

Learning with Neighbor Consistency for Noisy Labels

Ahmet Iscen¹ Jack Valmadre*² Anurag Arnab¹ Cordelia Schmid¹
¹Google Research ²AIML, University of Adelaide

Abstract

Recent advances in deep learning have relied on large, labelled datasets to train high-capacity models. However, collecting large datasets in a time- and cost-efficient manner often results in label noise. We present a method for learning from noisy labels that leverages similarities between training examples in feature space, encouraging the prediction of each example to be similar to its nearest neighbours. Compared to training algorithms that use multiple models or distinct stages, our approach takes the form of a simple, additional regularization term. It can be interpreted as an inductive version of the classical, transductive label propagation algorithm. We thoroughly evaluate our method on datasets evaluating both synthetic (CIFAR-10, CIFAR-100) and realistic (mini-WebVision, WebVision, Clothing1M, mini-ImageNet-Red) noise, and achieve competitive or state-of-the-art accuracies across all of them.

1. Introduction

While deep learning can achieve unprecedented accuracy in image classification tasks, it requires a large, supervised dataset that is often expensive to obtain. Unsupervised and semi-supervised learning seek to alleviate this requirement by incorporating unlabelled examples. However, these approaches cannot take advantage of the various sources of *noisy* labels in the modern world, such as images with hashtags in social media or images contained in web-pages retrieved by a textual query. Training algorithms that are robust to label noise are therefore highly attractive for deep learning.

The dominant approach to learning with noisy labels in recent work is to use the predictions of the model itself to reject or modify training examples (e.g. [1, 13, 22, 26, 29, 33, 37, 46]). This is inherently risky due to the ability of deep networks to fit arbitrary labels [47] and it is important to take significant measures against overfitting. Moreover, this paradigm often leads to complicated training procedures, such as maintaining multiple models or alternating between updating the model and updating the training set.

This paper proposes Neighbor Consistency Regularization (NCR) for the specific problem of learning with noisy labels, illustrated in Figure 1. Rather than adopting model predictions as pseudo-labels, NCR introduces an additional consistency loss that encourages each example to have similar predictions to its neighbors. Concretely, the neighbor consistency loss penalizes the divergence of each example’s prediction from a weighted combination of its neighbors’ predictions, with the weights determined by their similarity in feature space. The motivation of NCR is to enable incorrect labels to be improved or at least attenuated by the labels of their neighbors, relying on the assumption that the noise is sufficiently weak or unstructured so as not to overwhelm the correct labels. Compared to the popular approach of bootstrapping the model predictions [33], NCR can be seen as instead bootstrapping the learned feature representation, which may reduce its susceptibility to overfitting and improve its stability at random initialization.

NCR is inspired by label propagation algorithms for semi-supervised learning [15, 19, 49], which seek to transfer labels from supervised examples to neighboring unsupervised examples according to their similarity in feature space. However, whereas label propagation is typically performed in a batch setting over the entire dataset, our method effectively performs label propagation online within mini-batches during stochastic gradient descent. This results in a simple, single-stage training procedure. Moreover, whereas existing methods for label propagation represent *transductive* learning in that they only produce labels for the specific examples which are provided during training, NCR can be understood as an *inductive* form of label propagation in that it produces a model which can later be applied to classify unseen examples.

The key contributions of the paper are as follows.

- We propose Neighbor Consistency Regularization, a novel loss term for deep learning with noisy labels that encourages examples with similar feature representations to have similar predictions.
- We verify empirically that NCR achieves better accuracy than several important baselines at a wide range of noise levels from both synthetic and real distribu-

*Work done at Google.

tions, and is complementary to the popular regularization technique of mixup [48].

- We demonstrate that NCR achieves competitive or state-of-the-art accuracies on datasets evaluating both synthetic (CIFAR-10 and CIFAR-100) and realistic (mini-WebVision [22], mini-ImageNet-Red [18], Clothing1M [43]) noise scenarios.

2. Related work

This section reviews related work with a focus on image classification using deep learning and draws comparisons to our proposed approach. Note that, while there is a significant body of work examining the problem of noisy labels under the assumption that a small, trusted set of clean examples is available, this paper considers the variant where this is not the case.

Regularization. Regularization has an integral role under label noise as it limits the ability of a high-capacity model to fit arbitrary labels. Indeed, regularization alone can provide considerable robustness to label noise and many algorithms couple an effective regularization strategy with an explicit noise-handling technique to further robustness. One particularly effective form of regularization is mixup augmentation [48], which generates additional examples by linearly interpolating between pairs of examples in both the image and label space. It was originally demonstrated to provide robustness to synthetic label corruption on CIFAR-10.

Beyond the choice of regularization, there are numerous possible approaches to handle noisy labels. Early investigation into deep learning with label noise concentrated on either estimation of the label transition matrix [31, 35] or loss functions that are robust to outliers [11, 33, 45]. More recently, numerous methods have been proposed that re-weight or re-label the examples that are believed to have incorrect labels. We briefly review this body of work in the following paragraphs.

Model predictions for noisy labels. Many methods have used the predictions of the model itself during training to either generate pseudo-labels or to identify incorrect examples (or both). For deep learning, this often leverages the phenomenon that correct labels tend to be approximated earlier than incorrect ones [26]. Reed *et al.* [33] first proposed the bootstrapping loss, in which a fixed linear combination of the annotated label and the current prediction are adopted as the regression target for each example. Our method can be considered a bootstrapping approach which uses the learned *similarity* rather than the actual predictions. Extensions of [33] include setting the proportion of annotated and predicted labels adaptively using the model confidence [1, 46] and averaging model weights and predictions over time [26, 29]. Lukasik *et al.* [28] showed that some robustness to noise could be obtained by label smoothing,

which instead interpolates between the annotated label and a uniform distribution. To avoid overfitting, several works have proposed to split the training set in half and train two models, with each used to assess the training examples of the other [13, 22].

Identifying noisy examples using neighbors. Whereas our method, in effect, uses neighbors to modify the supervision of each example, several approaches have rather used neighbors with the aim of identifying and eliminating noisy examples. Both [41] and [12] proposed to identify and re-weight examples with incorrect labels using the local density of examples with the same annotation. Wu *et al.* [42] constructed a k -NN graph and kept only the examples which constituted the core of the largest connected component per class. Bahri *et al.* [2] eliminated examples whose labels did not agree with the prediction of a k -NN classifier. Multi-Objective Interpolation Training (MOIT) [30] identifies mislabelled examples by comparing the predictions of each example to the average prediction of its neighbors and then replaces the least reliable labels with the model’s current predictions. MOIT+ incorporates an additional stage of semi-supervised learning, in which the labels assumed to be noisy, are discarded.

Label propagation. Variants of label propagation [49] have been applied to various computer vision tasks, including retrieval [8, 17], semi-supervised learning [15, 19] and few-shot learning [9, 27, 34]. For semi-supervised learning in particular, [15] used it to obtain labels for unsupervised examples based on their neighbors in feature space. While most methods perform label propagation over large graphs containing many examples in a batch fashion, [27] and [34] considered label propagation within stochastic gradient descent for episodic few-shot learning [39]. This is a meta-learning approach which seeks a suitable feature representation with which to *later* perform label propagation during meta-testing. It is not directly applicable for learning with noisy labels and it does not scale to large sets of examples.

For learning with noisy labels, although neighbors have often been used to *identify* mislabelled examples (see above), few works have considered the use of neighbors to generate pseudo-labels. Iscen *et al.* [16] used graph convolution to propagate labels from a small set of examples with trusted labels to a large set of examples with noisy labels. By comparison, our approach does not require a clean set and is inductive in nature. Perhaps the most similar method to our own is that of [46], in which graph filtering is applied once to the global graph of all examples to refine the predictions of an initial model. In contrast, our consistency loss encourages similar examples to have similar labels throughout training and is much simpler to implement.

Consistency. The idea of constraining a network to produce consistent outputs as a way to leverage unlabelled data has appeared in several previous works. ICT [38] and Mix-

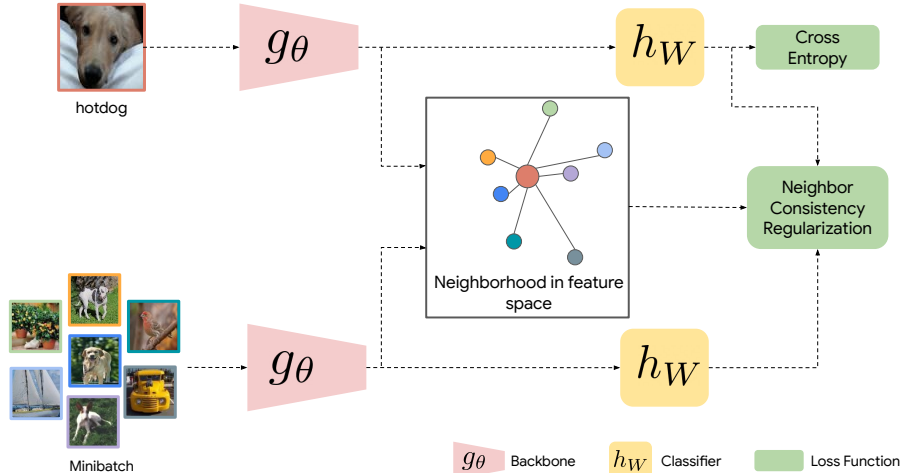


Figure 1. To address the problem of noisy labels in the training set, we propose Neighbor Consistency Regularization. This regularizer encourages examples with similar feature representations to have similar outputs, thus mitigating the impact of training examples with incorrect labels.

match [3] proposed variants of mixup for semi-supervised learning in which predictions replace the labels for unsupervised examples. Xie *et al.* [44] introduced Unsupervised Data Augmentation for semi-supervised image classification, where a model is encouraged to be robust to label-preserving transformations even when labels are not available by minimizing the divergence between predictions for transformed and non-transformed images. Most relevant to our work, [10] used prediction consistency with respect to image transformations for the express purpose of learning with noisy labels. While these forms of consistency are effective regularizers, neighbor consistency offers the ability to transfer supervision directly to mislabelled examples.

3. Preliminaries

We first define our notation and formulate the task of learning with noisy labels. We then describe label propagation, a graph-based semi-supervised learning method designed to work with closed datasets.

Problem formulation. We assume a dataset defined by $X := \{x_1, \dots, x_n\}$. Each example, *e.g.* an image, x_i , has a corresponding true label $\tilde{y}_i \in \mathcal{C}$. In our task, some of the labels, y_i , are noisy: $y_i \neq \tilde{y}_i$ and do not correctly reflect the visual content of the example x_i . During training, we do not know whether y_i is noisy ($y_i \neq \tilde{y}_i$) or clean ($y_i = \tilde{y}_i$). Our goal is to learn a model with the highest accuracy on the true labels, \tilde{y} , although an unknown number of labels in our training set are noisy.

We learn a convolutional neural network for classification. The network, denoted by $f_{\theta, W} : \mathcal{X} \rightarrow \mathbb{R}^c$, takes a dataset example x_i as the input and outputs *logits* for softmax classification. Its two learnable variables θ and W cor-

respond to the feature extractor and classifier, respectively. The feature extractor maps an image x_i to a d -dimensional vector $\mathbf{v}_i := g_\theta(x_i) \in \mathbb{R}^d$. The classifier maps the d -dimensional vector to class scores $\mathbf{z}_i := h_W(\mathbf{v}_i) \in \mathbb{R}^c$. Typically, the network parameters are learned by minimizing a loss function for supervised classification:

$$L_S(X, Y; \theta, W) := \frac{1}{m} \sum_{i=1}^m \ell(\sigma(\mathbf{z}_i), y_i), \quad (1)$$

where X and Y correspond to the set of examples in the mini-batch, $m = |X| = |Y|$ denotes the size of the mini-batch, σ is the softmax function and $\ell(\mathbf{q}, \mathbf{p})$ is the cross-entropy loss function for predictions \mathbf{q} . When the target distribution \mathbf{p} is a single label $y \in \mathcal{C}$, we adopt the shorthand $\ell(\mathbf{q}, y) = \ell(\mathbf{q}, \delta_y)$ for cross-entropy with a one-hot vector δ_y .

Label propagation is a graph-based technique used in semi-supervised learning [49]. We assume that we are given labeled and unlabeled examples in a dataset, and that the dataset is also defined by a graph which is either given or created from the k-NN of each example [8]. This method spreads the label information of each node to the other nodes based on the connectivity in the graph. This process is repeated until a global equilibrium state is achieved. Finally, unlabeled examples are assigned to the class from which they have received the most information. Formally, for a graph of dataset X , represented by an affinity matrix W , where $W_{ij} = \text{similarity}(x_i, x_j)$, Zhou *et al.* [49] show

that label propagation can be computed by minimizing

$$\mathcal{Q}(P) = \frac{1}{2}\mu \sum_{i=1}^n \|P_i - Y_i\|^2 + \frac{1}{2} \sum_{i,j=1}^n W_{ij} \left\| \frac{1}{\sqrt{D_{ii}}} P_i - \frac{1}{\sqrt{D_{jj}}} P_j \right\|^2, \quad (2)$$

where D is the degree matrix (a diagonal matrix with entries $D_{ii} = \sum_j W_{ij}$), $P \in \mathbb{R}^{n \times c}$ is the matrix of classification predictions, $Y \in \mathbb{R}^{n \times c}$ is the matrix of one-hot labels for all examples and μ is a regularization parameter. This objective function comprises two terms: a *fitting constraint*, which encourages the classification of each point to their assigned label, and a *smoothing* term, which encourages the outputs of nearby points in the graph to be similar.

One of the main limitations of label propagation is its *transductive* property. In transductive learning, the goal is to classify *seen* unlabeled examples. This is different to *inductive* learning, which learns a generic classifier to classify any unseen data. To apply label propagation on new test examples, a new graph W needs to be constructed each time a test example is seen. This makes it inefficient in practice.

Another requirement for label propagation is that the feature space needs to be fixed to compute the affinity matrix W . This requires the feature extractor to be learned beforehand, potentially from the noisy data. Existing work [15] has tried to overcome this issue by alternating between optimizing the feature space, and performing label propagation. However, this does not directly enforce smoothness, as the optimization of two components are done separately.

Our goal is to overcome the limitations of label propagation by 1) adapting it to an inductive setting 2) applying the smoothness constraint directly during optimization. In Section 4, we propose a simple and efficient approach which generalizes label propagation by enforcing smoothness in the form of a regularizer. As a result, we avoid constructing an explicit graph to propagate the information, and inference can be performed on any unseen test example.

4. Method

We now present our method *Neighbor Consistency Regularization* and compare it to classical label propagation. We then highlight its relationship to similar, online techniques.

4.1. Neighbor Consistency Regularization

When learning with noisy labels, the network is prone to overfit, or *memorize*, the mapping from x_i to a noisy label y_i for the training data [26]. This behavior typically results in a non-optimal classification performance in a clean evaluation set, as the network does not generalize well.

To overcome this issue, we propose *Neighbor Consistency Regularization* (NCR). Our main assumption is that the over-fitting occurs less dramatically before the classifier h_W . This is supported by MOIT [30], which shows that feature representations are robust enough to discriminate between noisy and clean examples when training a network. With that assumption, we can design a smoothness constraint similar to label propagation (2) when training the network. The overview of our method is shown in Figure 1.

Let us define the *similarity* between two examples by the cosine similarity of their feature representations, *i.e.* $s_{i,j} = \cos(\mathbf{v}_i, \mathbf{v}_j) = \mathbf{v}_i^T \mathbf{v}_j / (\|\mathbf{v}_i\| \|\mathbf{v}_j\|)$. Note that the feature representations contain non-negative values when obtained after a ReLU non-linearity, and therefore the cosine similarity is bounded in the interval $[0, 1]$. Our goal is to enforce *neighbor consistency regularization* by leveraging the structure of the feature space produced by g_θ to enhance the classifier h_W . More specifically, $h_W(\mathbf{v}_i)$ and $h_W(\mathbf{v}_j)$ should behave similarly if $s_{i,j}$ is high, regardless of their labels y_i and y_j . This would prevent the network from over-fitting to an incorrect mapping between an example x_i and a label y_i , if either (or both) y_i and y_j are noisy.

To enforce NCR, we design an objective function which minimizes the distance between logits \mathbf{z}_i and \mathbf{z}_j , if the corresponding feature representations \mathbf{v}_i and \mathbf{v}_j are similar:

$$L_{\text{NCR}}(X, Y; \theta, W) := \frac{1}{m} \sum_{i=1}^m D_{\text{KL}} \left(\sigma(\mathbf{z}_i/T) \left\| \sum_{j \in \text{NN}_k(\mathbf{v}_i)} \frac{s_{i,j}}{\sum_k s_{i,k}} \cdot \sigma(\mathbf{z}_j/T) \right. \right), \quad (3)$$

where D_{KL} is the KL-divergence loss to measure the difference between two distributions, T is the temperature and $\text{NN}_k(\mathbf{v}_i)$ denotes the set of k nearest neighbors of i in the feature space. We set $T = 2$ throughout our experiments. We normalize the similarity values so that the second term of the KL-divergence loss remains a probability distribution. We set the self-similarity $s_{i,i} = 0$ so that it does not dominate the normalized similarity. Gradients will be back-propagated to all inputs.

The objective (3) ensures that the output of x_i will be consistent with the output of its neighbors regardless of its potentially noisy label y_i . We combine it with the supervised classification loss function (1) to obtain the final objective to be minimized during the training:

$$L(X, Y; \theta, W) := (1 - \alpha) \cdot L_S(X, Y; \theta, W) + \alpha \cdot L_{\text{NCR}}(X, Y; \theta, W), \quad (4)$$

where the hyper-parameter $\alpha \in [0, 1]$ controls the impact of the each loss term. Similar to label propagation, the final loss objective (4) has two terms. The first term is the classification loss term L_S . This is analogous to the fitting

constraint in (2). The second term is the NCR loss L_{NCR} , which is also similar to the *smoothness* constraint in (2). We find that it sometimes helps to train the network with $\alpha = 0$ for several epochs (denoted by e in our experiments) before enabling the NCR term.

However, the main difference between label propagation and our method is that label propagation applies smoothness based on the graph edges W_{ij} computed over the entire dataset. On the other hand, our method is online, and does not require a global graph W . We enforce the NCR through the local neighborhood as the feature space is being learned. As a result, our method does not require a learned feature representation with noisy examples. It enriches the learned feature representation by reducing the negative impact of noisy examples.

Compared to standard training, NCR incurs an additional computational cost of order $O(m^2(d + c))$ where m is the batch size, d is the feature dimension and c is the number of classes. This arises in the computation of the similarity values and weighted predictions in (3). However, this operation is relatively fast to compute for moderate values of m because it is a dense matrix multiplication, for which modern GPUs are optimized.

4.2. Relation to other methods

Bootstrapping introduces an additional loss that adopts the model’s own predictions as labels [33]. Its motivation is to discourage the model from overfitting to examples which are difficult to fit. The overall loss combines the supervised loss and the bootstrap loss in fixed proportion:

$$L_{\text{B}}(X, Y; \theta, W) := \frac{1}{m} \sum_{i=1}^m (1 - \alpha) \cdot \ell(\sigma(\mathbf{z}_i), y_i) + \alpha \cdot \ell(\sigma(\mathbf{z}_i), \sigma_{\text{B}}(\mathbf{z}_i)), \quad (5)$$

where σ_{B} is the bootstrap activation function, which may be argmax or softmax (with optional temperature). NCR can be understood as bootstrapping from the neighborhood structure induced by the representation rather than from the model’s predictions. This eliminates the dependency on the classifier parameters W , which may be particularly advantageous as it has been shown that linear models can fit random labels given a sufficiently high-dimensional representation [26].

Label smoothing is a regularization method [36] that mixes the ground truth labels with a uniform distribution. It has been shown capable of denoising corrupted labels [28]. Under label smoothing, the supervised classification loss function becomes:

$$L_{\text{LS}}(X, Y; \theta, W) := \frac{1}{m} \sum_{i=1}^m (1 - \alpha) \cdot \ell(\sigma(\mathbf{z}_i), y_i) + \alpha \cdot \ell(\sigma(\mathbf{z}_i), \frac{1}{c} \mathbf{1}). \quad (6)$$

Note that the linear combination of losses is equivalent to a linear combination of labels due to the linearity of $\ell(\mathbf{q}, \mathbf{p}) = -\mathbf{p}^T \log \mathbf{q}$ with respect to \mathbf{p} . Our method can be considered a modified version of label smoothing where the uniform distribution is replaced with a distribution defined by the neighboring examples. It would collapse to label smoothing if the neighbors were random or if a high temperature T were used.

Mixup [48] bears some resemblance to our method in that it takes a convex combination of labels. However, mixup uses this combination as the regression target for a novel example obtained as the convex combination of the inputs, whereas NCR uses it as the target for an existing example.

5. Experiments

5.1. Experimental setup

We first perform ablation studies on datasets with synthetic noise, where the noise level can be varied, before considering datasets with organic noise, where the noise level is fixed and unknown. For experiments with synthetic noise, we use the standard variants of CIFAR-10 and -100 [20] as well as mini-ImageNet with “Blue” and “Red” noise [18]. The CIFAR and mini-ImageNet-Blue datasets are contaminated with uniform label noise, whereas mini-ImageNet-Red is constructed by replacing some examples in each class with false positives from an image search engine, representing more realistic noise. For experiments with organic noise, we use mini-WebVision [22], WebVision [25] and Clothing1M [43]. All datasets include a clean validation and/or final evaluation set. Exhaustive implementation details and hyperparameters are also included in the supplementary.

5.2. Ablation study

We first study the effect of the key hyperparameters of NCR at different noise levels using the CIFAR-10 validation set. Specifically, we investigate the impact of α , which controls the strength of the NCR term in (4), the number of neighbors k and the number of initialization epochs e . We first set the number of neighbors $k = m$ (the batch size) and the number of initialization epochs to zero while varying α . Subsequently, we select the optimal α for each noise level and vary k . Finally, we adopt the optimal k and vary e .

Figure 2 shows the validation accuracy for different noise ratios. The performance remains relatively stable over different hyperparameters for 0% and 20% noise ratios. It is optimal to set α to be high (0.9) for any non-zero noise ratio, indicating that the greater influence of NCR benefits these settings. Similarly, smaller k (e.g. $k = 10$) leads to higher accuracy when the noise ratio is above 0%. We also observe that $e = 0$ performs better for higher noise ratios (i.e. 40% and 80%). This shows that the NCR needs to be




Figure 2. **Ablation study.** Impact of hyperparameters α , k and e are evaluated on the CIFAR-10 validation set with ResNet-18.

enabled early in training to prevent the network from memorizing the noisy labels when the noise ratio is high.

5.3. Baseline comparison

We now compare NCR against the baselines defined in Section 4.2. The results are reported on the official validation set of mini-ImageNet- $\{\text{Red, Blue}\}$ datasets. We run each experiment five times and report the mean accuracy at the completion of training. We do not report the peak validation accuracy attained during training as the conclusions may be less likely to generalize to an unseen test set.

Table 1 shows the final accuracy for each method on the mini-ImageNet- $\{\text{Red, Blue}\}$ datasets across different noise splits. When compared with the *standard* baseline (eq. (1)), our method significantly improves the performance, up to 17.5% across all noise ratios. Furthermore, we show that our method is compatible with some of the existing baselines. Combining mixup with our method leads to further improvements in almost all scenarios. We observe that NCR improves the accuracy of the method even at 0% noise. This suggests that it has a general regularization effect. However, the improvement in accuracy is much more pronounced in the training sets which contain label noise.

Figure 3 presents further evidence that NCR inhibits memorization of noisy labels. After the training is complete, we perform one last forward-pass and obtain the confidence p that the model assigns to the annotated label for each training example. The top row shows that the baseline model overfits to the noisy labels, resulting in $p = 1$ for both clean and noisy images. On the other hand, NCR avoids overfitting and assigns low confidence $p = 0$ to most of the noisy training labels. Some of the noisy examples are still classified as their assigned label on mini-Imagenet-Red. This is likely due to the dataset containing realistic and correlated noise: the mislabelled class and the true class often have visual patterns in common.

For the synthetic noise (mini-Imagenet-Blue), NCR separates the clean and noisy examples up to 40% noise ratio. However, the model underfits with 80% noise, resulting in p being close to 0 for both clean and noisy examples. This results in a small improvement for NCR, see mini-ImageNet-Blue with 80% noise in Table 1.

Table 1 also includes an *oracle* that is obtained by ex-

cluding the mislabelled examples from the training set, reducing its size by 20%, 40% or 80% accordingly. Under *realistic* noise (mini-ImageNet-Red), the results show that NCR outperforms the oracle across all noise ratios when using mixup. While this may be surprising, it can be explained by the observation that the noisy examples for each class are often visually similar to the clean examples, and thus still contain some useful information. However, the performance of NCR is significantly less than the oracle under synthetic noise (mini-ImageNet-Blue), where the noisy examples are often entirely dissimilar to the clean examples.

5.4. Effect on feature embeddings

Using the datasets with known noise, we can compare the feature similarity of training examples that are correctly or incorrectly labelled as belonging to the same or different classes. In the ideal case, the distributions of within-class and between-class similarities for clean examples would have zero overlap, and would be identical to the true within-class and between-class similarities for mislabelled examples. Figure 4 presents the similarity distributions for mini-ImageNet-Blue and -Red. (Note that, for mini-ImageNet-Red, the set of clean examples is known but the true class of the mislabelled examples is not known, so we cannot obtain the true within-class similarity for mislabelled examples.) We compare the distributions for a baseline model and a model trained with NCR. While the distributions for the baseline model overlap, they are not identical, showing that the feature similarities contain some signal for NCR to take advantage of. Training with NCR is shown to achieve greater separation of the classes in feature space, which must be due to the backpropagation of gradients through the feature similarities in (3).

5.5. State-of-the-art comparison

In addition to Table 1, which compares to prior art on mini-ImageNet-Red, we compare NCR against the state-of-the-art on mini-WebVision, WebVision and Clothing1M (which evaluate realistic noise) in Table 2. Finally, we also compare to prior work on CIFAR-10 and -100 (which represents synthetic noise) in Table 3.

In Table 2, we also present a Data Augmentation (DA) variant of NCR, where we generate a second view of each

Table 1. **Baseline and oracle comparison.** Classification accuracy is reported on the mini-ImageNet-{Blue, Red} datasets with the ResNet-18 architecture for each individual noise ratio (0%, 20%, 40%, 80%). We present the mean accuracy and standard deviation from five trials. The oracle model is trained on only the known, clean examples in the training set using a cross-entropy loss.

Method	mini-ImageNet-Blue				mini-ImageNet-Red			
	0%	20%	40%	80%	0%	20%	40%	80%
BASELINES								
Standard	65.8±0.4	49.5±0.4	36.6±0.5	13.1±1.0	63.5±0.5	55.3±0.9	49.5±0.7	36.4±0.4
Mixup	67.4±0.4	60.1±0.2	51.6±0.8	21.0±0.5	65.5±0.5	61.6±0.5	57.2±0.6	43.7±0.3
Bootstrap	66.4±0.4	54.4±0.5	44.8±0.5	2.9±0.3	64.3±0.3	56.2±0.2	51.3±0.6	38.2±0.3
Bootstrap + Mixup	67.5±0.3	61.9±0.4	51.0±0.7	1.3±0.1	65.9±0.4	62.7±0.2	58.3±0.5	43.5±0.6
Label smoothing	67.5±0.8	60.2±0.5	50.2±0.4	20.9±0.8	65.7±0.5	59.7±0.4	54.0±0.6	39.7±0.5
Label smoothing + Mixup	68.6±0.3	63.3±1.0	57.1±0.2	14.4±0.3	66.9±0.2	63.4±0.4	59.2±0.4	45.5±0.7
OURS								
Ours: NCR	66.5±0.2	61.7±0.3	54.1±0.4	20.7±0.5	64.0±0.4	60.9±0.3	56.1±0.7	40.9±0.2
Ours: NCR + Mixup	67.9±0.6	64.3±0.1	59.2±0.6	14.2±0.4	66.3±0.5	64.6±0.6	60.4±0.3	45.4±0.4
OTHER WORKS								
D-Mix [22]	-	-	-	-	55.8	50.3	50.9	35.4
ELR [26]	-	-	-	-	57.4	58.1	50.6	41.7
MOIT [30]	-	-	-	-	64.7	63.1	60.8	45.9
ORACLE: CLEAN SUBSET								
Standard	65.8±0.4	63.9±0.5	60.6±0.4	45.4±0.8	63.5±0.5	61.7±0.1	58.4±0.3	41.5±0.5
Mixup	67.4±0.4	64.2±0.5	61.5±0.3	46.9±0.8	65.5±0.5	63.1±0.6	59.7±0.7	43.6±0.4

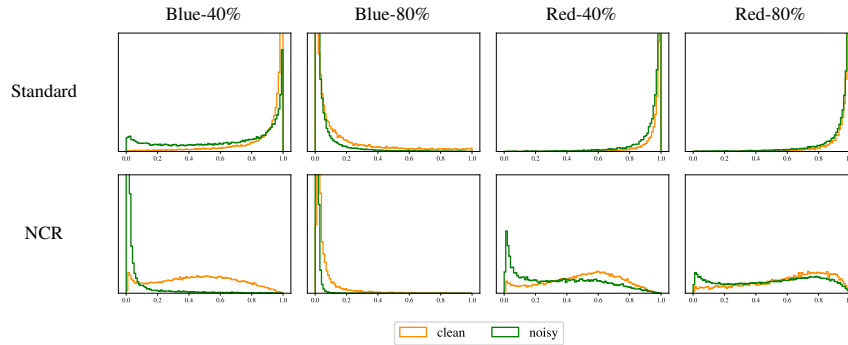


Figure 3. **Predicted confidence of training examples.** Whereas the baseline assigns similar confidence scores to clean and mislabelled examples (usually all high), NCR more often assigns lower confidence to mislabelled examples and higher confidence to correct examples.

example in the minibatch by applying random color jittering, motivated by [14]. Table 2 shows that NCR achieves state-of-the-art accuracy on datasets with realistic noise such as mini-WebVision, improving by 1.2% on the best previous result. NCR outperforms other methods on WebVision when ResNet50 is used, and is competitive against methods using Inception-ResNetV2, a much stronger architecture. Our method is also competitive on Clothing1M, where it is only 0.2% less than the state-of-the-art.

Note that NCR does not involve additional steps, such as dividing the dataset into multiple splits, learning multiple models, applying semi-supervised learning, or a second-stage with additional training as done in the other works [22, 26, 30] we compare to. While NCR is orthogonal to these strategies, we show that we achieve state of the art

results with minimal additional processing (*e.g.* data augmentation). Furthermore, NCR achieves higher performance compared to GJS [10] which applies a similar consistency regularization only on different augmentation of each example. This confirms that the neighbor consistency brings further improvements on top of augmentation consistency.

To show that our method is compatible with the existing approaches, we combine NCR with ELR [26] for CIFAR-10 and -100 comparisons in Table 3. We reproduce all the comparisons (by running the public code), except for MOIT [30] whose results are taken from the paper. Our method consistently improves over ELR, achieving state-of-the-art results in almost all noise ratios in CIFAR. For the noise ratios in which NCR does not obtain the highest result, it is typically the second-highest. Note that we fix the same hyperparam-

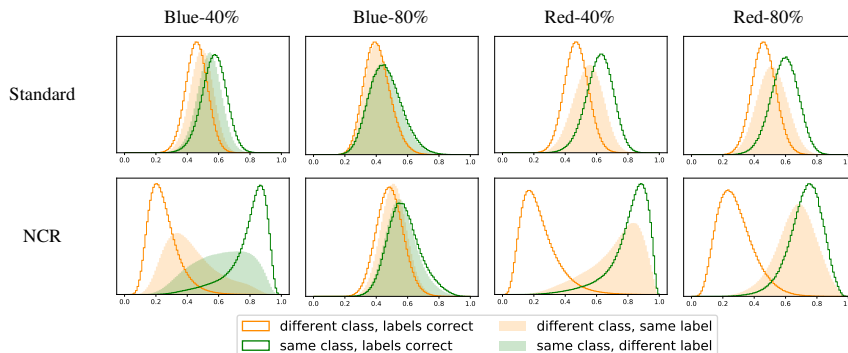


Figure 4. **Similarity distributions.** We compare the distribution of cosine similarities for training examples in mini-ImageNet that are correctly and incorrectly labelled as the same class or different classes. For mini-ImageNet-Blue, the features learned using NCR achieve significantly better class separation with 40% noise (or less, not pictured). For the more realistic mini-ImageNet-Red, NCR still achieves better separation of the clean examples but fails to separate examples that are incorrectly labelled as the same class.

Table 2. **State-of-the-art comparison with realistic noise.** We compare NCR and our baselines to other methods on mini-WebVision, WebVision and Clothing1M. All results use ResNet-50, except for those marked by \dagger which use Inception-ResNetV2.

	mini-WebVision	WebVision	Clothing1M
Standard	75.8	74.9	71.7
Mixup	77.2	75.5	72.2
Ours: NCR	77.1	75.7	74.4
Ours: NCR+Mixup	79.4	75.4	74.5
Ours: NCR+Mixup+DA	80.5	76.8	74.6
MLNT; 3 iter. [23]	–	–	73.5
CleanNet [21]	–	–	74.7
\mathcal{L}_{DMI} [40]	–	–	72.5
LongReMix [6]	–	–	73.0
ELR [26]	76.3 \dagger	–	–
ELR+ [26]	77.8 \dagger	–	74.8
DMix [22]	76.3	–	74.8
GJS [10]	79.3	–	–
MoPro [24]	–	73.9	–
MILe [32]	–	75.2	–
Heteroscedastic [5]	–	76.6 \dagger	–
CurrNet [12]	–	79.3 \dagger	–

Table 3. **State-of-the-art comparison with synthetic noise on CIFAR.** A-40% refers to 40% asymmetric noise. All of the other columns refer to symmetric noise.

	CIFAR-10						CIFAR-100					
	20%	40%	50%	80%	90%	A-40%	20%	40%	50%	80%	90%	A-40%
Standard	83.9	68.3	58.5	25.9	17.3	77.3	61.5	46.2	37.4	10.4	4.1	43.9
MOIT+ [30]	94.1	92.0	–	75.8	–	93.2	75.9	67.4	–	51.4	–	74.0
D-Mix [22]	95.1	94.2	93.6	91.4	74.5	91.8	76.7	74.6	73.1	57.1	29.7	72.1
ELR+ [26]	94.9	94.4	93.9	90.9	74.5	88.9	76.3	74.0	72.0	57.2	30.9	75.8
Ours+ [26]	95.2	94.5	94.3	91.6	75.1	90.7	76.6	74.2	72.5	58.0	30.8	76.3

eters across all noise ratios, unlike Divide-Mix [22], which is a more realistic scenario in practice.

6. Conclusion

This work introduced Neighborhood Consistency Regularization and demonstrated that it is an effective strategy for deep learning with label noise. While our approach draws inspiration from multi-stage training procedures for semi-supervised learning that employ transductive label propagation, it consists of a comparatively simple training procedure, requiring only that an extra loss be added to the objective which is optimized in stochastic gradient descent. The efficacy of NCR is emphasized by the fact that it achieved state-of-the-art results under both synthetic (CIFAR-10 and -100) and realistic (mini-WebVision) noise scenarios.

Limitations and future work. A limitation of NCR is that our proposed loss assumes that it has access to an adequate feature representation of the training data. We overcome this limitation in practice by first training the network for e epochs before applying the NCR loss, but future work is to remove this additional training hyperparameter.

Promising directions for future research including coupling NCR with a technique to rejecting out-of-distribution examples during training, as employed by other approaches in the literature [22, 30]. We also note that NCR could be applied to related problems, such as semi-supervised learning, as a regularization term.

Broader Impact. Our proposed method is suited to learning from noisy data, as could be obtained by automatic scraping of the internet (illustrated by our experiments on mini-ImageNet-Red and mini-WebVision). Data collected in this fashion may contain bias [4, 7] and a method which is able to more effectively learn from such data may inadvertently cause these biases to be amplified. Furthermore, when using training data automatically scraped from the web, it is possible that the data is being used for a purpose to which the original owner did not consent, potentially infringing on their privacy.

References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel O’Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. In *ICML*, pages 312–321. PMLR, 2019. 1, 2
- [2] Dara Bahri, Heinrich Jiang, and Maya Gupta. Deep k-NN for noisy labels. In *ICML*, pages 540–550. PMLR, 2020. 2
- [3] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. MixMatch: A holistic approach to semi-supervised learning. In *NeurIPS*, 2019. 3
- [4] Abeba Birhane and Vinay Uday Prabhu. Large image datasets: A pyrrhic win for computer vision? In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021. 8
- [5] Mark Collier, Basil Mustafa, Efi Kokiopoulou, Rodolphe Jenatton, and Jesse Berent. Correlated input-dependent label noise in large-scale image classification. In *CVPR*, 2021. 8
- [6] Filipe R Cordeiro, Ragav Sachdeva, Vasileios Belagiannis, Ian Reid, and Gustavo Carneiro. Longremix: Robust learning with high confidence samples in a noisy label environment. *arXiv preprint arXiv:2103.04173*, 2021. 8
- [7] Terrance de Vries, Ishan Misra, Changhan Wang, and Laurens van der Maaten. Does object recognition work for everyone? In *CVPR Workshops*, pages 52–59, 2019. 8
- [8] Michael Donoser and Horst Bischof. Diffusion processes for retrieval revisited. In *CVPR*, 2013. 2, 3
- [9] Matthijs Douze, Arthur Szlam, Bharath Hariharan, and Hervé Jégou. Low-shot learning with large-scale diffusion. In *CVPR*, 2018. 2
- [10] Erik Englesson and Hossein Azizpour. Generalized Jensen-Shannon divergence loss for learning with noisy labels. *arXiv preprint arXiv:2105.04522*, 2021. 3, 7, 8
- [11] Aritra Ghosh, Himanshu Kumar, and PS Sastry. Robust loss functions under label noise for deep neural networks. In *AAAI*, volume 31, 2017. 2
- [12] Sheng Guo, Weilin Huang, Haozhi Zhang, Chenfan Zhuang, Dengke Dong, Matthew R Scott, and Dinglong Huang. CurriculumNet: Weakly supervised learning from large-scale web images. In *ECCV*, pages 135–150, 2018. 2, 8
- [13] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor W. Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, NIPS’18, page 8536–8546, Red Hook, NY, USA, 2018. Curran Associates Inc. 1, 2
- [14] Elad Hoffer, Tal Ben-Nun, Itay Hubara, Niv Giladi, Torsten Hoefer, and Daniel Soudry. Augment your batch: Improving generalization through instance repetition. In *CVPR*, 2020. 7
- [15] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondrej Chum. Label propagation for deep semi-supervised learning. In *CVPR*, pages 5070–5079, 2019. 1, 2, 4
- [16] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, Ondřej Chum, and Cordelia Schmid. Graph convolutional networks for learning with few clean and many noisy labels. In *European Conference on Computer Vision*, pages 286–302. Springer, 2020. 2
- [17] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, Teddy Furon, and Ondrej Chum. Efficient diffusion on region manifolds: Recovering small objects with compact cnn representations. In *CVPR*, 2017. 2
- [18] 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. 2, 5
- [19] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017. 1, 2
- [20] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 5
- [21] Kuang-Huei Lee, Xiaodong He, Lei Zhang, and Linjun Yang. CleanNet: Transfer learning for scalable image classifier training with label noise. In *CVPR*, 2018. 8
- [22] Junnan Li, Richard Socher, and Steven C H Hoi. DivideMix: Learning with noisy labels as semi-supervised learning. *ICLR*, 2020. 1, 2, 5, 7, 8
- [23] Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S Kankanhalli. Learning to learn from noisy labeled data. In *CVPR*, pages 5051–5059, 2019. 8
- [24] Junnan Li, Caiming Xiong, and Steven CH Hoi. Mopro: Webly supervised learning with momentum prototypes. *arXiv preprint arXiv:2009.07995*, 2020. 8
- [25] Wen Li, Limin Wang, Wei Li, Eirikur Agustsson, and Luc Van Gool. Webvision database: Visual learning and understanding from web data. *arXiv preprint arXiv:1708.02862*, 2017. 5
- [26] Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. *NeurIPS*, 2020. 1, 2, 4, 5, 7, 8
- [27] Yanbin Liu, Juho Lee, Minseop Park, Saehoon Kim, Eunho Yang, Sung Ju Hwang, and Yi Yang. Learning to propagate labels: Transductive propagation network for few-shot learning. In *ICLR*, 2019. 2
- [28] Michal Lukasik, Srinadh Bhojanapalli, Aditya Menon, and Sanjiv Kumar. Does label smoothing mitigate label noise? In *ICML*, 2020. 2, 5
- [29] 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. *ICLR*, 2020. 1, 2
- [30] Diego Ortego, Eric Arazo, Paul Albert, Noel E O’Connor, and Kevin McGuinness. Multi-objective interpolation training for robustness to label noise. *arXiv preprint arXiv:2012.04462*, 2020. 2, 4, 7, 8
- [31] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *CVPR*, 2017. 2
- [32] Sai Rajeswar, Pau Rodriguez, Soumye Singhal, David Vazquez, and Aaron Courville. Multi-label iterated learning for image classification with label ambiguity. *arXiv preprint arXiv:2111.12172*, 2021. 8

- [33] Scott Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. Training deep neural networks on noisy labels with bootstrapping. *ICLR*, 2015. [1](#), [2](#), [5](#)
- [34] Pau Rodriguez, Issam Laradji, Alexandre Drouin, and Alexandre Lacoste. Embedding propagation: Smoother manifold for few-shot classification. In *ECCV*, 2020. [2](#)
- [35] Sainbayar Sukhbaatar, Joan Bruna, Manohar Paluri, Lubomir Bourdev, and Rob Fergus. Training convolutional networks with noisy labels. *arXiv preprint arXiv:1406.2080*, 2014. [2](#)
- [36] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, pages 2818–2826, 2016. [5](#)
- [37] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *CVPR*, pages 5552–5560, 2018. [1](#)
- [38] Vikas Verma, Alex Lamb, Juho Kannala, Yoshua Bengio, and David Lopez-Paz. Interpolation consistency training for semi-supervised learning. In *IJCAI*, pages 3635–3641, 7 2019. [2](#)
- [39] Oriol Vinyals, Charles Blundell, Tim Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *NeurIPS*, 2016. [2](#)
- [40] Xinshao Wang, Elyor Kodirov, Yang Hua, and Neil M Robertson. Derivative manipulation for general example weighting. In *arXiv preprint arXiv:1905.11233*, 2019. [8](#)
- [41] Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, and Shu-Tao Xia. Iterative learning with open-set noisy labels. In *CVPR*, 2018. [2](#)
- [42] Pengxiang Wu, Songzhu Zheng, Mayank Goswami, Dimitris Metaxas, and Chao Chen. A topological filter for learning with label noise. *NeurIPS*, 33, 2020. [2](#)
- [43] Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *CVPR*, 2015. [2](#), [5](#)
- [44] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. *NeurIPS*, 33, 2020. [3](#)
- [45] Yilun Xu, Peng Cao, Yuqing Kong, and Yizhou Wang. LDMI: A novel information-theoretic loss function for training deep nets robust to label noise. In *NeurIPS*, volume 32, 2019. [2](#)
- [46] Jingkan Yang, Litong Feng, Weirong Chen, Xiaopeng Yan, Huabin Zheng, Ping Luo, and Wayne Zhang. Webly supervised image classification with self-contained confidence. *arXiv preprint arXiv:2008.11894*, 2020. [1](#), [2](#)
- [47] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016. [1](#)
- [48] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *ICLR*, 2018. [2](#), [5](#)
- [49] Dengyong Zhou, Olivier Bousquet, Thomas Navin Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. In *NeurIPS*, 2003. [1](#), [2](#), [3](#)