

Instance-Dependent Noisy Label Learning via Graphical Modelling

Arpit Garg^{*1}, Cuong Nguyen¹, Rafael Felix¹, Thanh-Toan Do², and Gustavo Carneiro^{1,3}

¹Australian Institute for Machine Learning, University of Adelaide, Australia

²Department of Data Science and AI, Monash University, Australia

³Centre for Vision, Speech and Signal Processing, University of Surrey, United Kingdom,

Abstract

Noisy labels are unavoidable yet troublesome in the ecosystem of deep learning because models can easily overfit them. There are many types of label noise, such as symmetric, asymmetric and instance-dependent noise (IDN), with IDN being the only type that depends on image information. Such dependence on image information makes IDN a critical type of label noise to study, given that labelling mistakes are caused in large part by insufficient or ambiguous information about the visual classes present in images. Aiming to provide an effective technique to address IDN, we present a new graphical modelling approach called InstanceGM, that combines discriminative and generative models. The main contributions of InstanceGM are: i) the use of the continuous Bernoulli distribution to train the generative model, offering significant training advantages, and ii) the exploration of a state-of-the-art noisy-label discriminative classifier to generate clean labels from instance-dependent noisy-label samples. InstanceGM is competitive with current noisy-label learning approaches, particularly in IDN benchmarks using synthetic and real-world datasets, where our method shows better accuracy than the competitors in most experiments¹.

1. Introduction

The latest developments in deep neural networks (DNNs) have shown outstanding results in a variety of applications ranging from computer vision [31] to natural language processing [48] and medical image analysis [47]. Such success is strongly reliant on high-capacity models, which in turn, require a massive amount of correctly-annotated data for training [34, 67]. Anno-

tating a large amount of data is, however, arduous, costly and time-consuming, and therefore is often done via crowdsourcing [56] that generally produces low-quality annotations. Although that brings down the cost and scales up the process, the trade-off is the mislabelling of the data, resulting in a deterioration of deep models' performance [3, 35] due to the *memorisation effect* [2, 35, 44, 70]. This has, therefore, motivated the research of novel learning algorithms to tackle the label noise problem where data might have been mislabelled.

Early work in label noise [17] was carried out under the assumption that label noise was instance-independent (IIN), i.e., mislabelling occurred regardless of the information about the visual classes present in images. In IIN, we generally have a transition matrix that contains a predefined probability of flipping between pairs of labels (e.g., any image showing a *cat* has a high priori probability of being mislabelled as a *dog* and low a priori probability of being mislabelled as a *car*). This type of noise can also be divided into two sub-types: *symmetric*, where a true label is flipped to another label with equal probability across all classes, and *asymmetric*, where a true label is more likely to be mislabeled into one of some particular classes [17]. Nevertheless, the IIN assumption is impractical for many real-world datasets because we can intuitively argue that mislabellings mostly occur because of insufficient or ambiguous information about the visual classes present in images. As a result, recent studies have gradually shifted their focus toward the more realistic scenario of instance-dependent noise (IDN), where label noise depends on both the true class label and the image information [62].

Many methods have been introduced to handle not only IIN, but also IDN problems. Those include, but are not limited to, *sample selection* [12, 27, 33, 61, 72] that detects clean and noisy labels and applies semi-supervised learning methods on the processed data, *robust losses* [1, 38, 46] that can work well with either clean or noisy labels, and *probabilistic approaches* [66] that model the data genera-

^{*}arpit.garg@adelaide.edu.au

¹Code is available at <https://github.com/arpit2412/InstanceGM>.

Supported by Australian Research Council through grants DP180103232 and FT190100525.

tion process, including how a noisy label is created. Despite some successes, most methods are often demonstrated in IIN settings with simulated symmetric and asymmetric noise. However, their performance is degraded when evaluated on IDN problems, which include real-world and synthetic datasets. Although there are a few studies focusing on the IDN setting [10, 26, 62, 66, 74], their relatively inaccurate classification results suggest that the algorithms can be improved further.

In this paper, we propose a new method to tackle the IDN problem, called InstanceGM. Our method is designed based on a graphical model that considers the clean label Y as a latent variable and introduces another latent variable Z representing the image feature to model the generation of a label noise \hat{Y} and an image X . InstanceGM integrates generative and discriminative models, where the generative model is based on a variational auto-encoder (VAE) [28], except that we replace the conventional mean squared error (MSE) when modelling the likelihood of reconstructed images by a *continuous Bernoulli distribution* [40] that facilitates the training process since it avoids tuning additional hyper-parameters. For the discriminative model, to mitigate the problem of only using clean label data during the training process, which is a common issue present in the similar graphical model methods [66], we rely on DivideMix [33] that uses both clean and noisy-label data for training by exploring semi-supervised learning via MixMatch [5]. DivideMix is shown to be a reasonably effective discriminative classifier for our InstanceGM. In summary, the main contributions of the proposed method are:

- InstanceGM follows a graphical modelling approach to generate both the image X and its noisy label \hat{Y} with the true label Y and image feature Z as latent variables. The modelling is associated with the continuous Bernoulli distribution to model the generation of instance X to facilitate the training, avoiding tuning of additional hyper-parameters (see Remark 3).
- For the discriminative classifier of InstanceGM, we replace the commonly used co-teaching, which is a dual model that relies only on training samples classified as clean, with DivideMix [33] that uses all training samples classified as clean and noisy.
- InstanceGM shows state-of-the-art results on a variety of IDN benchmarks, including simulated and real-world datasets, such as CIFAR10 and CIFAR100 [30], Red Mini-ImageNet from Controlled Noisy Web Labels (CNWL) [65], ANIMAL-10N [53] and CLOTHING-1M [64].

2. Related work

As DNNs have been shown to easily fit randomly labelled training data [68], they can also overfit a noisy-label dataset, which eventually results in poor generalisation to a clean-label testing data [2, 35, 44, 70]. Several studies have, therefore, been conducted to investigate supervised learning under the label noise setting, including robust loss function [41, 58], sample selection [53, 55, 59], robust regularisation [14, 23, 43, 60] and robust architecture [11, 16, 29, 64]. Below, we review methods dealing with noisy labels, especially IDN, without the reliance on clean validation sets [22, 50, 57].

Let us start with methods designed to handle “any” type of label noise, including IDN and IIN. An important technique for both of these label noise types is sample selection [53, 55, 59], which aims to select clean-label samples automatically for training. Although it is well-motivated and often effective, it suffers from the cumulative error caused by mistakes in the selection process, mainly when there are numerous unclear classes in the training data. Consequently, sample selection methods often rely on multiple clean-label sample classifiers to increase their robustness against such cumulative error [33]. In addition, semi-supervised learning (SSL) [4, 12, 27, 33, 53, 72] have also been integrated with sample selection and multiple clean-label classifiers to enable the training from clean and noisy-label samples. In particular, SSL methods use clean and noisy samples by treating them as labelled and unlabelled data, respectively, with a MixMatch approach [5]. These methods above have been designed to handle “any” type of label noise, so they are usually assessed in synthetic IIN benchmarks and real-world IDN benchmarks.

Given that real-world datasets do not, in general, contain IIN, more recently proposed methods aim to address IDN benchmarks [6, 10, 37, 62, 66, 73]. In these benchmarks, the task of differentiating between hard clean-labelled samples and noisy-label samples pose a major challenge. Such issue is noted by Song et al. [54], who state that the model performance in IDN can degrade significantly compared to other types of noises.

One direct way of addressing IDN problems relies on a graphical model approach that has random variables representing the observed noisy label, the image, and the latent clean label. This model also has a generative process to produce an image given the (clean and noisy) label information [32]. Another approach examines a graphical model using a discriminative process [49], where the model attempts to explain the posterior probability of the observed noisy label by averaging the posterior probabilities of the clean class label. Yao et al. [66] developed a new causal model to address IDN that also uses the same variables as the methods above plus a latent image feature variable, which relies on generative models to produce the image from the clean




Figure 1. The proposed graphical model of the generation process that produces the observable (shaded nodes) data X and noisy label \hat{Y} from hidden (non-shaded nodes) data representation Z and clean label Y .

label and image feature, and to produce the noisy label from the image feature and clean label. That approach [66], however, did not produce competitive results compared with state of the art. We argue that the model’s poor performance is mostly due to the co-teaching [17] that is trained with a small set of samples classified as clean, which can inadvertently contain noisy-label samples – this is an issue that can cause a cumulative error, particularly in IDN problems.

Our work is motivated by the graphical model approaches mentioned above, that aim to address IDN problems. The main difference in our approach is the use of a more effective clean sample identifier that replaces co-teaching [66] by DivideMix [33], which considers the whole training set, instead of only the samples classified as clean. Moreover, we propose a more effective training of the image generative model based on the continuous Bernoulli distribution [40].

3. Methodology

3.1. Problem definition

We denote X as an observed random variable representing an image, Y as a latent random variable corresponding to the clean label of X , Z as a latent random variable denoting an image feature representation for of X , and \hat{Y} as the observed random variable for the noisy label. The training set is represented by $\mathcal{D} = \{(\mathbf{x}_i, \hat{\mathbf{y}}_i)\}_{i=1}^{|\mathcal{D}|}$, where the image is represented by $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^{H \times W \times 3}$ (with 3 color channels and size $H \times W$ pixels) and the noisy label $\hat{\mathbf{y}} \in \mathcal{Y} \in \{0, 1\}^{|\mathcal{Y}|}$ denoted by a one-hot vector. In the conventional supervised learning, \mathcal{D} is used to train a model $f_\theta : \mathcal{X} \rightarrow \Delta^{|\mathcal{Y}|-1}$ (where $\Delta^{|\mathcal{Y}|-1}$ represents the probability simplex), parameterised by $\theta \in \Theta$, that can predict the labels of testing images. The aim is to exploit the noisy data (X, \hat{Y}) from a training set to infer a model f_θ that can accurately predict the clean labels Y of data in a testing set.

3.2. Probabilistic noisy label modelling

We follow a similar approach presented in [66] to model the process that generates samples with noisy labels via the graphical model shown in Fig. 1, where the clean label Y and image feature representation Z are latent variables. Under this modelling assumption, a noisy-label sample $(\mathbf{x}, \hat{\mathbf{y}})$ can be generated as follows:

1. sample a clean label from its prior: $\mathbf{y} \sim p(Y)$,
2. sample a representation from its prior: $\mathbf{z} \sim p(Z)$,
3. sample an input data from its continuous Bernoulli distribution: $\hat{\mathbf{x}} \sim \mathcal{CB}(X; \lambda(\mathbf{z}, \mathbf{y}))$,
4. sample the corresponding noisy label from its categorical distribution: $\hat{\mathbf{y}} \sim \text{Cat}(\hat{Y}; \gamma(\hat{\mathbf{x}}, \mathbf{y}))$

Remark 1 *Conventionally, the process of generating data \mathbf{x} in step 3 above is often modelled as a Bernoulli distribution or multivariate normal distribution, corresponding to the binary cross-entropy (BCE) or MSE reconstruction losses, respectively. Such modelling, however, leads to a pervasive error [40] since the image pixels are in $[0, 1]$ instead of $\{0, 1\}$ (Bernoulli distribution)² or $(-\infty, +\infty)$ (multivariate normal distribution). We therefore adopt the continuous Bernoulli distribution [40] which has a support in $[0, 1]$ to correctly model this image generation process.*

Note that the parameters of the continuous Bernoulli and categorical distributions are conditioned on Z , X and Y , and modelled as the outputs of two DNNs:

$$\lambda = f_{\theta_x}(\mathbf{z}, \mathbf{y}) \quad \text{and} \quad \gamma = f_{\theta_y}(\hat{\mathbf{x}}, \mathbf{y}), \quad (1)$$

where f denotes the neural network, and θ_x, θ_y represent the network parameters. Following the convention in machine learning, we call $f_{\theta_x}(\cdot)$ the *decoder* and $f_{\theta_y}(\cdot)$ the *noisy label classifier*.

To solve the label noise problem that has data generated from the process above, we need to infer the posterior $p(Z, Y|X, \hat{Y})$. However, due to the complexity of the graphical model in Fig. 1, exact inference for the posterior $p(Z, Y|X, \hat{Y})$ is intractable, and therefore, the estimation must rely on an approximation. Motivated by [66], we employ variational inference to approximate the true posterior $p(Z, Y|X, \hat{Y})$ by a variational ‘‘posterior’’ $q(Z, Y|X, \hat{Y})$. Such posterior can be obtained by minimising the following Kullback-Leibler (KL) divergence:

$$\min_q \text{KL} \left[q(Z, Y|X, \hat{Y}) \parallel p(Z, Y|X, \hat{Y}) \right], \quad (2)$$

where the variational posterior $q(Z, Y|X, \hat{Y})$ can be factorised following the product rule of probability. We assume that the posterior of the clean label Y is independent from the noisy label \hat{Y} , given the instance X : $q(Y|X, \hat{Y}) = q(Y|X)$. In addition, the variational posterior of feature representation is independent from the noisy label given the clean label and input data: $q(Z|X, \hat{Y}, Y) = q(Z|X, Y)$. The variational posterior of interest can, therefore, be written as:

$$\begin{aligned} q(Z, Y|X, \hat{Y}) &= q(Z|X, \hat{Y}, Y) q(Y|X, \hat{Y}) \\ &= q(Z|X, Y) q(Y|X). \end{aligned} \quad (3)$$

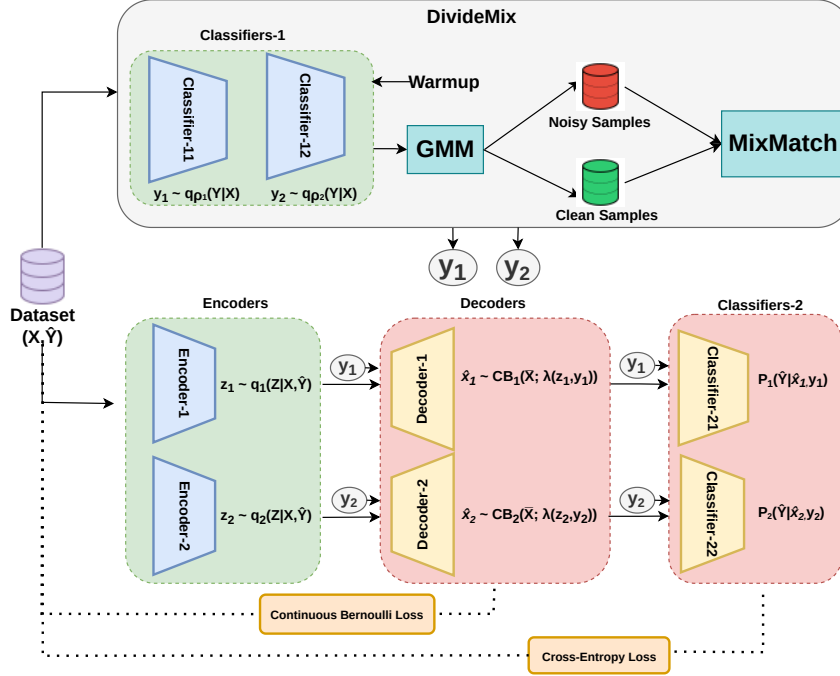


Figure 2. The proposed InstanceGM trains the *Classifiers* to output clean labels for instance-dependent noisy-label samples. We first warmup our two classifiers (*Classifier-11,12*) using the classification loss, and then with classification loss we train the GMM to separate clean and noisy samples with the semi-supervised model MixMatch [5] from the DivideMix [33] stage. Additionally, another set of encoders (*Encoder-1,2*) are used to generate the latent image features as depicted in the graphical model from Fig. 1. Furthermore, for image reconstruction, the decoders (*Decoder-1,2*) are used by utilizing the continuous Bernoulli loss, and another set of classifiers (*Classifier-21,22*) helps to identify the original noisy labels using the standard cross-entropy loss.

The objective function in (2) can then be expanded as:

$$\begin{aligned}
 L^{(vi)} = & \mathbb{E}_{q(Z|X,Y)q(Y|X)} [-\ln p(X|Z, Y)] \\
 & + \mathbb{E}_{q(Y|X)} [-\ln p(\hat{Y}|X, Y)] \\
 & + \text{KL}[q(Y|X)||p(Y)] \\
 & + \mathbb{E}_{q(Y|X)} [\text{KL}[q(Z|X, Y)||p(Z)]] . \quad (4)
 \end{aligned}$$

Remark 2 The objective function $L^{(vi)}$ in (4) shares similarity with the loss in variational auto-encoder [28]. In particular, the first two terms in (4) are analogous to the reconstruction loss, while the remaining terms are analogous to the KL loss that regularises the deviation between the posterior q and its prior.

To optimise the objective in (4), both the posteriors $q(Z|X, Y)$ and $q(Y|X)$ and priors $p(Z)$ and $p(Y)$ must be specified. We assume $q(Z|X, Y)$ to be a multivariate normal distribution with a diagonal covariance matrix and $q(Y|X)$ to be a categorical distribution:

$$\begin{aligned}
 q(Z|X = \mathbf{x}, \hat{Y} = \hat{\mathbf{y}}) &= \mathcal{N}(Z; \boldsymbol{\mu}(\mathbf{x}, \hat{\mathbf{y}}), \text{diag}(\boldsymbol{\sigma}^2(\mathbf{x}, \hat{\mathbf{y}})) \\
 q(Y|X = \mathbf{x}) &= \text{Cat}(Y; \boldsymbol{\rho}(\mathbf{x})), \quad (5)
 \end{aligned}$$

²Except for black and white images.

where the parameters of these distributions are modelled as the outputs of two DNNs. Hereafter, we call the network that models $q(Y|X)$ the *clean label classifier*, and the model $q(Z|X, \hat{Y})$, the *encoder*.

For the priors, we follow the convention in generative models, especially VAE, to assume $p(Z)$ as a standard normal distribution, while $p(Y)$ is a uniform distribution.

Given such assumptions, we can minimise the loss function $L^{(vi)}$ in (4) w.r.t. the parameters of the two classifiers, the encoder and decoder in (5) and (1). The obtained clean label classifier that models $q(Y|X)$ will then be used as the final classifier to evaluate data in the testing set.

Remark 3 Optimising the objective function in (4) often requires the definition of hyper-parameters to weight the KL divergences [15]. However, such weighting mechanism depends on the estimation of the KL divergences weights that is usually achieved with a grid-search using a validation set, making solutions dependent on the dataset. The reason for such weighting mechanism lies at the log-likelihoods used as reconstruction losses. For example, $-\ln p(X|Z, Y)$ is simply replaced by the corresponding loss functions, such as MSE, without taking the normalisation constants of those likelihood functions into account, resulting in an incorrect balance between reconstruction loss and regularisation. In

Algorithm 1 Graphical model approach for learning with label noise

```
1: procedure INSTANCEGM( $\mathcal{D}, T, \tau$ )
2:    $\triangleright \mathcal{D} = \{(\mathbf{x}_i, \hat{\mathbf{y}}_i)\}_{i=1}^{|\mathcal{D}|}$ : noisy dataset ◁
3:    $\triangleright T$ : total number of epochs ◁
4:    $\triangleright \tau$ : threshold to decide clean or noisy samples used in DivideMix ◁
5:    $q_1(Y|X), q_2(Y|X) \leftarrow \text{WARMUP}(\mathcal{D})$  ◁  $\triangleright$  Warm-up training of 2 clean-label classifiers on noisy dataset
6:   for  $e = 1 : T$  do
7:      $\mathcal{L}_1, \mathcal{U}_1, \mathcal{L}_2, \mathcal{U}_2 \leftarrow \text{CO-DIVIDE}(\mathcal{D}, q_1(Y|X), q_2(Y|X), \tau)$ 
8:      $\triangleright$  Apply Gaussian mixture model on loss values and filter out clean and noisy with a threshold on the likelihood ◁
9:      $\triangleright \mathcal{L}_{1:2}$  are labelled sets (mostly clean) ◁
10:     $\triangleright \mathcal{U}_{1:2}$  are unlabelled sets (mostly noisy) ◁
11:     $\mathcal{L}_{(1)}^{(\text{dm})} \leftarrow \text{DIVIDEMIX LOSS}(\mathcal{L}_1, \mathcal{U}_1, q_1(Y|X), q_2(Y|X))$  ◁  $\triangleright$  Calculate training loss in DivideMix
12:     $\mathcal{L}_{(2)}^{(\text{dm})} \leftarrow \text{DIVIDEMIX LOSS}(\mathcal{L}_2, \mathcal{U}_2, q_2(Y|X), q_1(Y|X))$ 
13:    for  $k = 1:2$  do ◁  $\triangleright$  Calculate loss on each one of the 2 models
14:      for each  $(\mathbf{x}_i, \hat{\mathbf{y}}_i) \in \mathcal{L}_k$  do
15:        Compute each instance loss:  $\mathcal{L}_i^{(\text{vi})} \leftarrow \text{VARIATIONAL-FREE ENERGY}(\mathbf{x}_i, \hat{\mathbf{y}}_i, q_k, p_k)$ 
16:         $\triangleright q_k$  is the variational posterior ◁
17:         $\triangleright p_k$  denotes prior and data generation ◁
18:        Compute average loss:  $\mathcal{L}_{(k)}^{(\text{vi})} = 1/|\mathcal{L}_k| \sum_{i=1}^{|\mathcal{L}_k|} \mathcal{L}_i^{(\text{vi})}$ 
19:        Update model parameters by minimizing  $\mathcal{L}_{(k)} = \mathcal{L}_{(k)}^{(\text{vi})} + \mathcal{L}_{(k)}^{(\text{dm})}$  ◁  $\triangleright$  Eq. (6)
20:    return  $q_1(Y|X)$  ◁  $\triangleright$  clean-label classifier

21: function VARIATIONAL-FREE ENERGY( $\mathbf{x}, \hat{\mathbf{y}}, q, p$ ) ◁  $\triangleright$  Calculate loss in Eq. (4)
22:   Sample  $\mathbf{y} \sim q(Y|X = \mathbf{x})$  ◁  $\triangleright$  Sample a clean label from its variational posterior
23:   Sample  $\mathbf{z} \sim q(Z|X = \mathbf{x}, \hat{Y} = \hat{\mathbf{y}})$  ◁  $\triangleright$  Sample a feature representation from its variational posterior
24:   Compute the 1st term in Eq. (4):  $-\ln p(X = \mathbf{x}|Z = \mathbf{z}, Y = \mathbf{y})$  ◁  $\triangleright$  image reconstruction loss
25:   Compute the 2nd term in Eq. (4):  $-\ln p(\hat{Y} = \hat{\mathbf{y}}|X = \mathbf{x}, Y = \mathbf{y})$  ◁  $\triangleright$  noisy-label cross-entropy loss
26:   Compute the remaining terms in Eq. (4)
27:   Compute  $\mathcal{L}^{(\text{vi})}$  as the sum of the above terms as specified in Eq. (4)
28:   return  $\mathcal{L}^{(\text{vi})}$ 
```

this paper, we propose the use of the correct form of the log-likelihood, namely the continuous Bernoulli distribution for $p(X|Z, Y)$ and categorical distribution for $p(\hat{Y}|X, Y)$, with their normalisation constants. Hence, we no longer need the weighting of the KL divergences, making our proposed method simpler to train.³

3.3. Practical implementation

In practice, the small loss hypothesis is often used to effectively identify the clean samples in a training set [17, 33]. However, naively implementing such hypothesis using a single model might accumulate error due to sample selection bias. One way to avoid such scenario is to train two models simultaneously where each model is updated using only the clean samples selected by the other model. In this paper, we integrate a similar approach into our modelling presented in Section 3.2 to solve the label noise problem.

In particular, we propose to train two models in parallel, resulting in four classifiers (two for the clean label classifier $q(Y|X)$ and the other two for noisy labels $p(\hat{Y}|X, Y)$), two encoders $q(Z|X, Y)$ and two decoders $p(X|Z, Y)$.

In CausalNL [66], co-teaching is used as a way to integrate the small loss hypothesis to regularise the clean label classifiers. Co-teaching might, however, limit the capability of the modelling since it only uses samples classified as clean and ignores the other samples classified as noisy. In addition, co-teaching is initially designed for IIN problems, while our focus is on IDN problems. Hence, we propose to integrate DivideMix [33], a method based on the small loss hypothesis as shown in Fig. 2. This method starts with a warmup stage, and utilizes all training samples after classifying them as clean and noisy (co-divide) using a two-component Gaussian mixture model (GMM). The training samples are used by MixMatch [5] – a semi-supervised classification technique that considers clean samples as labelled and noisy samples as unlabeled. DivideMix shows a rea-

³More detailed information mentioned in Appendix B

sonable efficacy for IDN problems, as shown in Table 1.

Remark 4 *Other instance-dependent methods similar to DivideMix [33], such as Contrast-to-Divide [72], ELR+ [35], can also be integrated into our proposed framework. The reason that DivideMix is used is due to its remarkable performance, especially on the IDN setting, and its publicly available implementation.*

In general, the loss function for training the proposed model consists of two losses: one is the loss $L^{(vi)}$ from the graphical modelling in (4), and the other is the loss to train DivideMix [33, Eq. (12)], denoted as $L^{(dm)}$. The whole loss is represented as:

$$L = L^{(vi)} + L^{(dm)} \quad (6)$$

and the training procedure is summarised in Algorithm 1 and depicted in Fig. 2.

4. Experiments

In this section, we show the results of extensive experiments on two standard benchmark datasets with IDN, CIFAR10 [30] and CIFAR100 [30] at various noise rates⁴, and three real-world datasets, ANIMAL-10N [53], Red Mini-Imagenet from CNWL [24] and CLOTHING-1M [64]. In Section 4.1, we explain all datasets mentioned above. In Section 4.2, we discuss all models and their parameters. We compare our approach with state-of-the-art models in IDN benchmarks and real-world datasets in Section 4.3.

4.1. Datasets

In both CIFAR10 and CIFAR100, there are 50k training images and 10k testing images with each images of size $32 \times 32 \times 3$ pixels, where CIFAR10 consists of 10 classes, CIFAR100 has 100 classes and both datasets are class-balanced. As CIFAR10 and CIFAR100 datasets do not include label noise by default, we added IDN with noise rates in $\{0.2, 0.3, 0.4, 0.45, 0.5\}$ following the setup proposed by Xia et al. [62].

Red Mini-Imagenet from CNWL [24] is a real-world dataset where images and their corresponding labels are crawled from internet at various controllable label noise rates. This dataset is proposed to study real-world noise in controlled settings. In this work, we focus on Red Mini-ImageNet since it shows a realistic type of label noise. Red Mini-ImageNet has 100 classes, with each class containing 600 images sampled from the ImageNet dataset [51]. The images are resized to 32×32 pixels from the original size of 84×84 to have a fair comparison with [12, 65]. The noise rates vary from 0% to 80%, but we use the rates 20%, 40%, 60% and 80% to be consistent with the literature [12, 65, 66].

⁴Performance degradation at high IDN is presented in Appendix C.

ANIMAL-10N is another real-world dataset proposed by Song et al. [53], which contains 10 animals with 5 pairs having similar appearances (e.g., wolf and coyote, hamster and guinea pig, etc.). The estimated rate of label noise is 8%. There are 50k training images 10k test images. No data augmentation is used, hence the setup is identical to the one proposed in [53].

CLOTHING-1M [64] is a real-world dataset that comprises 1 million training apparel images taken from 14 categories of online shopping websites. The labels in this dataset are generated from surrounding texts, with an estimated noise of 38.5%. Due to the inconsistency in image sizes, we follow the standard setup in the literature [12, 18, 33] and resize the images to 256×256 pixels. This dataset additionally includes 50k, 14k, and 10k manually validated clean training, validation, and testing data, respectively. During training, the clean training and validation sets are not used and only the clean testing set is used for assessment.

4.2. Implementation

All the methods are implemented in PyTorch [45], and run on NVIDIA RTX A6000. For the baseline model DivideMix, all the default hyperparameters are considered as mentioned in original paper by Li et al. [33]. All hyperparameter values mentioned below are from CausalNL [66] and DivideMix [33] unless otherwise specified. The size of the latent representation Z is fixed at 25 for CIFAR10, CIFAR100 and Red Mini-Imagenet, 64 for ANIMAL-10N, and 100 for CLOTHING-1M. For CIFAR10, CIFAR100 and Red Mini-Imagenet, we used non-pretrained PreAct-ResNet-18 (PRN18) [21] as an encoder. VGG-19 is used as an encoder for ANIMAL-10N, following SELFIE [53] and PLC [71]. For CLOTHING-1M, we used ImageNet-pretrained ResNet-50. *Clean data is not used for training.*

The training of the model used stochastic gradient descent (SGD) for DivideMix stage with momentum of 0.9, batch size of 64 and an L2 regularisation whose parameter is 5×10^{-4} . Additionally, Adam is used to train the VAE part of the model. The training runs for 300 epochs for CIFAR10 (~ 30 hrs), CIFAR100, Red Mini-Imagenet and ANIMAL-10N. The learning rate is 0.02 which is reduced to 0.002 at half of the number of training epochs. The WarmUp stage lasts for 10 epochs for CIFAR10, 30 for CIFAR100, ANIMAL-10N and Red Mini-Imagenet. For CLOTHING-1M, the WarmUp stage lasts 1 epoch with batch size of 32, and training runs for 80 epochs and learning rate of 0.01 decayed by a factor of 10 after the 40th epoch.

For CIFAR10, CIFAR100 [30], Red Mini-Imagenet [24] and ANIMAL-10N [53], the encoder has a similar architecture as CausalNL [66], with 4 hidden convolutional layers and feature maps containing 32, 64, 128 and 256

⁵Implementation details are present in Appendix A

Table 1. Test accuracy (%) of different methods on CIFAR10 and CIFAR100 [30] under various IDN noise rates. Most results are extracted from [66], while results with * are reported in their respective papers. Results taken from kMEIDTM [10] are presented with †.

Model	IDN - CIFAR10					IDN - CIFAR100				
	0.20	0.30	0.40	0.45	0.50	0.20	0.30	0.40	0.45	0.50
CE [66]	75.81	69.15	62.45	51.72	39.42	30.42	24.15	21.45	15.23	14.42
Mixup [69]	73.17	70.02	61.56	56.45	48.95	32.92	29.76	25.92	23.13	21.31
Forward [46]	74.64	69.75	60.21	48.81	46.27	36.38	33.17	26.75	21.93	19.27
T-Revision [63]	76.15	70.36	64.09	52.42	49.02	37.24	36.54	27.23	25.53	22.54
Reweight [36]	76.23	70.12	62.58	51.54	45.46	36.73	31.91	28.39	24.12	20.23
PTD-R-V [62]*	76.58	72.77	59.50	–	56.32	65.33†	64.56†	59.73†	–	56.80†
Decoupling [42]	78.71	75.17	61.73	58.61	50.43	36.53	30.93	27.85	23.81	19.59
Co-teaching [17]	80.96	78.56	73.41	71.60	45.92	37.96	33.43	28.04	25.60	23.97
MentorNet [25]	81.03	77.22	71.83	66.18	47.89	38.91	34.23	31.89	27.53	24.15
CausalNL [66]	81.79	80.75	77.98	79.53	78.63	41.47	40.98	34.02	33.34	32.13
HOC [74]*	90.03	–	85.49	–	–	68.82	–	62.29	–	–
CAL [73]*	92.01	–	84.96	–	–	69.11	–	63.17	–	–
kMEIDTM [10]*	92.26	90.73	85.94	–	73.77	69.16	66.76	63.46	–	59.18
DivideMix [33]	94.80	94.60	94.53	94.08	93.04	77.07	76.33	70.80	57.78	58.61
InstanceGM	96.68	96.52	96.36	96.15	95.90	79.69	79.21	78.47	77.49	77.19

Table 2. Test accuracy (%) for Red Mini-Imagenet (CNWL) [24]. Other model results are as presented in FaMUS [65] and PropMix [12]. We presented our proposed results with our proposed InstanceGM and with inclusion of self-supervision [8] in proposed algorithm (InstanceGM-SS).

Method	Noise rate			
	0.2	0.4	0.6	0.8
CE [65]	47.36	42.70	37.30	29.76
MixUp [69]	49.10	46.40	40.58	33.58
DivideMix [33]	50.96	46.72	43.14	34.50
MentorMix [24]	51.02	47.14	43.80	33.46
FaMUS [65]	51.42	48.06	45.10	35.50
InstanceGM	58.38	52.24	47.96	39.62
With self-supervised learning				
PropMix [12]	61.24	56.22	52.84	43.42
InstanceGM-SS⁵	60.89	56.37	53.21	44.03

features. In the decoding stage, we use 4 hidden layer transposed-convolutional network and the feature maps have 256, 128, 64 and 32 features. In Red Mini-Imagenet, we use a similar architecture as CIFAR100 with and without self-supervision [8]. For CLOTHING-1M [64], we use encoder networks with 5 convolutional layers, and the feature maps contain 32, 64, 128, 256 and 512 features. The decoder networks have 5 transposed-convolutional layers and the feature maps have 512, 256, 128, 64 and 32 features.

4.3. Comparison with Baselines and Measurements

In this section, we compare our proposed InstanceGM on baseline IDN benchmark datasets in Section 4.3.1, and

Table 3. Test accuracy (%) of different methods evaluated on ANIMAL-10N [53] where only noisy data are used to train models. Other models’ results are as presented in Nested-CE [9], and results with * are reported in their respective papers

Method	Test Accuracy (%)
CE [71]	79.4
Nested-Dropout [9]	81.3
CE+Dropout [9]	81.3
SELFIE [53]*	81.8
PLC [71]*	83.4
Nested-CE [9]	84.1
InstanceGM	84.6

we also validate our proposed model on various real-world noisy datasets in Section 4.3.2.

4.3.1 Instance-Dependent Noise Benchmark Datasets

The comparison between our InstanceGM and recently proposed approaches on CIFAR10 and CIFAR100 IDN benchmarks is shown in Table 1. Note that the proposed approach achieves considerable improvements in both datasets at various IDN noise rates ranging from 20% to 50%. Given that CausalNL represents the main reference for our method, it is important to compare the performance of the two approaches. For CIFAR10, our method is roughly 15% better in all noise rates, and for CIFAR100, our method is between 38% and 45% better. Compared to the current state-of-the-art methods in this benchmark (kMEIDTM [10] and DivideMix [33]), our method is around 2% better in CIFAR10 and between 2% to almost 20% better in CIFAR100.

Table 4. Test accuracy (%) for competing methods on CLOTHING-1M [64]. The accuracy of the baseline models (CausalNL and DivideMix) are in italics. Results of other models are from their respective papers. In the experiments only noisy labels are use for training. Top results with 1% accuracy are highlighted in bold.

Method	Test Accuracy (%)
CausalNL [66]	72.24
IF-F-V [26]	72.29
<i>DivideMix [33]</i>	<i>74.76</i>
Nested-CoTeaching [9]	74.90
InstanceGM	74.40

4.3.2 Real-world Noisy Datasets

In Tables 2 to 4, we present the results on ANIMAL-10N, Red Mini-Imagenet and CLOTHING-1M, respectively. In general the results show that InstanceGM outperforms or is competitive with the present state-of-the-art models for large-scale web-crawled datasets and small-scale human-annotated noisy datasets. Table 3 reports the classification accuracy on ANIMAL-10N. We can observe that InstanceGM achieves slightly better performance than all other baselines. For the other real-world datasets Red Mini-Imagenet and CLOTHING-1M, InstanceGM is competitive, as shown in Tables 2 and 4, demonstrating its ability to handle real-world IDN problems. In particular, Table 2 shows the results on Red Mini-Imagenet using two set-ups: 1) without pre-training (top part of the table), and 2) with self-supervised (SS) pre-training (bottom part of the table). The SS pre-training is based on DINO [8] with the unlabelled Red Mini-Imagenet dataset, allowing a fair comparison with PropMix [12], which uses a similar SS pre-training. Without SS pre-training, our InstanceGM is substantially superior to recently proposed approaches. With SS pre-training, results show that InstanceGM can improve its performance, allowing us to achieve state-of-the-art results on Red Mini-Imagenet.

5. Ablation Study

We show the ablation study of our proposed method on CIFAR10 [30], under IDN noise rate of 0.5 and ANIMAL-10N [53]. On Table 5, the performance of CausalNL [66] is relatively low, which can be explained by the small number of clean samples used by co-teaching [17], and the use of MSE for image reconstruction loss⁶. We argue that replacing co-teaching [17] by DivideMix [33] will improve classification accuracy because it allows the use of the whole training set. To demonstrate that, we take CausalNL [66] and replace its co-teaching by DivideMix, but keep the MSE reconstruction loss – this model is named CausalNL + DivideMix (w/o continuous Bernoulli). Note that this allows a

⁶line 80 in <https://github.com/a5507203/IDLN/blob/main/causalNL.py>

Table 5. This ablation study shows the test accuracy % on CIFAR10 under IDN at noise rate 0.5. First, we show the result of CausalNL [66]. Second, we show the result of CausalNL [66] with Co-teaching [17] replaced by DivideMix [33] (without Continuous Bernoulli reconstruction). Then at last we show the results of our proposed algorithm InstanceGM.

Method	Test Accuracy (%)
CausalNL [66]	78.63
CausalNL [66] + DivideMix [33]	88.62
* w/o continuous Bernoulli	
InstanceGM	95.90

Table 6. This ablation study shows the test accuracy % on ANIMAL-10N using various architectures (without self-supervision), including ResNet [20], VGG [52] and ConvNeXt [39] with InstanceGM. Table 3, reported the results of VGG [52] to provide a fair comparison with other methods.

Method	Test Accuracy (%)
InstanceGM with ResNet [20]	82.2
InstanceGM with VGG [52]	84.6
InstanceGM with ConvNeXt [39]	84.7

≈ 10% accuracy improvement from CausalNL, but the use of MSE reconstruction loss can still limit classification accuracy. Hence, by replacing the MSE loss by the continuous Bernoulli loss for image reconstruction, we notice a further ≈ 7% accuracy improvement.

For ANIMAL-10N [53], we test InstanceGM with various backbone networks (VGG [52], ResNet [20], and ConvNeXt [39]) and the results are displayed in Table 6. Due to the architectural differences, ConvNeXt [39] performed best on our proposed algorithm, but for a fair comparison with the other models, we use the VGG backbone [52] results in Table 3.

6. Conclusion

In this paper, we presented an instance-dependent noisy label learning algorithm method, called InstanceGM. InstanceGM explores generative and discriminative models [66], where for the generative model, we replace the usual MSE image reconstruction loss by the continuous Bernoulli reconstruction loss [40] that improves the training process, and for the discriminative model, we replace co-teaching by DivideMix [33] to enable the use of clean and noisy samples during training. We performed extensive experiments on various IDN benchmarks, and our results on CIFAR10, CIFAR100, Red Mini-Imagenet, ANIMAL-10N outperform the results of state-of-the-art methods, particularly in high noise rates and are competitive for CLOTHING-1M. The ablation study clearly shows the importance of the new continuous Bernoulli reconstruction loss [40] and DivideMix [33], with both improving classification accuracy from CausalNL [66].

References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel O’Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. In *International Conference on Machine Learning*, pages 312–321. PMLR, 2019. 1
- [2] Devansh Arpit, Stanisław Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, and Simon Lacoste-Julien. A closer look at memorization in deep networks. In *International Conference on Machine Learning*, volume 70, pages 233–242. PMLR, 2017. 1, 2
- [3] HeeSun Bae, Seungjae Shin, JoonHo Jang, Byeonghu Na, Kyungwoo Song, and Il-Chul Moon. From noisy prediction to true label: Noisy prediction calibration via generative model. In *International Conference on Machine Learning*, 2022. 1
- [4] Mélanie Bernhardt, Daniel C Castro, Ryutaro Tanno, Anton Schwaighofer, Kerem C Tezcan, Miguel Monteiro, Shruthi Bannur, Matthew P Lungren, Aditya Nori, and Ben Glocker. Active label cleaning for improved dataset quality under resource constraints. *Nature Communications*, 13(1):1–11, 2022. 2
- [5] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019. 2, 4, 5
- [6] Antonin Berthon, Bo Han, Gang Niu, Tongliang Liu, and Masashi Sugiyama. Confidence scores make instance-dependent label-noise learning possible. In *International Conference on Machine Learning*, pages 825–836. PMLR, 2021. 2
- [7] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33:9912–9924, 2020. 12
- [8] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *International Conference on Computer Vision*, pages 9650–9660, 2021. 7, 8, 12
- [9] Yingyi Chen, Xi Shen, Shell Xu Hu, and Johan AK Suykens. Boosting co-teaching with compression regularization for label noise. In *Conference on Computer Vision and Pattern Recognition*, pages 2688–2692, 2021. 7, 8
- [10] De Cheng, Tongliang Liu, Yixiong Ning, Nannan Wang, Bo Han, Gang Niu, Xinbo Gao, and Masashi Sugiyama. Instance-dependent label-noise learning with manifold-regularized transition matrix estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16630–16639, 2022. 2, 7
- [11] Lele Cheng, Xiangzeng Zhou, Liming Zhao, Dangwei Li, Hong Shang, Yun Zheng, Pan Pan, and Yinghui Xu. Weakly supervised learning with side information for noisy labeled images. In *European Conference on Computer Vision*, pages 306–321. Springer, 2020. 2
- [12] Filipe R Cordeiro, Vasileios Belagiannis, Ian Reid, and Gustavo Carneiro. Propmix: Hard sample filtering and proportional mixup for learning with noisy labels. In *British Machine Vision Conference*, 2021. 1, 2, 6, 7, 8, 12
- [13] Antonia Creswell, Kai Arulkumaran, and Anil A Bharath. On denoising autoencoders trained to minimise binary cross-entropy. *arXiv preprint arXiv:1708.08487*, 2017. 13
- [14] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015. 2
- [15] Alex Graves. Practical variational inference for neural networks. In *Advances in Neural Information Processing Systems*, volume 24, 2011. 4
- [16] Bo Han, Jiangchao Yao, Gang Niu, Mingyuan Zhou, Ivor Tsang, Ya Zhang, and Masashi Sugiyama. Masking: A new perspective of noisy supervision. In *Advances in Neural Information Processing Systems*, volume 31, 2018. 2
- [17] 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 *Advances in Neural Information Processing Systems*, volume 31, 2018. 1, 3, 5, 7, 8
- [18] Jiangfan Han, Ping Luo, and Xiaogang Wang. Deep self-learning from noisy labels. In *International Conference on Computer Vision*, pages 5138–5147, 2019. 6
- [19] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 12
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 8
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 6
- [22] Dan Hendrycks, Mantas Mazeika, Duncan Wilson, and Kevin Gimpel. Using trusted data to train deep networks on labels corrupted by severe noise. In *Advances in Neural Information Processing Systems*, volume 31, 2018. 2
- [23] Simon Jenni and Paolo Favaro. Deep bilevel learning. In *European Conference on Computer Vision*, pages 618–633, 2018. 2
- [24] Lu Jiang, Di Huang, Mason Liu, and Weilong Yang. Beyond synthetic noise: Deep learning on controlled noisy labels. In *International Conference on Machine Learning*, pages 4804–4815. PMLR, 2020. 6, 7
- [25] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International Conference on Machine Learning*, pages 2304–2313. PMLR, 2018. 7
- [26] Zhimeng Jiang, Kaixiong Zhou, Zirui Liu, Li Li, Rui Chen, Soo-Hyun Choi, and Xia Hu. An information fusion approach to learning with instance-dependent label noise.

- In *International Conference on Learning Representations*, 2021. 2, 8, 13
- [27] Taehyeon Kim, Jongwoo Ko, JinHwan Choi, and Se-Young Yun. FINE samples for learning with noisy labels. In *Advances in Neural Information Processing Systems*, volume 34, 2021. 1, 2
- [28] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2013. 2, 4
- [29] Shuming Kong, Yanyan Shen, and Linpeng Huang. Resolving training biases via influence-based data relabeling. In *International Conference on Learning Representations*, 2021. 2
- [30] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 2, 6, 7, 8
- [31] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, volume 25, 2012. 1
- [32] Neil Lawrence and Bernhard Schölkopf. Estimating a kernel fisher discriminant in the presence of label noise. In *International Conference on Machine Learning*, pages 306–306. Morgan Kaufmann, 2001. 2
- [33] Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *International Conference on Learning Representations*, 2020. 1, 2, 3, 4, 5, 6, 7, 8, 13
- [34] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen Awm Van Der Laak, Bram Van Ginneken, and Clara I Sánchez. A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42:60–88, 2017. 1
- [35] Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. In *Advances in Neural Information Processing Systems*, volume 33, pages 20331–20342, 2020. 1, 2, 6
- [36] Tongliang Liu and Dacheng Tao. Classification with noisy labels by importance reweighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(3):447–461, 2015. 7
- [37] Yang Liu. Understanding instance-level label noise: Disparate impacts and treatments. In *International Conference on Machine Learning*, pages 6725–6735. PMLR, 2021. 2
- [38] Yang Liu and Hongyi Guo. Peer loss functions: Learning from noisy labels without knowing noise rates. In *International conference on machine learning*, pages 6226–6236. PMLR, 2020. 1
- [39] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A ConvNet for the 2020s. In *Conference on Computer Vision and Pattern Recognition*, 2022. 8
- [40] Gabriel Loaiza-Ganem and John P Cunningham. The continuous Bernoulli: fixing a pervasive error in variational autoencoders. In *Advances in Neural Information Processing Systems*, volume 32, 2019. 2, 3, 8, 13
- [41] Xingjun Ma, Hanxun Huang, Yisen Wang, Simone Romano, Sarah Erfani, and James Bailey. Normalized loss functions for deep learning with noisy labels. In *International Conference on Machine Learning*, pages 6543–6553. PMLR, 2020. 2
- [42] Eran Malach and Shai Shalev-Shwartz. Decoupling “when to update” from “how to update”. In *Advances in Neural Information Processing Systems*, volume 30, 2017. 7
- [43] Aditya Krishna Menon, Ankit Singh Rawat, Sashank J Reddi, and Sanjiv Kumar. Can gradient clipping mitigate label noise? In *International Conference on Learning Representations*, 2019. 2
- [44] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, volume 30, 2017. 1, 2
- [45] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 6
- [46] 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 *Conference on Computer Vision and Pattern Recognition*, pages 1944–1952, 2017. 1, 7
- [47] Rui Qian, Xin Lai, and Xirong Li. 3D object detection for autonomous driving: A survey. *Pattern Recognition*, page 108796, 2022. 1
- [48] Rahul Ragesh, Sundararajan Sellamanickam, Arun Iyer, Ramakrishna Bairi, and Vijay Lingam. Hetegcn: Heterogeneous graph convolutional networks for text classification. In *ACM International Conference on Web Search and Data Mining*, pages 860–868, 2021. 1
- [49] Vikas C Raykar, Shipeng Yu, Linda H Zhao, Gerardo Hermosillo Valadez, Charles Florin, Luca Bogoni, and Linda Moy. Learning from crowds. *Journal of Machine Learning Research*, 11(4), 2010. 2
- [50] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *International conference on machine learning*, pages 4334–4343. PMLR, 2018. 2
- [51] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. 6
- [52] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015. 8
- [53] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. Selfie: Refurbishing unclean samples for robust deep learning. In *International Conference on Machine Learning*, pages 5907–5915. PMLR, 2019. 2, 6, 7, 8

- [54] Hwanjun Song, Minseok Kim, Dongmin Park, and Jae-Gil Lee. How does early stopping help generalization against label noise? In *ICML 2020 Workshop on Uncertainty and Robustness in Deep Learning*, 2019. [2](#)
- [55] Hwanjun Song, Minseok Kim, Dongmin Park, Yooju Shin, and Jae-Gil Lee. Robust learning by self-transition for handling noisy labels. In *ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 1490–1500, 2021. [2](#)
- [56] Hwanjun Song, Minseok Kim, Dongmin Park, Yooju Shin, and Jae-Gil Lee. Learning from noisy labels with deep neural networks: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. [1](#)
- [57] Andreas Veit, Neil Alldrin, Gal Chechik, Ivan Krasin, Abhinav Gupta, and Serge Belongie. Learning from noisy large-scale datasets with minimal supervision. In *Conference on Computer Vision and Pattern Recognition*, pages 839–847, 2017. [2](#)
- [58] Xinshao Wang, Yang Hua, Elyor Kodirov, and Neil M Robertson. IMAE for noise-robust learning: Mean absolute error does not treat examples equally and gradient magnitude’s variance matters. *arXiv preprint arXiv:1903.12141*, 2019. [2](#)
- [59] Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, and Shu-Tao Xia. Iterative learning with open-set noisy labels. In *Conference on Computer Vision and Pattern Recognition*, pages 8688–8696, 2018. [2](#)
- [60] Hongxin Wei, Lue Tao, Renchunzi Xie, and Bo An. Open-set label noise can improve robustness against inherent label noise. In *Advances in Neural Information Processing Systems*, volume 34, 2021. [2](#)
- [61] Xiaobo Xia, Tongliang Liu, Bo Han, Mingming Gong, Jun Yu, Gang Niu, and Masashi Sugiyama. Sample selection with uncertainty of losses for learning with noisy labels. *arXiv preprint arXiv:2106.00445*, 2021. [1](#)
- [62] Xiaobo Xia, Tongliang Liu, Bo Han, Nannan Wang, Mingming Gong, Haifeng Liu, Gang Niu, Dacheng Tao, and Masashi Sugiyama. Part-dependent label noise: Towards instance-dependent label noise. In *Advances in Neural Information Processing Systems*, volume 33, pages 7597–7610, 2020. [1](#), [2](#), [6](#), [7](#), [13](#)
- [63] Xiaobo Xia, Tongliang Liu, Nannan Wang, Bo Han, Chen Gong, Gang Niu, and Masashi Sugiyama. Are anchor points really indispensable in label-noise learning? In *Advances in Neural Information Processing Systems*, volume 32, 2019. [7](#)
- [64] Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *Conference on Computer Vision and Pattern Recognition*, pages 2691–2699, 2015. [2](#), [6](#), [7](#), [8](#)
- [65] Youjiang Xu, Linchao Zhu, Lu Jiang, and Yi Yang. Faster meta update strategy for noise-robust deep learning. In *Conference on Computer Vision and Pattern Recognition*, pages 144–153, June 2021. [2](#), [6](#), [7](#)
- [66] Yu Yao, Tongliang Liu, Mingming Gong, Bo Han, Gang Niu, and Kun Zhang. Instance-dependent label-noise learning under a structural causal model. In *Advances in Neural Information Processing Systems*, volume 34, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [67] Yu Yao, Tongliang Liu, Bo Han, Mingming Gong, Jiankang Deng, Gang Niu, and Masashi Sugiyama. Dual T: Reducing estimation error for transition matrix in label-noise learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 7260–7271, 2020. [1](#)
- [68] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017. [2](#)
- [69] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2017. [7](#)
- [70] Yivan Zhang, Gang Niu, and Masashi Sugiyama. Learning noise transition matrix from only noisy labels via total variation regularization. In *International Conference on Machine Learning*, pages 12501–12512. PMLR, 2021. [1](#), [2](#)
- [71] Yikai Zhang, Songzhu Zheng, Pengxiang Wu, Mayank Goswami, and Chao Chen. Learning with feature-dependent label noise: A progressive approach. In *International Conference on Learning Representations*, 2021. [6](#), [7](#)
- [72] 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 *Winter Conference on Applications of Computer Vision*, pages 1657–1667, 2022. [1](#), [2](#), [6](#)
- [73] Zhaowei Zhu, Tongliang Liu, and Yang Liu. A second-order approach to learning with instance-dependent label noise. In *Conference on Computer Vision and Pattern Recognition*, pages 10113–10123, June 2021. [2](#), [7](#)
- [74] Zhaowei Zhu, Yiwen Song, and Yang Liu. Clusterability as an alternative to anchor points when learning with noisy labels. In *International Conference on Machine Learning*, pages 12912–12923. PMLR, 2021. [2](#), [7](#)