

EvidentialMix: Learning with Combined Open-set and Closed-set Noisy Labels

Ragav Sachdeva¹, Filipe R. Cordeiro², Vasileios Belagiannis³, Ian Reid¹, and Gustavo Carneiro¹

¹University of Adelaide, Adelaide, Australia

²Universidade Federal Rural de Pernambuco, Recife, Brazil

³Universität Ulm, Ulm, Germany

Abstract

The efficacy of deep learning depends on large-scale data sets that have been carefully curated with reliable data acquisition and annotation processes. However, acquiring such large-scale data sets with precise annotations is very expensive and time-consuming, and the cheap alternatives often yield data sets that have noisy labels. The field has addressed this problem by focusing on training models under two types of label noise: 1) closed-set noise, where some training samples are incorrectly annotated to a training label other than their known true class; and 2) open-set noise, where the training set includes samples that possess a true class that is (strictly) not contained in the set of known training labels. In this work, we study a new variant of the noisy label problem that combines the open-set and closed-set noisy labels, and introduce a benchmark evaluation to assess the performance of training algorithms under this setup. We argue that such problem is more general and better reflects the noisy label scenarios in practice. Furthermore, we propose a novel algorithm, called EvidentialMix, that addresses this problem and compare its performance with the state-of-the-art methods for both closed-set and open-set noise on the proposed benchmark. Our results show that our method produces superior classification results and better feature representations than previous state-of-the-art methods. The code is available at <https://github.com/ragavsachdeva/EvidentialMix>.

1. Introduction

Deep learning has achieved outstanding results in several important classification problems, using large and well-curated training data [12, 3]. However, most of the interesting data sets available in the society are orders of magnitude larger, but poorly curated, which means that the data may contain acquisition and labeling mistakes that can lead to poor generalisation [26]. Therefore, one of the important

challenges of the field is the development of methods that can cope with such noisy label data sets. Lately, researchers have greatly fostered the development of this field by studying controlled synthetic label noise and discovering theories or methodologies that can then be applied to real-world noisy data sets.

The types of label noise investigated thus far can be classified into two categories – *closed-set* and *open-set* noise. Although these terms (‘closed-set’ and ‘open-set’) were coined only recently by Wang et al. in [24] where they introduced the open-set noisy label problem, the closed-set noisy label problem has been extensively studied since much before. When handling closed-set label noise, the majority of the learning algorithms assume a fixed set of training labels [14, 22]. In this setting, some of the training samples are annotated to an incorrect label, while their true class is present in the training label set. These mistakes can be completely random, where labels are flipped arbitrarily to an incorrect class, or consistent, when the annotator is genuinely confused about the annotation of a particular sample. A less studied label noise is the open-set noisy label problem [24], where we incorrectly sample some data observations, such that their true annotation is *not* contained within the set of known training labels. A hyperbolic example of such a setting could be the presence of a *horse* image in the training set for modelling a *cats vs dogs* binary classifier. As evident from their definitions, these two types of label noise are mutually exclusive, i.e., a given noisy label cannot be closed-set and open-set at the same time.

It is quite easy to substantiate that both open-set and closed-set noise are likely to co-occur in real-world data sets. For instance, recent methods for large scale data collection propose the use of querying commercial search engines (e.g., Google Images), where the search keywords serve as the labels of the queried images. It is evident from Figure 1 that collecting images using such methods can lead to both open-set and closed-set noise. However, thus far, no systematic study with *controlled* label noise has been presented, where the training data set contains both types of la-

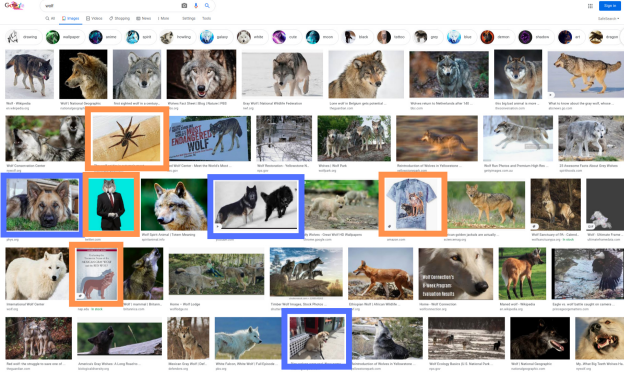


Figure 1: Results of a search engine query to collect data for a wolf-vs-dog binary classifier. The search keyword used here is “wolf”. The images bounded by an orange box are open-set noise (i.e. neither wolf nor dog) and the ones bounded by a blue box are closed-set noise (i.e. labelled as wolf but are actually a dog).

bel noise simultaneously. Even though there have been papers that evaluated their proposed methods on both [13, 24], the training data sets have been exclusively corrupted with either closed-set noise or open-set noise, but never in a combined fashion.

In this paper, we formulate a novel benchmark evaluation to address the noisy label learning problem that consists of a combination of closed-set and open-set noise. This proposed benchmark evaluation is defined by three variables: 1) the total proportion of label noise in the training set, represented by $\rho \in [0, 1]$; 2) the proportion of closed-set noise *within the set of samples containing noisy labels*, denoted by $\omega \in [0, 1]$ (this implies that $\rho \times \omega\%$ samples of the entire data set have a closed-set noisy label and $\rho \times (1 - \omega)\%$ samples of the entire data set have an open-set noisy label); and 3) the source of open-set noisy label data. Note that this setup generalises both types of label noise as it can collapse to one of the two noise types when $\omega \in \{0, 1\}$.

The state-of-the-art (SOTA) approaches that aim to solve the closed-set noisy label problem focus on methods that identify the samples that were incorrectly annotated and update their labels with semi-supervised learning (SSL) approaches [14] for the next training iteration. This strategy is likely to fail in the open-set problem because it assumes that there exists a correct class in the training labels for every training sample, which is not the case. On the other hand, the main approach addressing the open-set noise problem targets the identification of noisy samples to reduce their weights in the learning process [24]. Such strategy is inefficient in closed-set problems because the closed-set noisy label samples are still very meaningful during the SSL stage. Hence, to be robust in the scenarios where both closed-set and open-set noise samples are present, the learning algorithm must be able to identify the type of label noise affecting each training sample, and then either update the label,




Figure 2: Our proposed method, EvidentialMix, relies on two models, NetD and NetS. Both models are initially trained for a few epochs (see (1) WarmUp) using just a simple training mechanism that does not include any approach to handle label noise. Next in (2), we fit a ψ -component Gaussian Mixture Model (GMM) to the loss distribution of the samples from NetS in order to separate the training set into clean, closed-set and open-set samples. Following this, NetD is trained with the SSL learning mechanism [14] using the predicted clean and closed-set samples only. Finally, NetS is trained on the entire training set to minimise the subjective logic loss [21] using the labels estimated from NetD, and the process in (2) repeats.

if it is closed-set noise, or reduce its weight, if it is open-set noise. To achieve this, we propose a new learning algorithm, called EvidentialMix (EDM) – see Fig. 2. The key contributions of our proposed algorithm are the following:

- EDM is able to accurately distinguish between clean, open-set and closed-set samples, thus allowing it to exercise different learning mechanisms depending on the type of label noise. In comparison, previous methods [14, 24] can only separate clean samples from noisy ones, but not closed-noise from the open-noise samples.
- We show that our method can learn superior feature representations than previous methods as evident from the t-SNE plot in Figure 4, where our method has a unique cluster for each of the known label/class and another *separate* cluster for open-set samples. In comparison, previous methods have shown to largely overfit the open-set samples and incorrectly cluster them to one of the known classes.
- We experimentally show that EDM produces classification accuracy that is comparable or better than the previous methods on various label noise rates (including the extreme case where $\omega \in \{0, 1\}$).

2. Prior Work

There is an increasing interest in the study of modelling deep learning classifiers with noisy labels. For the closed-set noise, Reed et al. [19] proposed one of the first approaches that uses a transition matrix to learn how labels switch between different classes. The use of transition matrices has been further explored in many different

ways [18, 5], but none of them show competitive results, likely because they do not include mechanisms to identify and handle samples containing noisy labels. Data augmentation approaches [27] have been successively explored by closed-set noisy label methods, where the idea is that it can naturally increase the training robustness to label noise. Meta-learning is another technique explored in closed-set noisy label problems [15, 20], but the need for clean validation sets or artificial new training tasks makes this technique relatively unexplored. The use of curriculum learning (CL) [7] for closed-set problems has been explored to re-label training samples dynamically during training, based on their loss values. This approach has been extended with the training of multiple models [17, 25] that aim to focus the training on samples with small loss that are inconsistently classified by the multiple models. Recently, the explicit identification of noisy samples using negative learning has been explored by Kim et al. [9], with competitive results. Another important approach in handling label noise is model ensembling, as proposed by Tarvainen et al. [23]. The use of robust generative classifiers (RoG) to improve the performance of discriminative classifiers has been explored by Lee et al. [13], where they build an ensemble of robust linear discriminative models using features extracted from several layers of the trained discriminative model – in principle, this approach has the potential to improve the performance of any method and has been successively tested in closed-set and open-set scenarios.

The learning with open-set noisy labels has only recently been explored by Wang et al. [24], where the idea is to identify the samples containing noisy labels and reduce their weight in the training process since they almost certainly belong to a class not represented in the training set. Given that they are the only ones to explicitly address the open-set noise, their method is the main baseline for that problem.

The current SOTA for closed-set noisy label approaches are SELF [22] and DivideMix [14] – both consisting of methods that combine several of the approaches described above. SELF [22] combines model ensembling, re-labelling, noisy sample identification, and data augmentation; while DivideMix [14] uses multiple model training, noisy sample identification, and data augmentation [1]. These two approaches are likely to be vulnerable to open-set noise since they assume that training samples *must* belong to one of the training classes – an assumption that is not correct for open-set noise.

3. Method

3.1. Problem Definition

We define the training set as $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{|\mathcal{D}|}$, where the RGB image $\mathbf{x} : \Omega \rightarrow \mathbb{R}^3$ (Ω represents the image lattice), the set of training labels is denoted by \mathcal{Y} , which forms the standard basis of $|\mathcal{Y}|$ dimensions, $\mathbf{y} \in \{0, 1\}^{|\mathcal{Y}|}$




Figure 3: Per-sample loss distribution using subjective logic loss (Eq. (2)) with total noise ratio $\rho = 0.6$ and closed-set noise rate $\omega = 0.25$ (i.e. open-set noise rate is $1 - \omega = 0.75$) after WarmUp, with \mathcal{D} from CIFAR10 [11] and \mathcal{D}' from Imagenet32 [2].

($\sum_{c=1}^{|\mathcal{Y}|} \mathbf{y}(c) = 1$, representing a multi-class problem). Note that \mathbf{y}_i is the noisy label for \mathbf{x}_i and the hidden clean label is represented by \mathbf{y}_i^* .

For the **closed-set noise problem with noise rate** $\zeta \in [0, 1]$, we assume that $(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{D}$ is labelled as $\mathbf{y}_i = \mathbf{y}_i^*$ with probability $1 - \zeta$, and $\mathbf{y}_i \sim r(\mathcal{Y})$ with probability ζ , with $r(\mathcal{Y}, \theta_r)$ representing a random function that picks one of the labels in \mathcal{Y} following a particular distribution parameterised by θ_r .

For the **open-set noise problem with noise rate** $\eta \in [0, 1]$, we need to define a new training set \mathcal{D}' (with $\mathcal{D}' \cap \mathcal{D} = \emptyset$), where the label set for \mathcal{D}' is represented by \mathcal{Y}' (with $\mathcal{Y}' \cap \mathcal{Y} = \emptyset$) – this means that the images in \mathcal{D}' no longer have labels in \mathcal{Y} . In such open-set problem, a proportion of $1 - \eta$ samples is drawn with $(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{D}$ with $\mathbf{y}_i = \mathbf{y}_i^*$, while a proportion η of samples are obtained with $(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{D}'$ with $\mathbf{y}_i \sim r(\mathcal{Y}, \theta_r)$.

The **combined closed-set and open-set problem with rates** $\rho, \omega \in [0, 1]$ is defined by mixing the two types of noise above. More specifically, $1 - \rho$ of the training set contains images $(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{D}$ annotated with $\mathbf{y}_i = \mathbf{y}_i^*$, while $\omega \times \rho$ images are sampled as $(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{D}$ with label $\mathbf{y}_i \sim r(\mathcal{Y}, \theta_r)$, and $(1 - \omega) \times \rho$ images belong to \mathcal{D}' and labelled with $\mathbf{y}_i \sim r(\mathcal{Y}, \theta_r)$.

3.2. Noise Classification

The main impediment when dealing with this problem is the need to identify closed-set and open-set noisy samples since they must be dealt differently by the method. One possible way of dealing with this problem is by associating closed-set samples with high losses computed from confident but incorrect classification [14], and open-set samples with uncertain classification. To achieve this, we propose the use of the subjective logic (SL) loss function [21] that relies on the theory of evidential reasoning and SL to quantify classification uncertainty. The SL loss makes use of the Dirichlet distribution to represent subjective opinions, encoding belief and uncertainty. A network trained with the SL loss tries to learn parameters of a predictive posterior as a Dirichlet density function for the classification of the training samples. The resulting output for a given sample

is considered as the evidence for the classification of that sample over a set of class labels [8]. Figure 3 shows the per-sample loss distribution of training samples from a network trained using the SL loss. The separation between clean, closed-set and open-set samples is easy to capture.

3.3. EvidentialMix

Our proposed EvidentialMix simultaneously trains two networks: NetS, which uses the SL loss [21], and NetD, which uses the SSL training mechanism and the DivideMix (DM) loss [14]. Broadly speaking, the ability of the SL loss to estimate classification uncertainty allows NetS to divide the training set into clean-set, open-set and closed-set samples. The predicted clean-set and closed-set samples are then used to train NetD using MixMatch as outlined in [14], while the predicted open-set samples are discarded for that epoch. Following this, NetD re-labels the *entire* training data set (including predicted open-set samples) that are then used to train NetS.

As NetS iteratively learns from the labels predicted by NetD, it gets better at splitting the data into the three sets. This is so because the labels from NetD become more accurate over the training process given that it is only trained on clean and closed-set samples, and never on predicted open-set samples. The two networks thus complement each other to produce accurate classification results for the combined closed-set and open-set noise problem. A detailed explanation is outlined below, while Alg. 1 delineates the full training algorithm.

Algorithm 1 trains NetD, represented by $f_{\theta(D)}(c|\mathbf{x})$, and NetS, denoted by $f_{\theta(S)}(c|\mathbf{x})$ (with $c \in \{1, \dots, |\mathcal{Y}|\}$) – both of which return a logit in $\mathbb{R}^{|\mathcal{Y}|}$. In the warm-up stage (see line WarmUp(\mathcal{D})), we train both models for a limited number of epochs using the cross entropy loss for $f_{\theta(D)}(c|\mathbf{x})$, where probability is obtained by applying a softmax activation $\sigma(\cdot)$ to form $p_{\theta(D)}(c|\mathbf{x}) = \sigma(f_{\theta(D)}(c|\mathbf{x}))$, and the following (SL) loss for $f_{\theta(S)}(c|\mathbf{x})$ [21]:

$$\mathcal{L}^{(S)} = \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} \ell^{(S)}(\mathbf{x}_i, \mathbf{y}_i, \theta^{(S)}), \quad (1)$$

with [21]:

$$\ell^{(S)}(\mathbf{x}_i, \mathbf{y}_i, \theta^{(S)}) = \sum_{c=1}^{|\mathcal{Y}|} (\mathbf{y}_i(c) - \alpha_{ic}/S_i)^2 + \frac{\alpha_{ic}(S_i - \alpha_{ic})}{S_i^2(S_i + 1)}, \quad (2)$$

where $\alpha_{ic} = \varphi(f_{\theta(S)}(c|\mathbf{x}_i)) + 1$ for class $c \in \{1, \dots, |\mathcal{Y}|\}$, with $\varphi(\cdot)$ representing the ReLU activation function, and $S_i = \sum_{c=1}^{|\mathcal{Y}|} \alpha_{ic}$.

The classification of samples into the clean, closed-set, and open-set is performed using the SL loss values from Eq. (2) for the entire training set \mathcal{D} . More specifically, we take the set of losses $\{\ell^{(S)}(\mathbf{x}_i, \mathbf{y}_i, \theta^{(S)})\}_{i=1}^{|\mathcal{D}|}$ and fit a ψ -component Gaussian mixture model (GMM) using the

Algorithm 1: EvidentialMix (EDM)

Input: $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{|\mathcal{D}|}$, number of augmentations M , temperature sharpening T , loss weights $\lambda^{(\mathcal{U})}$ and $\lambda^{(reg)}$, MixMatch parameter α , number of epochs E .

```

1  $f_{\theta(D)}(c|\mathbf{x}), f_{\theta(S)}(c|\mathbf{x}) = \text{WarmUp}(\mathcal{D})$ 
2 while  $e < E$  do
3    $\mathcal{W}, \mathcal{W}^{\text{op}}, \mathcal{W}^{\text{cl}} = \text{GMM}(\mathcal{D}, f_{\theta(S)}(c|\mathbf{x}))$ 
   // Train NetD
4    $\mathcal{X} = \{(\mathbf{x}_i, \mathbf{y}_i, w_i) | (\mathbf{x}_i, \mathbf{y}_i, w_i) \in (\mathcal{D}, \mathcal{W}), w_i >$ 
      $\max(w_i^{\text{op}}, w_i^{\text{cl}})\}$ 
5    $\mathcal{U} = \{\mathbf{x}_i | (\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{D}, w_i^{\text{cl}} > \max(w_i, w_i^{\text{op}})\}$ 
6   for  $\text{iter}=1$  to  $\text{num\_iters}$  do
7      $\{(\mathbf{x}_b, \mathbf{y}_b, w_b)\}_{b=1}^B \subset \mathcal{X}$  // randomly
       pick  $B$  samples from  $\mathcal{X}$ 
8      $\{\mathbf{u}_b\}_{b=1}^B \subset \mathcal{U}$  // randomly pick  $B$ 
       samples from  $\mathcal{U}$ 
9     for  $b=1$  to  $B$  do
10      for  $m=1$  to  $M$  do
11         $\hat{\mathbf{x}}_{b,m} = \text{DataAugment}(\mathbf{x}_b)$ 
12         $\hat{\mathbf{u}}_{b,m} = \text{DataAugment}(\mathbf{u}_b)$ 
13      end
14      for  $c=1$  to  $|\mathcal{Y}|$  do
15         $\mathbf{p}_b(c) = \frac{1}{M} \sum_m p_{\theta(D)}(c|\hat{\mathbf{x}}_{b,m})$ 
16         $\mathbf{q}_b(c) = \frac{1}{M} \sum_m p_{\theta(D)}(c|\hat{\mathbf{u}}_{b,m})$ 
17      end
18       $\hat{\mathbf{y}}_b =$ 
        TempSharpen $_T(w_b \mathbf{y}_b + (1 - w_b) \mathbf{p}_b)$ 
19       $\hat{\mathbf{q}}_b = \text{TempSharpen}_T(\mathbf{q}_b)$ 
20    end
21     $\hat{\mathcal{X}} = \{(\hat{\mathbf{x}}_{b,m}, \hat{\mathbf{y}}_b)\}_{b \in (1, \dots, B), m \in (1, \dots, M)}$ 
22     $\hat{\mathcal{U}} = \{(\hat{\mathbf{u}}_{b,m}, \hat{\mathbf{q}}_b)\}_{b \in (1, \dots, B), m \in (1, \dots, M)}$ 
23     $\mathcal{X}', \mathcal{U}' = \text{MixMatch}_{\alpha}(\cdot)$ 
24     $\theta^{(D)} = \text{SGD}(\mathcal{L}^{(D)}, \theta^{(D)}, \mathcal{X}', \mathcal{U}')$ 
25  end
  // Train NetS
26  for  $i=1$  to  $|\mathcal{D}|$  do
27     $\hat{c}_i =$ 
       $\arg \max_{c \in \mathcal{Y}} [(w_i^{\text{cl}}) p_{\theta(D)}(c|\mathbf{x}_i) + (1 - w_i^{\text{cl}}) \mathbf{y}_i(c)]$ 
28     $\hat{\mathbf{y}}_i = \text{onehot}(\hat{c}_i)$ 
29  end
30   $\theta^{(S)} = \text{SGD}(\mathcal{L}^{(S)}, \theta^{(S)}, \{(\mathbf{x}_i, \hat{\mathbf{y}}_i)\}_{i=1}^{|\mathcal{D}|})$ 
31 end
```

Expectation-Maximization algorithm. The idea we explore in this paper lies in the fact that the model output for the **clean samples** will tend to be *confident and at the same time agree with its original label, producing a small loss*. The model output for **closed-set noise samples** will also tend to be *confident but at the same time disagree with its origi-*

nal label, generating a large loss value. The model output for **open-set noise samples**, however, will *not produce a confident output, resulting in a loss value that is neither large nor small*. Therefore, the multi-component GMM will capture each of these sets, where the clean probability w_i is the posterior probability $p(\mathcal{G}|\ell^{(S)}(\mathbf{x}_i, \mathbf{y}_i, \theta^{(S)}))$, where \mathcal{G} denotes the set of Gaussian components with mean $\leq \mu_{min}$ (i.e., small losses), the closed-set posterior probability $w_i^{cl} = p(\mathcal{G}^{cl}|\ell^{(S)}(\mathbf{x}_i, \mathbf{y}_i, \theta^{(S)}))$ is computed from the components with mean $\geq \mu_{max}$ (i.e., large losses), and the open-set probability w_i^{op} is the posterior of the remaining components with means $\in (\mu_{min}, \mu_{max})$ – these posterior probabilities form the three sets $\mathcal{W} = \{w_i\}_{i=1}^{|\mathcal{D}|}$, $\mathcal{W}^{cl} = \{w_i^{cl}\}_{i=1}^{|\mathcal{D}|}$ and $\mathcal{W}^{op} = \{w_i^{op}\}_{i=1}^{|\mathcal{D}|}$. Using these posteriors, we can build the set of clean samples, represented by \mathcal{X} with samples where the clean posterior probability w_i is larger than the other two probabilities, and the closed-set, denoted by \mathcal{U} , containing samples that have the closed-set posterior probability w_i^{cl} larger than the other two.

Next, we train NetD with the clean set \mathcal{X} and closed-set \mathcal{U} , defined above. A mini-batch is sampled from \mathcal{X} and \mathcal{U} , and we augment each sample in each set M times [14]. The average classification probabilities for the clean and closed-set samples are then computed using the M augmented samples, which after temperature sharpening, denoted by $\text{TempSharpen}_T(\cdot)$ with T denoting the temperature, form the ‘new’ samples and labels for the clean and closed-set samples, $\hat{\mathcal{X}} = \{(\hat{\mathbf{x}}_{b,m}, \hat{\mathbf{y}}_b)\}_{b,m=1}^{B,M}$ and $\hat{\mathcal{U}} = \{(\hat{\mathbf{u}}_{b,m}, \hat{\mathbf{q}}_b)\}_{b,m=1}^{B,M}$, respectively. The last stage before stochastic gradient descent (SGD) is the MixMatch process [1], where samples from the $\hat{\mathcal{X}}$ and $\hat{\mathcal{U}}$ are linearly combined to form \mathcal{X}' and \mathcal{U}' . SGD minimises the DM loss that combines the following functions [14]:

$$\mathcal{L}^{(D)} = \mathcal{L}^{(\mathcal{X})} + \lambda^{(\mathcal{U})} \mathcal{L}^{(\mathcal{U})} + \lambda^{(reg)} \mathcal{L}^{(reg)}, \quad (3)$$

where $\lambda^{(\mathcal{U})}$ denotes the weight of the loss associated with the unlabelled data set and $\lambda^{(reg)}$ weights the regularisation loss. The loss terms in Eq. (3) are defined by

$$\mathcal{L}^{(reg)} = \sum_{c=1}^{|\mathcal{Y}|} \frac{1}{|\mathcal{Y}|} \log \left(\frac{1}{|\mathcal{X}'| + |\mathcal{U}'|} \sum_{\mathbf{x} \in (\mathcal{X}' \cup \mathcal{U}')} p_{\theta^{(D)}}(c|\mathbf{x}) \right), \quad (4)$$

and

$$\begin{aligned} \mathcal{L}^{(\mathcal{X})} &= -\frac{1}{|\mathcal{X}'|} \sum_{(\hat{\mathbf{x}}, \hat{\mathbf{y}}) \in \mathcal{X}'} \sum_{c=1}^{|\mathcal{Y}|} \hat{\mathbf{y}}(c) \log(p_{\theta^{(D)}}(c|\hat{\mathbf{x}})), \\ \mathcal{L}^{(\mathcal{U})} &= \frac{1}{|\mathcal{U}'|} \sum_{(\hat{\mathbf{u}}, \hat{\mathbf{q}}) \in \mathcal{U}'} \|\hat{\mathbf{q}} - p_{\theta^{(D)}}(\cdot|\hat{\mathbf{u}})\|_2^2, \end{aligned} \quad (5)$$

where $p_{\theta^{(D)}}(\cdot|\hat{\mathbf{u}}) \in [0, 1]^{|\mathcal{Y}|}$ represents the model output of all $|\mathcal{Y}|$ labels for input $\hat{\mathbf{u}}$. After train-

ing NetD, we train the NetS model $f_{\theta^{(S)}}(c|\mathbf{x})$ by minimising the SL loss Eq. (1) with an updated training set, represented by $\{(\mathbf{x}_i, \hat{\mathbf{y}}_i)\}_{i=1}^{|\mathcal{D}|}$, formed by $\hat{c}_i = \arg \max_c [(w_i^{cl})p_{\theta^{(D)}}(c|\mathbf{x}_i) + (1 - w_i^{cl})\mathbf{y}_i(c)]$, for $i \in \{1, \dots, |\mathcal{D}|\}$ that produces the new labels $\hat{\mathbf{y}}_i = \text{onehot}(\hat{c}_i)$.

The inference for a test sample \mathbf{x} relies entirely on the NetD classifier, as follows: $c^* = \arg \max_{c \in \mathcal{Y}} p_{\theta^{(D)}}(c|\mathbf{x})$.

3.4. Implementation

We train an 18-layer PreAct Resnet [6] (for both NetS and NetD) using stochastic gradient descent (SGD) with momentum of 0.8, wight decay of 0.0005 and batch size of 64. The learning rate is 0.02 for WarmUp and for 100 epochs in the main training process, which is reduced to 0.002 afterwards. The WarmUp stage lasts for 10 and 30 epochs for NetD and NetS, respectively, where NetD is trained with a cross-entropy loss (i.e., $\mathcal{L}^{(\mathcal{X})}$ in Eq. (5) using the unchanged training set \mathcal{D}) while NetS is trained with the subjective logic loss $\mathcal{L}^{(S)}$ in Eq. (1) also using \mathcal{D} . After WarmUp, both models are trained for $E = 200$ epochs. Similar to [14], the number of data augmented samples is $M = 2$, the sharpening temperature is $T = 0.5$, the Mix-Match parameter is $\alpha = 4$, and the regularisation weight for the DM loss in Eq. (3) is $\lambda^{(reg)} = 1$. However, unlike [14] that manually selects the value of $\lambda^{(\mathcal{U})}$ based on the value of ρ , we set $\lambda^{(\mathcal{U})} = 25$ for all our experiments. For the GMM, we use $\psi = 20$ components, $\mu_{min} = 0.3$, and $\mu_{max} = 0.7$ since these values produced stable results.

4. Experiments

Following the prior work on closed-set and open-set noise problems [14, 24, 13], we conduct our experiments on the CIFAR-10 data set [11] for closed-set noise [14, 13]; and include the CIFAR-100 (small-scale) [11] and ImageNet32 (large-scale) [2] data sets for the open-set noise scenario [24, 13]. CIFAR-10 has 10 classes with 5000 32×32 pixel training images per class (forming a total of 50000 training images), and a testing set with 10000 32×32 pixel images with 1000 images per class. CIFAR-100 has 100 classes with 5000 32×32 pixel images per class and ImageNet32 is a down-sampled variant of ImageNet [3], with 1281149 images and 1000 classes, but resized to 32×32 pixels per image. All data sets above have been set up with curated labels, so below we introduce a new noisy label benchmark evaluation that combines closed-set and open-set synthetic label noise.

4.1. Combined Open-set and Closed-set Noisy Label Benchmark

The proposed benchmark is defined by the rate of label noise in the experiment, denoted by $\rho \in \{0.3, 0.6\}$, and the proportion of closed-set noise in the label noise, denoted by $\omega \in \{0, 0.25, 0.5, 0.75, 1\}$. The closed-set label noise is

		ρ	0.3					0.6				
		ω	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
ImageNet32	RoG [13]	Best	91.9	90.7	90.2	89.6	89.5	87.8	85.7	84.5	83.1	82.9
		Last	91.0	88.7	86.6	86.2	83.9	85.9	78.1	70.3	64.7	59.8
	ILON [24]	Best	91.8	90.7	88.0	86.5	85.8	87.7	83.4	81.2	78.7	77.3
		Last	90.6	86.9	82.0	77.3	72.7	85.5	72.6	58.9	54.4	46.5
	DivideMix [14]	Best	92.4	92.5	93.4	93.9	94.3	92.5	92.8	93.2	93.9	94.7
		Last	92.0	92.5	93.0	93.7	94.1	92.5	92.2	92.8	93.2	94.6
	EDM (Ours)	Best	93.2	94.4	94.7	95.1	95.2	91.2	93.7	94.0	94.1	94.1
		Last	92.5	93.7	94.5	94.7	94.8	90.9	93.1	93.4	93.9	94.1
CIFAR-100	RoG [13]	Best	91.4	90.9	89.8	90.4	89.9	88.2	85.2	84.1	83.7	83.1
		Last	89.8	87.4	85.9	84.9	84.5	82.1	72.9	66.3	62.0	59.5
	ILON [24]	Best	90.4	88.7	87.4	87.2	86.3	83.4	82.6	80.5	78.4	77.1
		Last	87.4	84.3	80.0	74.6	73.8	78.0	67.9	55.2	48.7	45.6
	DivideMix [14]	Best	89.3	90.5	91.5	93.0	94.3	89.0	90.6	91.8	93.4	94.4
		Last	88.7	90.1	90.9	92.8	94.0	88.7	89.8	91.5	93.0	94.3
	EDM (Ours)	Best	92.9	93.8	94.5	94.8	95.3	90.6	92.9	93.4	93.7	94.3
		Last	91.9	93.1	94.0	94.5	95.1	89.4	91.4	92.8	93.4	94.0

Table 1: Benchmark results of all competing methods and our proposed EDM. Clean data was sampled from CIFAR10, while the open-set noise came from ImageNet32 and CIFAR-100. The total noise in the training data is represented by $\rho \in \{0.3, 0.6\}$, where the closed-set proportion of this noise is $\omega \in \{0, 0.25, 0.5, 0.75, 1\}$ and open-set proportion is $1 - \omega$.

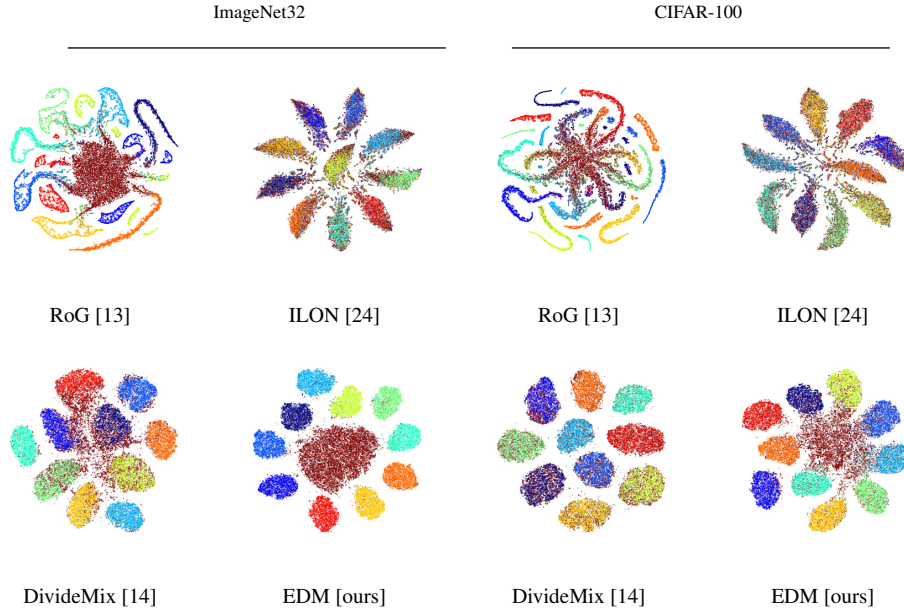


Figure 4: t-SNE plots of the related methods and our proposed EDM, where the total noise rate is $\rho = 0.6$ with the closed-set proportion being $\omega = 0.5$, and CIFAR-100 and ImageNet32 representing open-set data sets. The **brown** samples represent the open-set noise, while the other colours denote the true CIFAR-10 classes.

simulated by randomly selecting $\rho \times \omega \times 100\%$ of the training samples from CIFAR-10, and symmetrically shuffling their label, similarly to the synthetic label noise used in [14]. The open-set label noise is simulated by randomly selecting $\rho \times (1 - \omega) \times 100\%$ of the training images from CIFAR-

10 and replacing them with images randomly selected from either CIFAR-100 [11] or ImageNet32 [2], where a CIFAR-10 label is randomly assigned to each one of these images. Results are based on the classification accuracy on the clean testing set from CIFAR-10, using the benchmark proposed

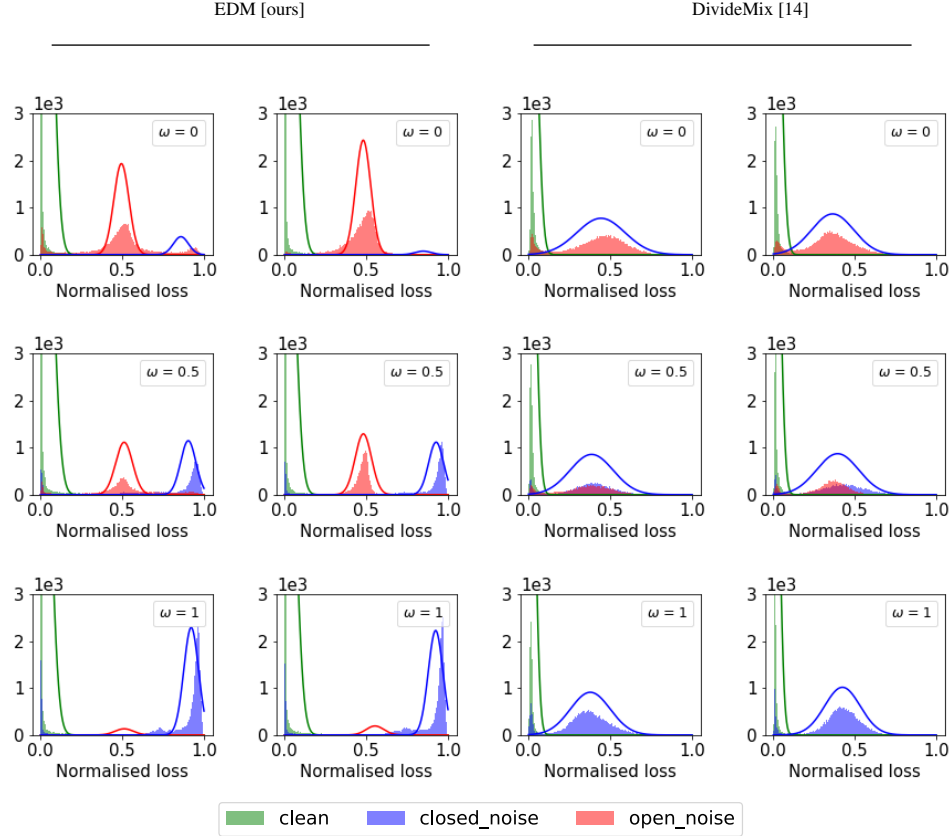


Figure 5: Per-sample loss distributions for the open-set (orange), closed set (blue) and clean set (green) samples produced by EDM (left) and DivideMix (right) at epoch $e = 100$, where $\rho = 0.6$ and $\omega \in \{0, 0.5, 1\}$, with open-set datasets $\in \{\text{CIFAR-100, ImageNet32}\}$. We also show the estimated GMM posterior probabilities for the clean, open-set and closed-set noise samples using our EDM, and for the clean and noisy samples using DivideMix [14].

above. We also show a comparison of the sample distribution between EDM and other related approaches in the feature space using t-SNE [16], and the effectiveness of EDM to separate clean, closed-set and open-set noisy samples.

4.2. Related Approaches for Comparison

We compare our proposed approach with the three methods listed below:

DivideMix¹ [14] is the current SOTA method that converts the problem of closed-set noisy label learning into a semi-supervised learning problem. It follows a multiple-model approach that splits the training data into clean and noisy subsets by fitting a 2-component Gaussian Mixture Model (GMM) to the loss values of the training samples at each epoch. Next, the framework discards the labels of the predicted noisy samples and uses MixMatch [1] to train the model.

ILON² [24] introduces the open-set noisy label learning

¹We used the publicly available code provided by the authors of the paper to produce our results.

²As the authors did not make their code publicly available, we imple-

mented their method from scratch and trained a Siamese network of 18-layer PreAct Resnet to produce our results.

problem, where the proposed approach is based on an iterative solution that re-weights the samples based on the outlier measure of the Local Outlier Factor Algorithm (LOF) [10]. **RoG**¹ [13] builds an ensemble of generative classifiers formed from features extracted from multiple layers of the ResNet model. The authors of RoG tested their approach on both closed-set noise and open-set noise *separately* which makes it an important benchmark to consider in our *combined* setup.

4.3. Results and Discussion

Classification accuracy: Table 1 shows the results computed from the benchmark evaluation of the proposed EDM, in comparison with the results of RoG [13], ILON [24], and DivideMix [14]. The evaluation relies on different rates of label noise ($\rho \in \{0.3, 0.6\}$) and closed-set noise ($\omega \in \{0, 0.25, 0.5, 0.75, 1\}$) using CIFAR-100 and ImageNet32 as open-set data sets. Results show that our method EDM outperforms all the competing approaches for 17 out

of the 20 noise settings and is a close second on the remaining 3. For $\rho = 0.3$, EDM produces better results than all competing methods for all values of ω and choice of open-set data sets, with an improvement of more than 3% over the next best method in some cases. On the other hand, both RoG and ILON perform significantly worse than EDM and DivideMix, particularly for $\rho = 0.6$ where the difference in accuracy is over 15% in some cases. In general, RoG and ILON are observed to perform worse when the proportion of closed-set noise increases, while the converse is true for DivideMix and EDM. It is also apparent that EDM is more robust to open-set noise than DivideMix, as evident from classification results when ω is small.

Feature representations: We show the t-SNE plots [16] in Fig. 4 for all the methods for the case where total noise rate $\rho = 0.6$, with the closed-set proportion $\omega = 0.5$ using CIFAR-100 and ImageNet32 as open-set data sets. In particular, the features for all methods are extracted from the last layer of the models (in our case, we use the features from the NetD, which is the one used for classification, as explained in Sec. 3.3). In the visualisation, the brown samples are from the open-set data sets, while all other colours represent the true CIFAR-10 classes. This clearly shows that our proposed EDM is quite effective at separating open-set samples from the other clean and closed-set training samples, while DivideMix and ILON largely overfit these samples, as evident from the spread of open-set samples around CIFAR-10 classes. Interestingly, RoG also shows good separation, but with apparently more complex distributions than EDM.

Noise classification: Fig. 5 shows the distribution of loss function values for the clean, open-set and closed-set samples in cases where the noise rate $\rho = 0.6$, with closed-set rates $\omega \in \{0, 0.5, 1\}$ using samples from both CIFAR-100 and ImageNet32 as open-set noise. From these graphs, it is clear that the SL loss in Eq. (2) is able to successfully distinguish samples from each one of the three sets above, even when only one of the noise types is available, such as the case when $\omega \in \{0, 1\}$. This suggests that the exploration of uncertainty in the SL loss to identify samples belonging to open-set noise is effective. Among the methods tested in this paper, DivideMix [14] also tries to separate the training samples into clean and noisy sets using the loss in (3). However, the resulting distribution seems inadequate to allow for a clear separation between the three sets because the open-set and closed-set noisy labels are basically indistinguishable. Consequently, DivideMix is able to separate clean samples from noisy samples, but not closed-set noise from open-set noise, thus forcing it to treat both noise types similarly during the training (i.e., both types are treated as closed-set noise). This is not ideal given that the open-set samples will be allocated to one of the incorrect training labels, which can ultimately cause the training to overfit these samples.

5. Conclusion

In this paper, we investigate a variant of the noisy label problem that combines open-set [24, 13] and closed-set noisy labels [14, 13]. To test various methods for this new problem, we propose a new benchmark that systematically changes the total noise rate and the proportion of closed-set and open-set noise. The open-set samples were sourced from either a small-scale data set (CIFAR100) or a large-scale data set (ImageNet32) such that the true label of these samples is not contained in the primary data set (CIFAR10). We argue that such a problem setup is more general and similar to real-life noisy label scenarios. We then propose the EvidentialMix algorithm to successfully address this new noise type with the use of the subjective logic loss [21] that produces low loss for clean samples, high loss for closed-set noisy samples, and mid-range loss for open-set samples. The clear division of the training data allows us to (1) identify and thereby remove the open-set samples from training to avoid overfitting them, given that they do not belong to any of the known classes, and (2) learn from the predicted closed-set samples in a semi-supervised fashion as in [14]. The evaluation shows that our proposed EDM is more effective to address this new combined open-set and closed-set label noise problem than the current state-of-the-art approaches for closed-set problems [14, 13] and open-set problems [24, 13].

Future work: The motivation for introducing this problem was to open the dialogue in the research community to investigate the combined open-set and closed-set label noise. Moving forward, we aim to explore more challenging noise settings such as incorporating asymmetric [19] and semantic noise [13] in the proposed combined label noise problem. Since we are the first ones to address this problem in a *controlled* setup, there is no precedent on how these more challenging noise scenarios could be meaningfully incorporated. For instance, even though asymmetric closed-set noise has previously been studied in the literature [19], it is not obvious what its counterpart, asymmetric open-set noise entails; for instance, it is not immediately clear how to build a noise transition matrix between CIFAR10 and ImageNet classes. In addition, we see merit in investigating other types of uncertainty to identify open-set noise, such as with Bayesian learning [4], and aim to explore such methods.

Acknowledgements: IR and GC gratefully acknowledge the support of the Australian Research Council through the Centre of Excellence for Robotic Vision CE140100016 and Future Fellowship (to GC) FT190100525. GC acknowledges the support by the Alexander von Humboldt-Stiftung for the renewed research stay sponsorship. RS acknowledges the support by the Playford Trust Honours Scholarship.

References

- [1] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. MixMatch: A Holistic Approach to Semi-Supervised Learning. *arXiv e-prints*, page arXiv:1905.02249, May 2019.
- [2] Patryk Chrabaszcz, Ilya Loshchilov, and Frank Hutter. A downsampled variant of imagenet as an alternative to the cifar datasets. *arXiv preprint arXiv:1707.08819*, 2017.
- [3] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [4] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- [5] Jacob Goldberger and Ehud Ben-Reuven. Training deep neural-networks using a noise adaptation layer. 2017.
- [6] 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.
- [7] 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, 2018.
- [8] Lance Kaplan, Federico Cerutti, Murat Sensoy, Alun Preece, and Paul Sullivan. Uncertainty Aware AI ML: Why and How. *arXiv e-prints*, page arXiv:1809.07882, Sept. 2018.
- [9] Youngdong Kim, Junho Yim, Juseung Yun, and Junmo Kim. Nlnl: Negative learning for noisy labels. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 101–110, 2019.
- [10] Hans-Peter Kriegel, Peer Kröger, Erich Schubert, and Arthur Zimek. Loop: local outlier probabilities. In *Proceedings of the 18th ACM conference on Information and knowledge management*, pages 1649–1652, 2009.
- [11] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [12] Yann LeCun, Y. Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521:436–44, 05 2015.
- [13] Kimin Lee, Sukmin Yun, Kibok Lee, Honglak Lee, Bo Li, and Jinwoo Shin. Robust inference via generative classifiers for handling noisy labels. *arXiv preprint arXiv:1901.11300*, 2019.
- [14] Junnan Li, Richard Socher, and Steven C. H. Hoi. DivideMix: Learning with Noisy Labels as Semi-supervised Learning. *arXiv e-prints*, page arXiv:2002.07394, Feb. 2020.
- [15] Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S Kankanhalli. Learning to learn from noisy labeled data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5051–5059, 2019.
- [16] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [17] Eran Malach and Shai Shalev-Shwartz. Decoupling” when to update” from” how to update”. In *Advances in Neural Information Processing Systems*, pages 960–970, 2017.
- [18] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1944–1952, 2017.
- [19] Scott Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. Training Deep Neural Networks on Noisy Labels with Bootstrapping. *arXiv e-prints*, page arXiv:1412.6596, Dec. 2014.
- [20] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. *arXiv preprint arXiv:1803.09050*, 2018.
- [21] Murat Sensoy, Lance Kaplan, and Melih Kandemir. Evidential deep learning to quantify classification uncertainty, 2018.
- [22] Duc Tam Nguyen, Chaithanya Kumar Mummadi, Thi Phuong Nhung Ngo, Thi Hoai Phuong Nguyen, Laura Beggel, and Thomas Brox. SELF: Learning to Filter Noisy Labels with Self-Ensembling. *arXiv e-prints*, page arXiv:1910.01842, Oct. 2019.
- [23] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*, pages 1195–1204, 2017.
- [24] Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, and Shu-Tao Xia. Iterative Learning with Open-set Noisy Labels. *arXiv e-prints*, page arXiv:1804.00092, Mar. 2018.
- [25] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor W Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? *arXiv preprint arXiv:1901.04215*, 2019.
- [26] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
- [27] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.