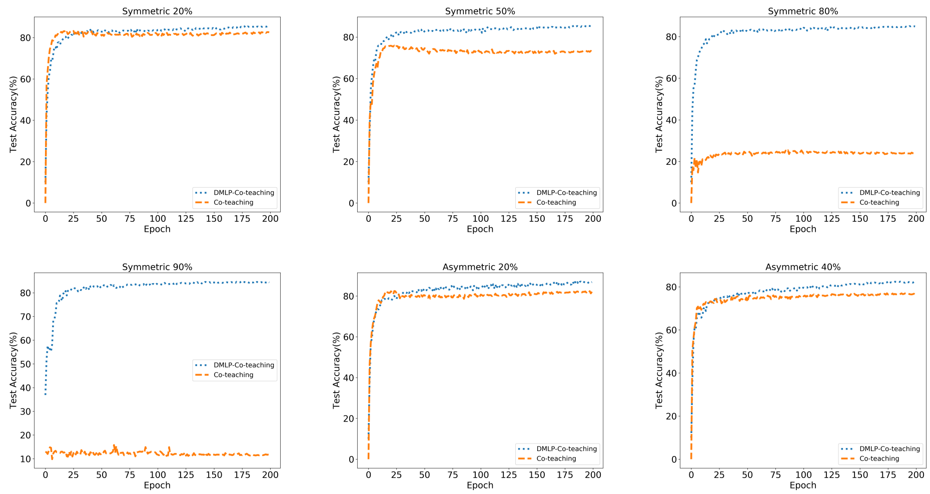

Figure 12: Accuracy curve of DMLP-Co-teaching and Co-teaching on CIFAR-10 under different noise settings.

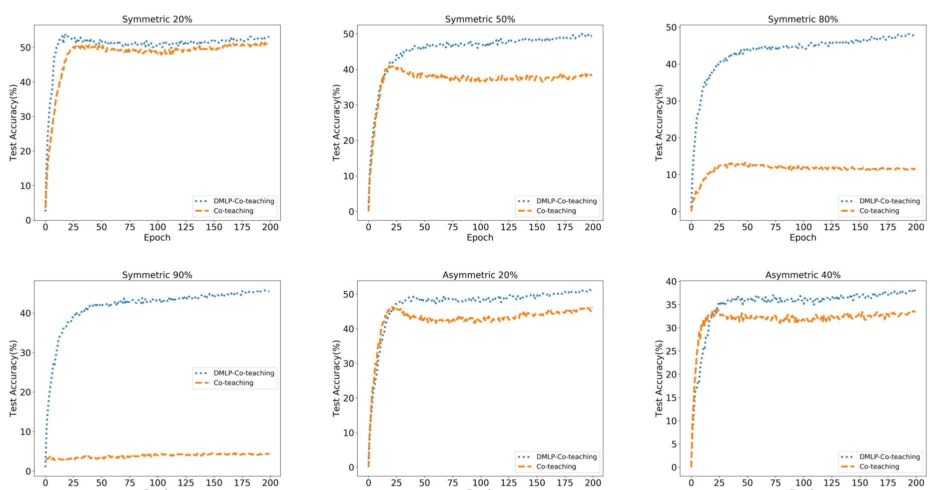

Figure 13: Accuracy curve of of DMLP-Co-teaching and Co-teaching on CIFAR-100 under different noise settings.