

# Instance-dependent Noisy-label Learning with Graphical Model Based Noise-rate Estimation

Arpit Garg<sup>1</sup>, Cuong Nguyen<sup>3</sup>, Rafael Felix<sup>1</sup>, Thanh-Toan Do<sup>2</sup>, and Gustavo Carneiro<sup>3</sup>

<sup>1</sup> Australian Institute for Machine Learning, University of Adelaide, SA, Australia

<sup>2</sup> Department of Data Science and AI, Monash University, Australia

<sup>3</sup> Centre for Vision, Speech and Signal Processing, University of Surrey, UK

**Abstract.** Deep learning faces a formidable challenge when handling noisy labels, as models tend to overfit samples affected by label noise. This challenge is further compounded by the presence of instance-dependent noise (IDN), a realistic form of label noise arising from ambiguous sample information. To address IDN, Label Noise Learning (LNL) incorporates a sample selection stage to differentiate clean and noisy-label samples. This stage uses an arbitrary criterion and a pre-defined curriculum that initially selects most samples as noisy and gradually decreases this selection rate during training. Such curriculum is sub-optimal since it does not consider the actual label noise rate in the training set. This paper addresses this issue with a new noise-rate estimation method that is easily integrated with most state-of-the-art (SOTA) LNL methods to produce a more effective curriculum. Synthetic and real-world benchmarks’ results demonstrate that integrating our approach with SOTA LNL methods improves accuracy in most cases.<sup>4</sup>

**Keywords:** Noisy-labels · Instance-dependent noise · Label noise Learning

## 1 Introduction

Deep neural networks (DNNs) have demonstrated their effectiveness in various domains, including vision [47], language [42], and medicine [37]. However, its efficacy is largely dependent on high-quality training data, which can be resource-intensive to obtain [43]. Although cost-effective labelling techniques, such as data mining [11] and crowdsourcing [39], offer cheaper alternatives, they often compromise label quality [39]. Consequently, this can lead to erroneous labels in real-world datasets [21]. This is a relevant issue because even slight inaccuracies in the labels can significantly affect the performance of DNN due to their inherent memorisation capabilities [34, 53]. This has prompted the development of algorithms resilient to noisy-label learning, with the aim of training models

<sup>4</sup> Code is available at <https://github.com/arpit2412/NoiseRateLearning>.

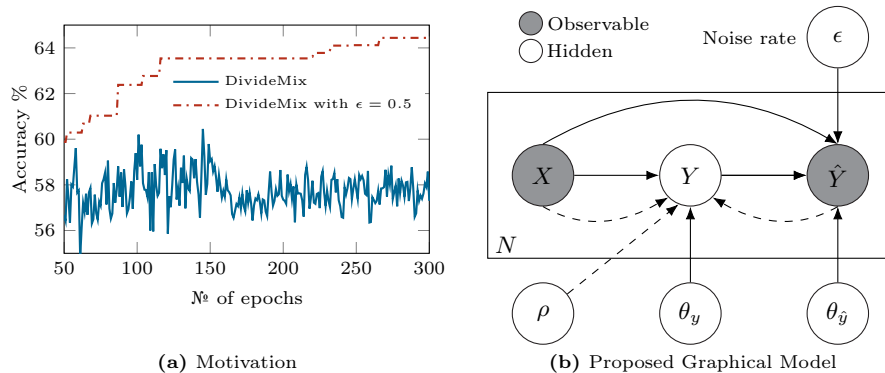
Supported by the Engineering and Physical Sciences Research Council (EPSRC) through grant EP/Y018036/1 and the Australian Research Council (ARC) through grant FT190100525.

despite inaccuracies in training data labels. Although various strategies exist, our work uniquely proposes an approach centered on the estimation of the label noise rate from the training data, which is notably absent from current methodologies.

There are two types of label noise, namely instance-independent noise (IIN) [16] and instance-dependent noise (IDN) [44]. Such type of label noise typically influences the design principles of noisy-label learning algorithms. For example, IIN focuses on mislabellings that are independent of sample information [16], where estimating the underlying label transition matrix is a common way of handling this type of noise [51]. On the other hand, in the more realistic IDN, mislabelling is due to both sample information and true class labels [44], which generally require the combination of many label noise learning techniques, such as robust loss functions [28, 56], and noisy-label sample selection [26, 58]. Of the strategies mentioned above, sample selection approaches that classify training data into clean and noisy samples have produced competitive results on many benchmarks [9, 11, 15, 21, 26]. Such sample selection techniques require the definition of a classification criterion and a selection curriculum. Many studies on this topic focus on developing new sample selection criteria, such as the small-loss hypothesis [26], which states that noisy-label samples have higher loss values than clean-label samples, particularly at the early stage of training [1]. Another criterion type is the feature-based one. An example of such criterion is FINE [21] that discriminates clean and noisy-label samples via the distance to class-specific eigenvectors. In this technique, clean-label samples tend to lie closer to the class-specific dominant eigenvector of the latent representations than noisy-label samples. Another type of criterion is proposed in SSR [11], which introduces a selection criterion based on the K-nearest-neighbor (KNN) classification in the feature space. Furthermore, CC [58] uses a two-stage sampling procedure, including class-level feature clustering followed by a consistency score. An equally important problem in sample selection is the definition of the curriculum for selecting clean training samples, but it has received comparatively less attention.

The sample selection curriculum defines a threshold to be used with one of the criteria listed above to classify the training samples as clean or noisy at each training epoch [45]. For example, the threshold can be fixed to an arbitrary clustering score that separates clean and noisy samples [26], but this strategy does not account for the proportion of label noise in the training set, nor does it consider the dynamics of the selection of noisy-label samples during the training. The consideration of such dynamics has been studied in [16, 49], which defined a curriculum of the noisy-label sampling rate  $R(t)$  as a function of the training epoch  $t \in \{1, \dots, T\}$ . The curriculum  $R(t)$  defines a sampling rate close to 100% of the training set at the beginning of the training, which is then reduced to arbitrarily low rates at the end of the training. In practice, the function  $R(t)$  is predefined [16] or learned by weighting a set of basis functions [49].

Although generally effective, these techniques do not consider the label noise rate estimated from the training set, making them vulnerable to over-fitting (if too many noisy-label samples are classified as clean) or under-fitting (if informative clean-label samples are classified as noisy). It can be argued that the estimation

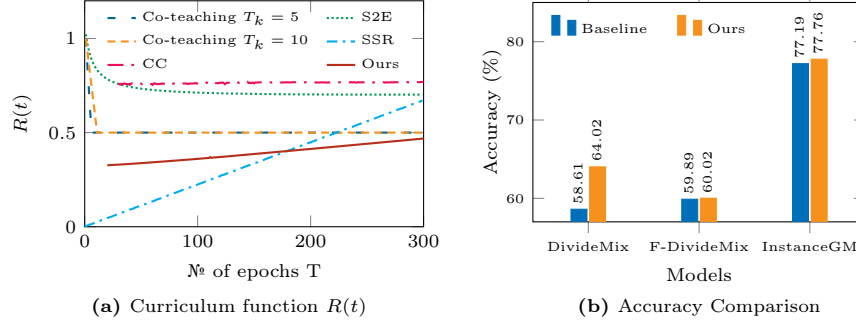


**Fig. 1:** (a) Comparison of test accuracy % (as a function of training epoch) between the original DivideMix [26] (solid, blue curve) and our modified DivideMix (dashed, red curve) that selects the clean and noisy data based on a fixed noise rate  $R(t) = 1 - \epsilon = 50\%$  using the small-loss criterion on CIFAR100 [22] at 0.5 IDN [44]; (b) The proposed probabilistic graphical model that generates noisy-label  $\hat{Y}$  conditioned on the image  $X$ , the latent clean-label  $Y$  and noise rate  $\epsilon$ , where forward pass (solid lines) is parameterised by  $\theta_y, \theta_{\hat{y}}$  and  $\epsilon$  representing the generation step, and the backward pass (dashed lines) is parameterised by  $\rho$ .

of the label transition matrix [6, 7, 48, 51] aims to recover the noise rate affecting pairwise label transitions. However, label transition matrix techniques follow a quite different strategy compared to sample selection methods, where their main challenge is the general underconstrained aspect of the matrix estimation, making them sensitive to large noise rates and not scalable to a high number of classes [39]. Although several methods address label noise, none estimates the label noise rate directly from the training data to guide sample selection. Addressing this gap is the primary challenge of our study.

To underscore the importance of using the noise rate for sample selection during training, we experiment with CIFAR100 [22] at an IDN rate  $\epsilon = 50\%$  [44] (noise rate specifications and other details are explained in Sec. 4). We use DivideMix [26] and replace its sample selection (based on an arbitrary clustering score [10, 32, 41]) by a thresholding process that classifies the  $R(t) = 1 - \epsilon = 50\%$  largest loss samples as noisy, and the remaining ones as clean in all training epochs  $t \in \{1, \dots, T\}$ . This sample selection is used for the semi-supervised learning of DivideMix [26]. As shown in Fig. 1a, the new sample selection approach based on the “manually provided” noise rate (dashed red curve) improves 6% in terms of prediction accuracy compared to the original DivideMix [26] (solid blue curve) which relies on arbitrary thresholding. Similar conclusions can be achieved with other methods that apply sample selection strategies to address the noisy-label learning problem, as shown in the experiments.

In this paper, we introduce a new sample selection strategy centered on estimating the label noise rate of the training set. Our strategy is based on our novel noisy-label learning graphical model illustrated in Fig. 1b. This model can be



**Fig. 2:** (a) Visual comparison of different  $R(t)$  on CIFAR100 [22] at 0.5 IDN [44]: (i) Co-teaching [16] with curriculum based on hyper-parameter  $T_k \in \{5, 10\}$ , where  $R(t) = 1 - \tau \cdot \min(t/T_k, 1)$  with  $\tau = \epsilon = 0.5$  being manually set; (ii) S2E [49], where  $R(t)$  is estimated with a bi-level optimisation; (iii) SSR [11], where  $R(t)$  is based on a relabelling pattern, (iv) CC [58] following the small-loss hypothesis [26] and cosine similarity for  $R(t)$ ; and (v) ours. (b) Accuracy of noisy-label learning methods on CIFAR100 [22] at 0.5 IDN [44], including DivideMix [26], FINE [21] and InstanceGM [15], without (left, blue) and (right, orange) integration of our proposed graphical model for estimating the noise rate  $\epsilon$  and sample selection rate based on  $\epsilon$ .

seamlessly integrated with state-of-the-art (SOTA) noisy-label learning techniques, providing them with a precise noise rate estimate and subsequently refining the sample selection curriculum. In particular, our model’s curriculum is not restrained by predefined  $R(t)$  functions [16, 49], but rather relies on a dynamically estimated noise rate sourced directly from the training set, as depicted in Fig. 2a. Our method dynamically estimates training set noise rates to mitigate overfitting and underfitting, seamlessly integrates with existing algorithms, and serves as a robust foundation for future noisy-label learning research. To summarise, our main contributions include:

- An innovative noisy-label learning graphical model, shown in Fig. 1b, that not only estimates but also leverages the noise rate from the training dataset to produce a refined sample selection curriculum for SOTA LNL methods.
- A simple and synergistic integration strategy of our novel graphical model with several SOTA noisy-label learning algorithms, such as DivideMix [26] and SSR [11], to improve their sample selection effectiveness, leading to an increased test accuracy, as shown in Fig. 2b.

Empirical evaluations show the role of our proposed sample selection methodology in improving the efficacy of leading noisy-label learning algorithms across various synthetic and real-world IDN benchmarks. Although our primary focus lies in addressing IDN problems, because of its more challenging and realistic nature, we have also applied our model to IIN problems, with detailed results in Tab. 6 and discussion in Appendix 3 of the supplementary material. With our innovative approach, we aim to improve the accuracy of SOTA LNL methods.

## 2 Related Work

It is well-known that DNNs suffer from overfitting when trained with noisy-labels [52], resulting in poor generalisation [28, 55]. To mitigate the problems related to noisy-label learning, several techniques have been developed, including *noise-robust loss functions* [30], *noise-label sample selection* [21, 26, 40, 58, 59], *meta-learning* [57], and *re-labeling* followed by *sample selection* [11]. These techniques can handle many types of label noise (e.g., symmetric and asymmetric), but the most challenging type, i.e., IDN [39], tends to be successfully addressed with methods that have a training stage based on sample selection techniques [21, 26, 40, 58, 59]. These sample selection techniques separate clean and noisy-label samples, where clean samples are treated as labelled, and noisy samples are discarded [8, 16, 18, 49] or treated as unlabeled samples [9, 15, 26, 35, 59] for semi-supervised learning [3]. Another solution consists of aligning clean and noisy samples with an information fusion method [19]. A major limitation of these approaches is the overly simplistic design of the curriculum to select clean and noisy-label samples during training. It is either predetermined [16, 18] or learned from a set of predetermined basis functions [49]. We argue that the design of a curriculum based on an estimated noise rate would benefit such methods, as motivated by Fig. 1a, but to the best of our knowledge, such estimation has not been explored by previous methods. The closest idea explored is based on the estimation of the type of noise instead of the noise rate [45]. It is possible to argue that the estimation of the label transition matrices [5, 7, 44, 60], composed of the learned noise rate that affects the transition between pairs of labels, is a way to estimate the label noise rate. However, they portray results of comparatively lower accuracy for large real-world datasets or for IDN problems [39, 44]. A possible reason for these poorer results is that label transition approaches suffer from identifiability issues [14], where any clean-label distribution assignment is acceptable as long as the distribution of observed labels can be reconstructed [13]. This makes the identification of the actual underlying clean-labels challenging. One solution is to consider the use of multiple annotations per sample to help analyse agreements and disagreements for improved identification of clean patterns [14, 31]. However, the annotation of datasets with a single label per training sample is already challenging; the complexity is significantly larger to acquire multiple labels per sample.

An alternative technique to handle IDN problems is based on graphical models that represent the relationship between various observed and latent variables [15, 45, 50]. The approaches in [15, 50] use graphical models that rely on a generative approach to produce noisy-labels from the respective image features and latent clean-labels. However, previous graphical models do not take into account the underlying noise rate parameter during modeling. Our work is the first graphical model approach to estimate the noise rate of the dataset. In addition, a close examination of existing literature reveals a trend: when traditional graphical models are integrated with SOTA noisy-label learning methods, the outcomes often fall short of expectations in terms of accuracy [2, 15, 50]. This observation underscores a strong point of our work. Our approach is uniquely tailored for

easy integration with current SOTA noisy-label learning methods for sample selection, which, in turn, improves over the SOTA classification accuracy.

### 3 Method

In this section, we present our new graphical model that estimates the noise rate, which will be used in the sample selection process. Let  $\mathcal{D} = \{(x_i, \hat{y}_i)\}_{i=1}^N$  be the noisy-label training set containing  $d$ -dimensional data vector  $x_i \in \mathcal{X} \subseteq \mathbb{R}^d$  and its respective  $C$ -dimensional one-hot encoded observed (potentially corrupted) label  $\hat{y}_i \in \hat{\mathcal{Y}} = \{\hat{y} : \hat{y} \in \{0, 1\}^C \wedge \mathbf{1}_C^\top \hat{y} = 1\}$ , where  $\mathbf{1}_C$  is a vector of ones with  $C$  dimensions. The aim is to estimate the label noise rate  $\epsilon$ , used for the generation of noisy-label training data from the observed training dataset  $\mathcal{D}$  and integrate this label noise rate into the sample selection strategy.

#### 3.1 Graphical Model

We portray the generation of noisy-label via the probabilistic graphical model shown in Fig. 1b. The observed random variables, denoted by shaded circles, are data  $X$  and the corresponding noisy-label  $\hat{Y}$ . We also have one latent variable, namely: the clean-label  $Y$ . Under our proposed modeling assumption, a noisy-label of a data instance can be generated as follows:

- sample an instance from  $p(X)$ , i.e. :  $x \sim p(X)$
- sample a clean-label from the clean-label distribution:  
 $y \sim \text{Categorical}(Y; f_{\theta_y}(x))$
- sample a noisy-label from the noisy-label distribution:  
 $\hat{y} \sim \text{Categorical}(\hat{Y}; \epsilon \times f_{\theta_{\hat{y}}}(x, f_{\theta_y}(x)) + (1 - \epsilon) \times y),$

where  $\text{Categorical}(\cdot)$  denotes a categorical distribution,  $f_{\theta_y} : \mathcal{X} \rightarrow \Delta_{C-1}$  and  $f_{\theta_{\hat{y}}} : \mathcal{X} \times \Delta_{C-1} \rightarrow \Delta_{C-1}$  denote two classifiers for the clean-label  $Y$  and noisy-label  $\hat{Y}$ , respectively, with  $\Delta_{C-1} = \{s : s \in [0, 1]^C \wedge \mathbf{1}_C^\top s = 1\}$  being the  $(C - 1)$ -dimensional probability simplex. According to the data generation process shown in Fig. 1b,  $\epsilon$  corresponds to  $\mathbb{E}_{(x, \hat{y}) \sim p(X, \hat{Y})} [P(\hat{y} \neq y | x)]$ , which is the label noise rate of the training dataset of interest. Our aim is to infer the parameters  $\theta_y, \theta_{\hat{y}}$  and  $\epsilon$  from a noisy-label dataset  $\mathcal{D}$  by maximising the following log-likelihood:

$$\max_{\theta_y, \theta_{\hat{y}}, \epsilon} \mathbb{E}_{(x_i, \hat{y}_i) \sim \mathcal{D}} [\ln p(\hat{y}_i | x_i; \theta_y, \theta_{\hat{y}}, \epsilon)] = \max_{\theta_y, \theta_{\hat{y}}, \epsilon} \mathbb{E}_{(x_i, \hat{y}_i) \sim \mathcal{D}} \left[ \ln \sum_{y_i} p(\hat{y}_i, y_i | x_i; \theta_y, \theta_{\hat{y}}, \epsilon) \right]. \quad (1)$$

Due to the presence of the clean-label  $y_i$ , it is difficult to directly evaluate the log-likelihood in Eq. (1). Therefore, we employ the *expectation - maximisation* (EM) algorithm [10] to maximise the log-likelihood. The main idea of the EM algorithm is to (i) construct a tight lower bound of the likelihood in Eq. (1) by estimating the posterior of the latent variable  $Y$  (known as *expectation step*) and (ii) maximise that lower bound (known as *maximisation step*). Formally, let

$q(y_i|x, \hat{y}; \rho)$  be an arbitrary distribution over a clean-label  $y_i$ . The evidence lower bound (ELBO) of the log-likelihood in Eq. (1) can be obtained through Jensen’s inequality and is presented as follows:

$$\begin{aligned} Q(\theta_y, \theta_{\hat{y}}, \epsilon, \rho) &= \mathbb{E}_{(x_i, \hat{y}_i) \sim \mathcal{D}} [\mathbb{E}_{q(y_i|x_i, \hat{y}_i; \rho)} [\ln p(\hat{y}_i|x_i, y_i; \theta_y, \theta_{\hat{y}}, \epsilon)] \\ &\quad - \text{KL}[q(y_i|x_i, \hat{y}_i; \rho) \| p(y_i|x_i; \theta_y, \theta_{\hat{y}}, \epsilon)]] \\ &= \mathbb{E}_{(x_i, \hat{y}_i) \sim \mathcal{D}} [\mathbb{E}_{q(y_i|x_i, \hat{y}_i; \rho)} [\ln p(\hat{y}_i|x_i, y_i; \theta_y, \theta_{\hat{y}}, \epsilon) \\ &\quad + \ln p(y_i|x_i; \theta_y)] + \mathbb{H}[q(y_i|x_i, \hat{y}_i; \rho)]], \end{aligned} \quad (2)$$

where  $\text{KL}[q \| p]$  is the Kullback – Leibler divergence between distributions  $q$  and  $p$  and  $\mathbb{H}(q)$  is the entropy of the distribution  $q$ . The EM algorithm is then carried out iteratively by alternating the following two steps:

*E step:* we maximise the ELBO in (2) w.r.t.  $q(y_i|x_i, \hat{y}_i; \rho)$ . Theoretically, such an optimisation results in  $\text{KL}[q(y_i|x_i, \hat{y}_i; \rho) \| p(y_i|x_i, \hat{y}_i)] = 0$  or  $q(y_i|x_i, \hat{y}_i; \rho) = p(y_i|x_i, \hat{y}_i)$ . This is equivalent to estimating the posterior of the clean-label  $y_i$  given noisy-label data  $(x_i, \hat{y}_i)$ . Obtaining the posterior  $p(y_i|x_i, \hat{y}_i)$  is, however, intractable for most deep-learning models. To mitigate such an issue, we follow the *variational EM* approach [33] by employing an approximate posterior  $q(y_i|x_i, \hat{y}_i; \rho^{(t)})$  that is the closest to the true posterior  $p(y_i|x_i, \hat{y}_i)$ , where:

$$\rho^{(t)} = \arg \max_{\rho} Q(\theta_y^{(t)}, \theta_{\hat{y}}^{(t)}, \epsilon^{(t)}, \rho), \quad (3)$$

with the superscript  $^{(t)}$  denoting the parameters at the  $t$ -th iteration. Although this results in a non-tight lower bound of the log-likelihood in Eq. (1), it does increase the variational bound  $Q$ .

*M step:* we maximise the ELBO in Eq. (2) w.r.t.  $\theta_y, \theta_{\hat{y}}$  and  $\epsilon$  given  $\rho^{(t)}$  obtained in the E step:

$$\theta_y^{(t+1)}, \theta_{\hat{y}}^{(t+1)}, \epsilon^{(t+1)} = \arg \max_{\theta_y, \theta_{\hat{y}}, \epsilon} Q(\theta_y, \theta_{\hat{y}}, \epsilon, \rho^{(t)}). \quad (4)$$

The estimated noise rate  $\epsilon$  can then be integrated into certain noisy-label algorithms to train the models of interest as mentioned in Sec. 1. In addition, the inference of noise rate  $\epsilon$  might be associated with the identifiability issue when estimating the clean-label  $Y$  [29], i.e., there exists multiple sets of  $\rho$  and  $\theta_y$ , where each set can explain the observed noisy-label data equally well. Such issues are addressed in the following subsection.

### 3.2 Sample Selection

The identifiability issue when inferring the clean-label  $Y$  from noisy-label data  $(X, \hat{Y})$  can be mitigated either by acquiring multiple noisy-labels [29] or introducing additional constraints, such as small loss hypothesis [16] or FINE [21]. Since requesting additional noisy-labels per training sample is not always possible, we

**Algorithm 1** Proposed noisy-label learning that relies on the estimation of noise rate  $\epsilon$  to build a sample selection curriculum.

---

```

1: procedure NOISE RATE ESTIMATION AND INTEGRATION( $\mathcal{D}, T, \lambda$ )
2:    $\triangleright \mathcal{D} = \{(x_i, \hat{y}_i)\}_{i=1}^N$ : training set with noisy-label data  $\triangleleft$ 
3:    $\triangleright T$ : number of epochs  $\triangleleft$ 
4:    $\triangleright \lambda$ : a hyper-parameter  $\triangleleft$ 
5:    $\triangleright$  Notation:  $\Theta^{(t)} = (\theta_y^{(t)}, \theta_{\hat{y}}^{(t)}, \epsilon^{(t)})$   $\triangleleft$ 
6:   Initialise  $\theta_y^{(1)}, \theta_{\hat{y}}^{(1)}, \epsilon^{(1)}$  and  $\rho^{(0)}$ 
7:    $\theta_y^{(1)} \leftarrow \text{WARM UP}(\mathcal{D}, \theta_y^{(1)})$ 
8:    $t \leftarrow 0$ 
9:   for  $n_{\text{epoch}} = 1 : T$  do
10:    for each mini-batch  $\mathcal{S}$  in SHUFFLE( $\mathcal{D}$ ) do
11:       $t \leftarrow t + 1$ 
12:       $\mathcal{S}_{\text{clean}}, \mathcal{S}_{\text{noisy}} \leftarrow \text{SAMPLE SELECTION}(\mathcal{S}, \theta_y^{(t)}, \epsilon^{(t)})$   $\triangleright \text{Eq. (7)}$ 
13:       $\triangleright$  Variational E-step as in Eq. (3)  $\triangleleft$ 
14:       $\rho^{(t)} \leftarrow \text{E STEP}(\mathcal{S}, \theta_y^{(t)}, \theta_{\hat{y}}^{(t)}, \epsilon^{(t)}, \rho^{(t-1)})$ 
15:       $\triangleright$  M-step as in Eq. (8)  $\triangleleft$ 
16:       $\theta_y^{(t+1)}, \theta_{\hat{y}}^{(t+1)}, \epsilon^{(t+1)} \leftarrow \text{M STEP}(\mathcal{S}_{\text{clean}}, \mathcal{S}_{\text{noisy}}, \theta_y^{(t)}, \theta_{\hat{y}}^{(t)}, \epsilon^{(t)}, \rho^{(t)}, \lambda)$ 
17: return  $\theta_y$   $\triangleright$  parameter of the clean-label classifier

```

---

follow the latter approach by imposing a constraint (scoring term), denoted as  $L(\theta_y, \epsilon^{(t)})$ , over  $\theta_y$  in the M step via a sample selection approach based on the estimated noise rate  $\epsilon^{(t)}$ . Formally, we propose a new curriculum when selecting samples as follows:

$$R(t) = 1 - \epsilon^{(t)}. \quad (5)$$

In the simplest case, such as Co-teaching [16] or FINE [21], the constraint for  $\theta_y$  can be written as:


$$L(\theta_y, \epsilon^{(t)}) = \sum_{(x_i, \hat{y}_i) \in \mathcal{S}_{\text{clean}}} \text{KL} [\text{Categorical}(Y; \hat{y}) \| \text{Categorical}(Y; f_{\theta_y}(x_i))], \quad (6)$$

where:

$$\begin{aligned}
\mathcal{S}_{\text{clean}} &= \{(x_i, \hat{y}_i) : (x_i, \hat{y}_i) \in \mathcal{D} \wedge z(x_i, y_i) \leq m\}, \\
m &\in \{m : \Pr(z(x_i, y_i) \leq m) \geq R(t) \wedge \Pr(z(x_i, y_i) \geq m) \geq 1 - R(t)\}, \\
\mathcal{S}_{\text{noisy}} &= \mathcal{D} \setminus \mathcal{S}_{\text{clean}}.
\end{aligned} \quad (7)$$

with  $R(t)$  defined in (5), and  $z(x, y)$  representing the score of a criterion (e.g., loss [26, 58, 59], distance to the largest eigenvectors [21], or KNN scores [11]). Intuitively, the loss in Eq. (6) is simply the cross-entropy loss on the  $\lfloor R(t) \times N \rfloor$  samples that have smallest scores (with  $\lfloor \cdot \rfloor$  being the floor function). One can also extend to other SOTA models by replacing the loss  $L$  accordingly. For example,





**Fig. 3:** Our training algorithm uses an existing noisy-label classifier (e.g., DivideMix [26]) parameterised by  $\theta_y$  as the clean-label model  $p(Y|X; \theta_y)$ . The generation of noisy-label (given  $X$  and  $Y$ ) is carried out by a model parameterised by noise rate  $\epsilon$  and  $\theta_{\hat{y}}$ . The noisy-label classifier is based on a sample-selection mechanism that uses a curriculum  $R(t) = 1 - \epsilon^{(t)}$ .

if DivideMix [26] is used as a base model to constrain  $\theta_y$ ,  $L$  will include two additional terms: loss on un-labeled data and regularisation using *mixup* [54].

### 3.3 Training and Testing

Given the sample selection approach in Sec. 3.2, the M step in Eq. (4) is slightly modified by integrating the constraint in Eq. (6), which can be written as

$$\theta_y^{(t+1)}, \theta_{\hat{y}}^{(t+1)}, \epsilon^{(t+1)} = \arg \max_{\theta_y, \theta_{\hat{y}}, \epsilon} Q(\theta_y, \theta_{\hat{y}}, \epsilon, \rho^{(t)}) - \lambda L(\theta_y, \epsilon^{(t)}), \quad (8)$$

where  $\lambda$  is a hyper-parameter and  $L$  is defined in Eq. (6). The training procedure is summarised in Algorithm 1 and visualised in Fig. 3. In the implementation, we integrate the proposed method into existing models, such as DivideMix [26] or FINE [21]. Note that the clean-label classifier  $f_{\theta_y}(\cdot)$  is also the clean classifier of the base model.

## 4 Experiments

We show extensive experiments in several noisy-label synthetic benchmarks with CIFAR100 [22], and real-world benchmarks, including CNWL’s red mini-ImageNet [17], Clothing1M [45] and mini-WebVision [27]. Sec. 4.2 describes implementation details. We evaluate our approach by plugging SOTA models into  $p(y|x; \theta_y)$ , defined in Sec. 3, with results being shown in Sec. 4.3, and ablation studies in Sec. 4.4.

#### 4.1 Dataset Descriptions

We follow the existing literature to generate the IDN labels [44] for **CIFAR100** [22]. The dataset contains 50,000 training images and 10,000 testing images of shape  $32 \times 32 \times 3$ . The dataset is class-balanced with 100 different categories and the IDN rates considered are 0.2, 0.3, 0.4 and 0.5 [44]. Another important IDN benchmark dataset is the **red mini-ImageNet**, a real-world dataset from CNWL [17]. In this dataset, there are 100 distinct classes, with each class containing 600 images. The images and their corresponding noisy-labels have been crawled from the internet at various controllable label noise rates. For a fair comparison with the existing literature [9, 15, 46], we have resized the images to  $32 \times 32$  pixels from the  $84 \times 84$  original pixel settings. We show results with 40%, 60% and 60% noise rates on this dataset. **Clothing1M** [45] is a real-world clothing classification dataset that contains 1 *million* images with an estimated noise rate of 38.5%. The dataset contains 14 different categories, and the labels are generated from surrounding texts. Images in this dataset are of different sizes, and we follow the resize structure suggested in [21, 26]. This dataset also contains 50k manually validated clean training images, 14k images for the validation set, and 10k testing images. We have not used the clean training, validation set, or any extra training images while training. Only the testing set is used for evaluation. **Mini-WebVision** [26] contains 65,944 images with their respective labels from the 50 different initial categories from the WebVision [27], with image size of  $256 \times 256$  pixels. Following the evaluation process commonly used in this benchmark, 50 categories from the ILSVRC12 [23] dataset are also used for testing.

#### 4.2 Implementation

All methods are implemented in Pytorch [36] and use one NVIDIA RTX 3090 card for training and testing. As mentioned in the original papers, hyperparameter settings are kept the same for the baselines used in the proposed algorithm. All classifier architectures are also kept the same as the baseline models. A random initialisation of noise rate parameter  $\epsilon$  with the sigmoid as its activation function is employed for all experiments to maintain the fairness of the comparisons with other approaches. The value of  $\lambda$  in Eq. (8) is set to 1 for all the cases. We integrate many SOTA approaches [11, 15, 21, 26, 58, 59] into our graphical model, as explained in Sec. 3.3. For CIFAR100 [22] with IDN [44], we integrate DivideMix [26], C2D [59], CC [58], and InstanceGM [15] into our model, given their superior performance across various noise rates. Additionally, we also use F-Dividemix from FINE [21], as shown in Fig. 2b. For red mini-ImageNet [17], we test our proposed approach with and without DINO self-supervision [4]. For the implementation without self-supervision, we use DivideMix [26] and InstanceGM [15], and for the self-supervised version, we only use InstanceGM [15]. The models for Clothing1M [45] and mini-WebVision [26] (validation on ImageNet [23]) are trained using DivideMix [26], SSR [11], C2D [59] and CC [58]. Additional implementation details, empirical analysis, and computational analysis are presented in the supplementary material in Appendices 1, 2 and 4, respectively.

**Table 1:** Test accuracy % on (*left*) CIFAR100 [22] under different IDN [45], and (*right*) red mini-ImageNet [17]. Other models’ results are from [7, 15, 46]. Here, we integrate DivideMix [26], C2D [59], CC [58] and InstanceGM [15] into our proposed model for CIFAR100 [22]. For red mini-ImageNet [17], we present the results with and without self-supervision [4]. We integrate DivideMix [26] and InstanceGM [15] into our model, with the latter tested with and without self-supervision.

Method	Noise Rates - IDN			
	0.2	0.3	0.4	0.5
CE [50]	30.42	24.15	21.45	14.42
PartT [44]	65.33	64.56	59.73	56.80
kMEIDTM [7]	69.16	66.76	63.46	59.18
DivideMix [26]	77.03	76.33	70.80	58.61
<b>DivideMix-Ours</b>	77.42	77.21	72.41	64.02
	(+0.39)	(+0.88)	(+1.61)	(+5.41)
C2D [59]	78.61	78.18	72.89	63.19
<b>C2D-Ours</b>	79.07	78.59	73.31	65.28
	(+0.46)	(+0.41)	(+0.42)	(+2.09)
CC [58]	79.61	77.56	76.58	63.19
<b>CC-Ours</b>	79.72	78.71	77.38	67.53
	(+0.11)	(+1.15)	(+0.80)	(+4.34)
InstanceGM [15]	79.69	79.21	78.47	77.19
<b>InstanceGM-Ours</b>	79.61	79.40	79.52	78.21
	(-0.08)	(+0.19)	(+1.05)	(+1.02)

Method	Noise rate		
	0.4	0.6	0.8
CE [46]	42.70	37.30	29.76
MixUp [54]	46.40	40.58	33.58
MentorMix [17]	47.14	43.80	33.46
FaMUS [46]	51.42	45.10	35.50
DivideMix [26]	46.72	43.14	34.50
<b>DivideMix-Ours</b>	50.70	45.11	37.44
	(+3.98)	(+1.97)	(+2.94)
InstanceGM [15]	52.24	47.96	39.62
<b>InstanceGM-Ours</b>	56.61	51.40	43.83
	(+4.37)	(+3.44)	(+4.21)
<b>With self-supervised learning</b>			
PropMix [9]	56.22	52.84	43.42
InstanceGM-SS [15]	56.37	53.21	44.03
<b>InstanceGM-SS-Ours</b>	58.29	53.60	45.47
	(+1.92)	(+0.39)	(+1.44)

**Table 2:** Final estimated noise rate  $\epsilon$  for the Tab. 1, (*left*) CIFAR100 [22] under various IDN [44], and (*right*) red mini-ImageNet [17] with and without self-supervision.

Estimated noise rates				
Method	Actual noise rate			
	0.2	0.3	0.4	0.5
DivideMix-Ours	0.18	0.34	0.47	0.53
C2D-Ours	0.19	0.33	0.38	0.52
CC-Ours	0.21	0.34	0.41	0.54
InstanceGM-Ours	0.23	0.37	0.42	0.47

Estimated noise rates			
Method	Actual noise rate		
	0.4	0.6	0.8
DivideMix-Ours	0.39	0.58	0.73
InstanceGM-Ours	0.38	0.55	0.74
InstanceGM-SS-Ours	0.48	0.53	0.69

### 4.3 Comparison with SOTA

This section compares our approach with SOTA methods on datasets with the IDN settings and noisy real-world settings. The bold text in tables indicate SOTA results, and our results are in the greyed rows.

**Synthetic Instance-Dependent Noise** The comparison between various baselines and our proposed method on CIFAR100 [22] with IDN [44] is shown in Tab. 1 (*left*) with the noise rate ranging from 0.2 to 0.5. It is worth noting that using our proposed model with DivideMix [26], C2D [59], CC [58] and InstanceGM [15] improve their performance in almost all cases, particularly with large noise rates. Tab. 2 (*left*) also shows the final noise rate  $\epsilon$  estimated by our model which is reasonably close to the simulated noise rate displayed in the table’s header.

**Table 3:** Test accuracy (%) of competing methods, and final estimated noise rate  $\epsilon$  on Clothing1M [45]. Also, we have not considered competing models that rely on a clean or validation set whilst training. We integrate DivideMix [26], C2D [59], SSR [11] and CC [58] into our model. DivideMix [26] and CC [58] shows the locally reproduced results. Mentioned noise rate by [45] is 0.385.

Method	Test accuracy (%)	Estimated noise rate
OT-Filter [12]	74.50	
ELR+ with C2D [59]	74.58	
AugDesc [35]	75.11	
NCE [25]	75.30	
DivideMix [26]	74.32	
<b>DivideMix-Ours</b>	74.41	(+0.09)
C2D [59]	74.58	
<b>C2D-Ours</b>	74.71	(+0.13)
SSR (class-imbalance) [11]	74.12	
<b>SSR-Ours</b>	74.20	(+0.08)
CC [58]	75.24	
<b>CC-Ours</b>	<b>75.31</b>	(+0.07)

**Real-World Noise** We also evaluate our proposed method on various real-world noisy settings regarding test accuracy and estimated noise rates  $\epsilon$  in Tabs. 1 to 4. Similarly to the synthetic IDN in Tab. 1 (*left*), the results show that existing noisy-label robust methods can be easily integrated with our model to outperform current SOTA results for real-world noisy-label datasets. Tabs. 1 and 2 (*right*) shows the results on red mini-ImageNet using two configurations, including cases without self-supervision (top part of the table) and with self-supervision (bottom part of the table). The self-supervision DINO pre-training [4] relies only on images from red mini-ImageNet to enable a fair comparison with existing baselines [9, 15]. Results from Tab. 1 (*right*) demonstrate that our approach improves the performance of SOTA methods by a considerable margin in all cases. In fact, using estimated noise rate  $\epsilon$  (Tab. 2 (*right*)) while training InstanceGM [15] without self-supervision shows better performance than existing self-supervised baselines at 0.4 noise rate. Moreover, DivideMix [26], C2D [59], SSR [11], and CC [58] are used as a baselines for Clothing1M [45] as shown in Tab. 3, with results showing slight improvements with the use of our method. For mini-WebVision [26] (validation on ImageNet [23]), we use DivideMix [26], C2D [59], SSR [11], and CC [58] as baselines (Tab. 4). It is worth noting that our results are better than the SOTA for all settings in Tab. 4.

#### 4.4 Ablation


We show an ablation study with testing accuracy (*top*) and training time (*bottom*) of our approach in Tab. 5 (*left*) on CIFAR100 [22] at 0.3 and 0.5 IDN [44] using DivideMix [26] as baseline. Detailed complexity analysis is shown in the supplementary material, Tab. 10. Initially, the accuracy result of baseline DivideMix [26] is 76.33%, 58.61% for 0.3 and 0.5 IDN, respectively. In the second row, we fix the noise rate  $\epsilon$  at 0.3 and 0.5 for DivideMix’s sample selection, as explained in Sec. 3.3 (without updating  $\epsilon$ ), then the results improved to 78.14% and 64.44%,

**Table 4:** Test accuracy (%) and final estimated noise rate  $\epsilon$  on mini-WebVision [26] and validation on ImageNet [23]. We integrate DivideMix [26], Contrast-to-Divide (C2D) [59], SSR [11] and CC [58] into our model, whilst models suffixed with **-Ours** denotes our proposed approach. ImageNet [23] is only considered for validation.

Dataset	mini-WebVision		ImageNet		Estimated noise rate
	Top-1	Top-5	Top-1	Top-5	
BtR [38]	80.88	92.76	75.96	92.20	
NCE [25]	79.50	<b>93.80</b>	76.30	94.10	
DivideMix [26]	77.32	91.64	75.20	91.64	
<b>DivideMix-Ours</b>	<b>78.51</b>	<b>92.03</b>	<b>76.11</b>	<b>93.24</b>	0.43
	(+1.19)	(+0.39)	(+0.91)	(+1.60)	
C2D [59]	79.42	92.32	78.57	93.04	
<b>C2D-Ours</b>	<b>80.20</b>	<b>92.82</b>	<b>79.16</b>	<b>93.12</b>	0.43
	(+0.78)	(+0.50)	(+0.59)	(+0.08)	
SSR [11]	80.92	92.80	75.76	91.76	
<b>SSR-Ours</b>	<b>81.68</b>	<b>93.80</b>	<b>76.91</b>	<b>93.05</b>	0.43
	(+0.76)	(+1.00)	(+1.15)	(+1.29)	
CC [58]	79.36	93.64	76.08	93.86	-
<b>CC-Ours</b>	<b>80.01</b>	<b>93.79</b>	<b>77.11</b>	<b>94.21</b>	0.44
	(+0.65)	(+0.15)	(+1.03)	(+0.35)	

**Table 5:** (*left*) Ablation results on CIFAR100 at 0.3 and 0.5 IDN. We display accuracy results under different configurations and compare the training time of our approach against DivideMix [26]. The implementation settings and detailed computational analysis are described in Sec. 4.2 and Appendix 4 respectively. (*right above*) Ablation results for initialisation with different  $\epsilon$  values on CIFAR100 [22] at a synthetic noise rate of 0.5 IDN [44] with DivideMix-Ours. (*right below*) Ablation graph showing the learning pattern of the noise rate parameter estimation under different initialisation values, i.e.,  $\epsilon = 0.001$  (blue),  $\epsilon = 0.1$  (orange),  $\epsilon = 0.9$  (red), when trained on DivideMix-Ours case at 0.5 IDN [44] on CIFAR100 [22].

Models	Accuracy (%)	
	0.3	0.5
DivideMix [26]	76.33	58.61
DivideMix with fixed $\epsilon$	78.14	64.44
Our method with pre-trained DivideMix	75.01	52.31
Our method with original DivideMix (no $\epsilon$ )	76.20	56.30
<b>DivideMix-Ours</b>	<b>77.21</b>	<b>64.02</b>
<b>Training Time (GPU-hours)</b>		
DivideMix [26]	7.20	
<b>DivideMix-Ours</b>	8.50	
<b>DivideMix-Ours (Mixed Precision)</b>	6.10	



which is the ideal case (i.e., a perfect noise rate estimation) that motivated our work. In the third case, we use the proposed graphical model with pre-trained DivideMix [26] that shows accuracy of 75.01% and 52.31%. In the next case, the proposed graphical model is trained together with DivideMix [26] without considering the estimated noise rate  $\epsilon$  for sample selection, which results in accuracy of 76.20% and 56.30%. In the last row, we show the training of the proposed model with DivideMix [26], together with the estimation of noise rate  $\epsilon$ , and the

selection of samples based on that, with  $\approx 1\%$  and  $\approx 8\%$  accuracy improvement for 0.3 and 0.5 IDN, which is quite close to our ideal case (second row). We also study the effect that the initialisation of the noise rate parameter  $\epsilon$  plays in our method. As described in Sec. 4.2, we randomly initialise  $\epsilon$  within the range  $(0, 1)$ . To test the robustness of our method to this initialisation, we perform an ablation study by initialising  $\epsilon$  with three different values, namely: 0.001, 0.1, and 0.9 on CIFAR100 [24] at 0.5 IDN [44] using the DivideMix-ours approach. Our results, presented in Tab. 5 (*right above and below*), demonstrate that the initial value of  $\epsilon$  has a limited effect on the model’s performance, highlighting the stability of our method across different initialisation of  $\epsilon$ . In the supplementary material, Tab. 7 (*left*) presents comparative analyses of our model against PartT [44] and SSR [11] on CIFAR100 [22] at 0.5 IDN [44], while Tab. 7 (*right*) showcases similar comparisons with baselines DivideMix [26] and InstanceGM [15] at a higher noise rate of 0.6 IDN [44]. Additionally, Tab. 8 provides further insights into the performance of DivideMix [26] on red mini-ImageNet [17] at lower noise rates of 0.1, 0.2 and 0.3. Tab. 9 illustrates noise rate estimation with standard deviation for our model integrated with DivideMix [26] and InstanceGM [15] on CIFAR100 [22] under IDN [44] settings at 0.2, 0.3, 0.4, and 0.5. Further statistical analysis is detailed in Appendix 5 with confidence interval presented in Tab. 12. Performance analysis with various architectures (Tab. 13) and lambda values (Tab. 14) are mentioned in Appendix 6. Threshold-based comparisons are in Appendix 7 and Tab. 15.

## 5 Discussion and Conclusion

In this work, we proposed an innovative graphical model for IDN noisy-label learning, focusing on the estimation of the label noise rate that is leveraged to introduce a robust sample selection curriculum. Given that SOTA IDN noisy-label learning approaches tend to rely on sample selection methods, our method can be seamlessly integrated to them to improve their performance in many synthetic and real-world benchmarks, including CIFAR100 [22], red mini-ImageNet [17], Clothing1M [45], mini-WebVision [26], and ImageNet [23].

From a societal perspective, our methodology offers a positive impact by mitigating biases inherent in noisy data, thereby promoting more fair and accurate machine learning outcomes. It also opens the door to future research avenues, including exploring advanced noise rate estimation methodologies and investigating noise-robust loss functions such as GCE [56] and ELR [28] with it. Our choice of using the same network architecture for clean and noisy labels enables compatibility with co-teaching techniques [16, 26]. However, investigating different network structures for clean and noisy labels remains a potential avenue for future research. Moreover, we aim to delve into its dynamics, uncovering substantial improvements for certain models while observing more marginal enhancements for others. For more detailed information, please refer to Appendix 8. Our approach also hints at potential breakthroughs in streamlining data annotation challenges.

## References

1. Arpit, D., Jastrzębski, S., Ballas, N., Krueger, D., Bengio, E., Kanwal, M.S., Maharaj, T., Fischer, A., Courville, A., Bengio, Y., Lacoste-Julien, S.: A closer look at memorization in deep networks. In: International Conference on Machine Learning. vol. 70, pp. 233–242. PMLR (2017)
2. Bae, H., Shin, S., Jang, J., Na, B., Song, K., Moon, I.C.: From noisy prediction to true label: Noisy prediction calibration via generative model. In: International Conference on Machine Learning (2022)
3. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. In: Advances in Neural Information Processing Systems. vol. 32 (2019)
4. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: International Conference on Computer Vision. pp. 9650–9660 (2021)
5. Chen, P., Ye, J., Chen, G., Zhao, J., Heng, P.A.: Beyond class-conditional assumption: A primary attempt to combat instance-dependent label noise. In: AAAI Conference on Artificial Intelligence. pp. 11442–11450 (2021)
6. Chen, Z., Song, A., Wang, Y., Huang, X., Kong, Y.: A noise rate estimation method for image classification with label noise. *Journal of Physics: Conference Series* **2433**(1), 012039 (feb 2023). <https://doi.org/10.1088/1742-6596/2433/1/012039>, <https://dx.doi.org/10.1088/1742-6596/2433/1/012039>
7. Cheng, D., Liu, T., Ning, Y., Wang, N., Han, B., Niu, G., Gao, X., Sugiyama, M.: Instance-dependent label-noise learning with manifold-regularized transition matrix estimation. In: Conference on Computer Vision and Pattern Recognition. pp. 16630–16639 (2022)
8. Cheng, H., Zhu, Z., Li, X., Gong, Y., Sun, X., Liu, Y.: Learning with instance-dependent label noise: A sample sieve approach. In: International Conference on Learning Representations (2020)
9. Cordeiro, F.R., Belagiannis, V., Reid, I., Carneiro, G.: PropMix: Hard sample filtering and proportional mixup for learning with noisy labels. In: British Machine Vision Conference (2021)
10. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. *Journal of the royal statistical society: series B (methodological)* **39**(1), 1–22 (1977)
11. Feng, C., Tzimiropoulos, G., Patras, I.: SSR: An efficient and robust framework for learning with unknown label noise. In: British Machine Vision Conference (2022)
12. Feng, C., Ren, Y., Xie, X.: Ot-filter: An optimal transport filter for learning with noisy labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16164–16174 (2023)
13. Frénay, B., Kabán, A.: A comprehensive introduction to label noise. In: European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (2014)
14. Frénay, B., Verleysen, M.: Classification in the presence of label noise: A survey. *IEEE Transactions on Neural Networks and Learning Systems* **25**(5), 845–869 (2013)
15. Garg, A., Nguyen, C., Felix, R., Do, T.T., Carneiro, G.: Instance-dependent noisy label learning via graphical modelling. In: Winter Conference on Applications of Computer Vision. pp. 2288–2298 (2023)

16. Han, B., Yao, Q., Yu, X., Niu, G., Xu, M., Hu, W., Tsang, I., Sugiyama, M.: Co-teaching: Robust training of deep neural networks with extremely noisy labels. In: *Advances in Neural Information Processing Systems*. vol. 31 (2018)
17. Jiang, L., Huang, D., Liu, M., Yang, W.: Beyond synthetic noise: Deep learning on controlled noisy labels. In: *International Conference on Machine Learning*. pp. 4804–4815. PMLR (2020)
18. Jiang, L., Zhou, Z., Leung, T., Li, L.J., Fei-Fei, L.: MentorNet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In: *International Conference on Machine Learning*. pp. 2304–2313. PMLR (2018)
19. Jiang, Z., Zhou, K., Liu, Z., Li, L., Chen, R., Choi, S.H., Hu, X.: An information fusion approach to learning with instance-dependent label noise. In: *International Conference on Learning Representations* (2022), <https://openreview.net/forum?id=ecH2FKaARUp>
20. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. *Advances in Neural Information Processing Systems* **33**, 18661–18673 (2020)
21. Kim, T., Ko, J., Choi, J., Yun, S.Y.: FINE samples for learning with noisy labels. In: *Advances in Neural Information Processing Systems*. vol. 34 (2021)
22. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Tech. rep., University of Toronto (2009)
23. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*. vol. 25 (2012)
24. Kye, S.M., Choi, K., Yi, J., Chang, B.: Learning with noisy labels by efficient transition matrix estimation to combat label miscorrection. In: *European Conference on Computer Vision*. pp. 717–738. Springer (2022)
25. Li, J., Li, G., Liu, F., Yu, Y.: Neighborhood collective estimation for noisy label identification and correction. In: *European Conference on Computer Vision*. pp. 128–145. Springer (2022)
26. Li, J., Socher, R., Hoi, S.C.: DivideMix: Learning with noisy labels as semi-supervised learning. In: *International Conference on Learning Representations* (2020)
27. Li, W., Wang, L., Li, W., Agustsson, E., Gool, L.V.: WebVision Database: Visual learning and understanding from web data. *CoRR* (2017)
28. Liu, S., Niles-Weed, J., Razavian, N., Fernandez-Granda, C.: Early-learning regularization prevents memorization of noisy labels. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 20331–20342 (2020)
29. Liu, Y., Cheng, H., Zhang, K.: Identifiability of label noise transition matrix. In: *International Conference on Machine Learning* (2023)
30. Ma, X., Huang, H., Wang, Y., Romano, S., Erfani, S., Bailey, J.: Normalized loss functions for deep learning with noisy labels. In: *International Conference on Machine Learning*. pp. 6543–6553. PMLR (2020)
31. Malach, E., Shalev-Shwartz, S.: Decoupling “when to update” from “how to update”. In: *Advances in Neural Information Processing Systems*. vol. 30 (2017)
32. McLachlan, G.J., Peel, D.: *Finite mixture models*. Wiley (1996)
33. Neal, R.M., Hinton, G.E.: A view of the EM algorithm that justifies incremental, sparse, and other variants. *Learning in graphical models* pp. 355–368 (1998)
34. Neyshabur, B., Bhojanapalli, S., McAllester, D., Srebro, N.: Exploring generalization in deep learning. In: *Advances in Neural Information Processing Systems*. vol. 30 (2017)



35. Nishi, K., Ding, Y., Rich, A., Hollerer, T.: Augmentation strategies for learning with noisy labels. In: Conference on Computer Vision and Pattern Recognition. pp. 8022–8031 (2021)
36. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: PyTorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems* **32** (2019)
37. Shamshad, F., Khan, S., Zamir, S.W., Khan, M.H., Hayat, M., Khan, F.S., Fu, H.: Transformers in medical imaging: A survey. *Medical Image Analysis* p. 102802 (2023)
38. Smart, B., Carneiro, G.: Bootstrapping the relationship between images and their clean and noisy labels. In: Winter Conference on Applications of Computer Vision. pp. 5344–5354 (2023)
39. Song, H., Kim, M., Park, D., Shin, Y., Lee, J.G.: Learning from noisy labels with deep neural networks: A survey. *IEEE Transactions on Neural Networks and Learning Systems* (2022)
40. Tanaka, D., Ikami, D., Yamasaki, T., Aizawa, K.: Joint optimization framework for learning with noisy labels. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
41. Titterton, D.M.: On the likelihood of mixtures of distributions. *Biometrika* **71**(3), 511–522 (1984). <https://doi.org/10.1093/biomet/71.3.511>
42. Trummer, I.: From BERT to GPT-3 codex: harnessing the potential of very large language models for data management. *VLDB Endowment* **15**(12), 3770–3773 (2022)
43. Tu, H., Menzies, T.: Debtfree: minimizing labeling cost in self-admitted technical debt identification using semi-supervised learning. *Empirical Software Engineering* **27**(4), 80 (2022)
44. Xia, X., Liu, T., Han, B., Wang, N., Gong, M., Liu, H., Niu, G., Tao, D., Sugiyama, M.: Part-dependent label noise: Towards instance-dependent label noise. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 7597–7610 (2020)
45. Xiao, T., Xia, T., Yang, Y., Huang, C., Wang, X.: Learning from massive noisy labeled data for image classification. In: Conference on Computer Vision and Pattern Recognition. pp. 2691–2699 (2015)
46. Xu, Y., Zhu, L., Jiang, L., Yang, Y.: Faster meta update strategy for noise-robust deep learning. In: Conference on Computer Vision and Pattern Recognition. pp. 144–153 (June 2021)
47. Yang, S., Jiang, L., Liu, Z., Loy, C.C.: Unsupervised image-to-image translation with generative prior. In: Conference on Computer Vision and Pattern Recognition. pp. 18332–18341 (2022)
48. Yang, S., Yang, E., Han, B., Liu, Y., Xu, M., Niu, G., Liu, T.: Estimating instance-dependent bayes-label transition matrix using a deep neural network. In: *International Conference on Machine Learning*. pp. 25302–25312. PMLR (2022)
49. Yao, Q., Yang, H., Han, B., Niu, G., Kwok, J.T.Y.: Searching to exploit memorization effect in learning with noisy labels. In: *International Conference on Machine Learning*. pp. 10789–10798. PMLR (2020)
50. Yao, Y., Liu, T., Gong, M., Han, B., Niu, G., Zhang, K.: Instance-dependent label-noise learning under a structural causal model. In: *Advances in Neural Information Processing Systems*. vol. 34 (2021)

51. Yao, Y., Liu, T., Han, B., Gong, M., Deng, J., Niu, G., Sugiyama, M.: Dual T: Reducing estimation error for transition matrix in label-noise learning. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 7260–7271 (2020)
52. Zhang, C., Bengio, S., Hardt, M., Recht, B., Vinyals, O.: Understanding deep learning requires rethinking generalization. In: *International Conference on Learning Representations* (2017)
53. Zhang, C., Bengio, S., Hardt, M., Recht, B., Vinyals, O.: Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM* **64**(3), 107–115 (2021)
54. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: *International Conference on Learning Representations* (2017)
55. Zhang, Y., Niu, G., Sugiyama, M.: Learning noise transition matrix from only noisy labels via total variation regularization. In: *International Conference on Machine Learning*. pp. 12501–12512. PMLR (2021)
56. Zhang, Z., Sabuncu, M.: Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in Neural Information Processing Systems* **31** (2018)
57. Zhang, Z., Pfister, T.: Learning fast sample re-weighting without reward data. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 725–734 (2021)
58. Zhao, G., Li, G., Qin, Y., Liu, F., Yu, Y.: Centrality and consistency: two-stage clean samples identification for learning with instance-dependent noisy labels. In: *European Conference on Computer Vision*. pp. 21–37. Springer (2022)
59. Zheltonozhskii, E., Baskin, C., Mendelson, A., Bronstein, A.M., Litany, O.: Contrast to divide: Self-supervised pre-training for learning with noisy labels. In: *Winter Conference on Applications of Computer Vision*. pp. 1657–1667 (2022)
60. Zhu, Z., Liu, T., Liu, Y.: A second-order approach to learning with instance-dependent label noise. In: *Conference on Computer Vision and Pattern Recognition*. pp. 10113–10123 (June 2021)