

A Framework using Contrastive Learning for Classification with Noisy Labels*

Madalina Ciortan^{§1}, Romain Dupuis^{§1}, and Thomas Peel¹

¹EURO NOVA BE

Abstract

We propose a framework using contrastive learning as a pre-training task to perform image classification in the presence of noisy labels. Recent strategies such as pseudo-labelling, sample selection with Gaussian Mixture models, weighted supervised contrastive learning have been combined into a fine-tuning phase following the pre-training. This paper provides an extensive empirical study showing that a preliminary contrastive learning step brings a significant gain in performance when using different loss functions: non robust, robust, and early-learning regularized. Our experiments performed on standard benchmarks and real-world datasets demonstrate that: i) the contrastive pre-training increases the robustness of any loss function to noisy labels and ii) the additional fine-tuning phase can further improve accuracy, but at the cost of additional complexity.

1 Introduction

Collecting large and well-annotated datasets for image classification tasks represents a challenge as human quality annotations are expensive and time-consuming. Alternative methods exist, such as web crawlers [24]. Nevertheless, these methods generate noisy labels decreasing the performance of deep neural networks. They tend to overfit to noisy labels due to their high capacity [41]. That is why developing efficient noisy-label learning (NLL) techniques is of great importance.

Various strategies have been proposed to deal with NLL: i) Noise transition matrix [8, 30, 38] estimates the noise probability and corrects the loss function, ii) a clean subset can help to avoid overfitting [13], iii) samples selection identifies true-labeled samples [9, 14, 19], and iv) robust loss functions solve the classification problem only by adapting the loss function to be less sensitive to noisy labels [23, 34, 43]. Methods also combine other strategies (eg. ELR+ [19], DivideMix [22]): two networks, semi-supervised learning,

label correction, or mixup. They show the most promising results but lead to a large number of hyperparameters. That is why we explore improvement strategies for robust loss functions. They are simpler to integrate and faster to train, but as illustrated in Figure 1, they tend to overfit and have lower performance for high noise ratios.




Figure 1: Overfitting patterns on 3 losses when training classifiers with noisy labels. The plot depicts top-1 accuracy for ResNet18 trained on the CIFAR-100 dataset with a symmetric noise (80%) for three losses: Cross Entropy (CE), Normalized Focal Loss + Reverse CE (NFL+RCE), and Early Learning Regularization (ELR).

Meanwhile, new self-supervised learning algorithms for image representations have been recently developed [4, 11]. Such algorithms extract representation (or features) in unsupervised settings. These representations can then be used for downstream tasks such as classification. Methods based on contrastive learning compete with fully supervised learning while fine-tuning only on a small fraction of all available labels. Therefore, using contrastive learning for NLL appears as promising. In this work, contrastive learning aims to pre-train the classifier to improve its robustness.

The key contributions of this work are:

- A framework increasing robustness of any loss function to noisy labels by adding a contrastive pre-training task.
- The adaptation of the supervised contrastive loss to use

*Submitted to CVPR L2ID 2021 Workshop - Under Review

[§]Equal contribution

sample weight values, representing the probability of correctness for each sample in the training set

- An extensive empirical study identifying and benchmarking additional state of the art strategies to boost the performance of pre-trained models: pseudo-labeling, sample selection with GMM, and weighted supervised contrastive learning.

2 Related works

Existing approaches dealing with NLL and contrastive learning in computer vision are briefly reviewed. Extra details can be found in Le-Khac et al. [18], Song et al. [33].

2.1 Noise tolerant classification

Sample Selection: This method identifies noisy and clean samples within the training data. Several strategies leverage the interactions between multiple networks to identify the probably correct labels [9, 14, 19]. Recent works [1, 32] exploit the small loss trick to identify clean and noisy samples by considering a certain number of small-loss training samples as true-labeled samples. This approach can be justified by the memorization effect: deep neural networks first fit the training data with clean labels during a so-called early learning phase, before overfitting the noisy samples during the memorization phase [2, 22].

Robust Loss Function: Commonly used loss functions, such as Cross Entropy (CE) or Focal Loss, are not robust to noisy labels. Therefore, new loss functions have been designed. Such robust loss functions can be easily incorporated into existing pipelines to improve performance regarding noisy labels. The symmetric cross entropy [34] has been proposed by adding a reverse CE loss to the initial CE. This combination improves the accuracy of the model compared to classical loss functions. Ma et al. [23] show theoretically that normalization can convert classical loss functions into loss functions robust to noise labels. The combination of two robust loss functions can also improve robustness. However, the performance of normalized loss functions remains quite low for high noise rates as illustrated in Figure 1.

Semi-supervised: Semi-supervised approaches deal with both labeled and unlabeled data. Recent works [19, 27, 35] combine sample selection with semi-supervised methods: the possibly noisy samples are treated as unlabeled and the possibly clean samples are treated as labeled. Semi-supervised approaches show competitive results. However, they use several hyperparameters that can be sensitive to changes in data or noise type [28, 33].

Contrastive learning: recent developments in self-supervised and contrastive learning [20, 28, 42] inspire new

approaches in NLL. Li et al. [20] employed features learned by contrastive learning to detect out-of-distribution samples.

2.2 Contrastive learning for vision data

Contrastive learning extracts features by comparing different samples. The central idea is to bring different instances of the same input image closer and spread instances from different images apart. The inputs are usually divided into positive (similar) and negative pairs (dissimilar). Data augmentation is typically used to create positive pairs. Frameworks have been recently developed, such as CPCv2 [12], SimCLR [4], Moco [11]. Once the self-supervised model is trained, the extracted representations can be used for downstream tasks, such as classification in this work.

Chen et al. [4] demonstrate that large sets of negatives (and large batches) are crucial in learning good representations. However, large batches are limited by GPU memory. Maintaining a memory bank accumulating a large number of negative representations is an elegant solution decoupling the batch size from the number of negatives [25]. Nevertheless, the representations get outdated in a few iterations. The Momentum Encoder [11] addresses the issues by generating a dynamic memory queue of representations. Other strategies aim at getting more meaningful negative samples to reduce the memory/batch size [15].

3 Preliminaries

Let $D = \{(\mathbf{x}_i, \bar{y}_i)\}_{i=1..n}$, $\mathbf{x}_i \in \mathbb{R}^d$, $\bar{y}_i \in \{1, \dots, K\}$ denote a noisy input dataset with an unknown number of samples incorrectly labelled. The associated true and unobservable labels are written y_i . The images \mathbf{x}_i are of size d and the classification problem has K classes. The goal is to train a deep neural network (DNN) f . Using a robust loss function for training consists of minimizing the empirical risk defined by robust loss functions in order to find the set of optimal parameters θ . The one-hot encoding of the label is denoted by the distribution $q(k|\mathbf{x})$ for a sample \mathbf{x} and a class k , such as $q(y_i|\mathbf{x}_i) = 1$ and $q(k \neq y_i|\mathbf{x}_i) = 0$, $\forall i \in \{1, \dots, n\}$. The probability vector of f is defined by the softmax function $p(k|\mathbf{x}) = \frac{e^{z_k}}{\sum_{j=1}^K e^{z_j}}$ where z_k denotes the logits output with respect to class k .

3.1 Classification with robust loss functions

Our method employs noise-robust losses to train the classifier in the presence of noisy labels. Such losses improve the classification accuracy compared to the commonly used CE, as illustrated in Figure 1. In this section, the general empirical risk for a given mini-batch is defined by

$L = \sum_{i=1}^N \mathcal{L}(f(x_i), \bar{y}_i) = \sum_{i=1}^N l_i$. The term l_i is modified by each loss function.

The classical CE is used as a baseline loss function not robust to noisy labels [7] and is defined as:

$$l_{ce} = - \sum_{k=1}^K q(k|\mathbf{x}_i) \log(p(k|\mathbf{x}_i)). \quad (1)$$

As presented in section 2, Ma et al. [23] introduce robust loss functions called Active Passive Losses that do not suffer from underfitting. We investigate the combination between the Normalized Focal Loss (NFL) and the Reversed Cross Entropy (RCE) called NFL+RCE. It shows promising results on various benchmarks. The NFL is defined as:

$$l_{nfl} = \frac{- \sum_{k=1}^K q(k|\mathbf{x}_i) (1 - p(k|\mathbf{x}_i))^\gamma \log(p(k|\mathbf{x}_i))}{- \sum_{j=1}^K \sum_{k=1}^K q(y=j|\mathbf{x}_i) (1 - p(k|\mathbf{x}_i))^\gamma \log(p(k|\mathbf{x}_i))}, \quad (2)$$

where $\gamma \geq 0$ is an hyperparameter. The RCE loss is:

$$l_{rce} = - \sum_{k=1}^K p(k|\mathbf{x}_i) \log(q(k|\mathbf{x}_i)). \quad (3)$$

where $q(k|\mathbf{x}_i) = 0$ is truncated to a small value such that $\log(q(k|\mathbf{x}_i)) = -4$. The combination $l_{nfl+rce}$ following the framework simply gives a different weights α and β to each loss: $\alpha \cdot l_{nfl} + \beta \cdot l_{rce}$. The hyperparameters α and β control the balancing between active and passive learning. For simplicity, they are set to 1.0 without any tuning.

Liu et al. [22] propose another framework to deal with NLL based on the ‘‘early learning’’ phase. The loss, called Early Learning Regularization (ELR), adds a regularization term to capitalize on early learning. ELR is not strictly speaking a robust loss but belongs to robust penalization and label correction methods. The penalization term corrects the CE based on estimated soft labels identified with semi-supervised learning techniques. It prevents memorization of false labels by steering the model towards these targets:

$$l_{elr} = l_{ce} + \frac{\lambda_{elr}}{N} \log \left(1 - \sum_{k=1}^K p(k|\mathbf{x}_i) t(k|\mathbf{x}_i) \right). \quad (4)$$

The target is not set equal to the model output but is estimated with a temporal ensembling from semi-supervised methods. Let $t(k|\mathbf{x}_i)^{(l)}$ denote the target for example \mathbf{x}_i at iteration l of training with a momentum β :

$$t(k|\mathbf{x}_i)^{(l)} = \beta t(k|\mathbf{x}_i)^{(l-1)} + (1 - \beta) p(k|\mathbf{x}_i)^{(l)}. \quad (5)$$

3.2 Contrastive learning

Contrastive learning methods learn representations by contrasting positive and negative examples. A typical framework is composed of several blocks [6]:

- **Data augmentation:** It decouples the pretext tasks from the network architecture. Chen et al. [4] study broadly the impact of data augmentation. We follow their suggestion combining random crop (and flip), color distortion, Gaussian blur, and gray-scaling.
- **Encoding:** The encoder extracts features (or representation) from augmented data samples. A classical choice for the encoder is the ResNet model [10] for image data. The final goal of the contrastive approach is to find correct weights for the encoder.
- **Loss function:** The loss function usually combines positive and negative pairs. The Noise Contrastive Estimation (NCE) and its variants are popular choices. The general formulation for such loss function is defined for the i -th pair as [37]:

$$L_i = -\log \frac{\exp(\mathbf{z}_i^T \mathbf{z}_{j(i)}/\tau)}{\sum_{a \in A(i)} \exp(\mathbf{z}_i^T \mathbf{z}_a/\tau)}, \quad \text{with } i \in I, \quad (6)$$

where \mathbf{z} is a feature vector, I is the set of indexes in the mini-batch, i is the index of the anchor, $j(i)$ is the index of an augmented version of the anchor source image, $A(i) = I \setminus \{i\}$, and τ is a temperature controlling the dot product. The denominator includes one positive and K negative pairs.

- **Projection head:** That step is not used in all frameworks. The projection head maps the representation to a lower-dimensional space and acts as an intermediate layer between the representation and the embedding pairs. Chen et al. [4, 5] show that the projection head helps to improve the representation quality.

4 A framework coupling contrastive learning and noisy labels

As illustrated in Figure 2, our method classifies noisy samples in a two phased process. First, a classifier pre-trained with contrastive learning produces train-set pseudo-labels (pre-training phase, in panel a), used during the training of a subsequent fine-tuning phase (panel b). The underlying intuition is that the predicted pseudo-labels are more accurate than the original noisy labels. The contrastive learning performed in the first phase (panel a1) improves the classifier (panel a2), sensitive to noisy labels; the resulting model can

be also used in a standalone way with a reduced number of hyperparameters, without the underlying fine-tuning phase.

The second phase leverages the pseudo-labels predicted by the pre-training in all underlying steps (b1-b3). To mitigate the effect of potentially incorrectly predicted pseudo-labels, a Gaussian Mixture Model (GMM, panel b1) with 2 components follows the small loss-trick to predict for each sample the probability of correctness. This value is used as a weight in a supervised contrastive step (panel b2), performed to improve the learned representations by taking advantage of the label information. A classification head is added to the contrastive model in order to produce the final predictions (panel b3). The fine-tuning phase can be seen as an adaptation of the pre-training phase to handle pseudo-labels.

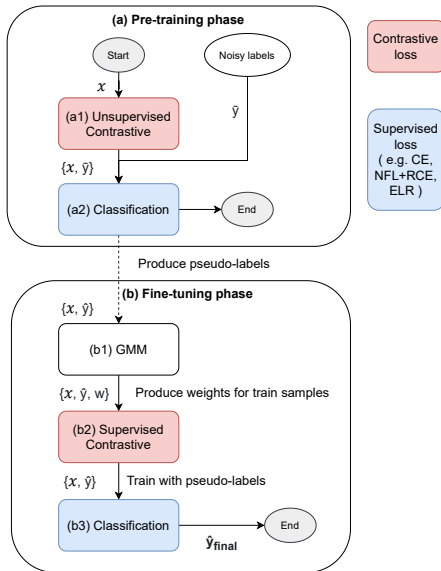


Figure 2: Overview of the framework consisting of two phases: pre-training (panel a) and fine-tuning (panel b). First, a contrastive learning phase (a1) is performed to pre-train the model in an unsupervised way for a classification phase (a2) which then predicts train-set pseudo-labels \hat{y} . The fine-tuning phase uses \hat{y} as a new ground truth. A GMM model (b1) predicts the probability of correctness for each sample, used as a corrective weight factor in a supervised contrastive training (panel b2). The final predictions \hat{y}_{final} are produced by the (b3) classifier.

To maximize the impact of the contrastive learning on the underlying classification, the supervised training is performed in 2 steps: a warm-up step, updating only the classifier layer (while keeping the encoder frozen) is followed by the full model training. We compared three different loss functions for the supervised classification: the classical CE, the robust NFL+RCE, and the ELR loss.

4.1 Sample selection and correction with pseudo-labels

Pseudo labels represent one hot encoded model’s predictions on the training set. Pseudo-labels were initially used in semi-supervised learning to produce annotations for unlabelled data; in the noisy label setting, various techniques (e.g. DivideMix, etc) identify a subset with a high likelihood of correctness and treat the remaining samples as the unlabeled counterpart in semi-supervised learning. In this work, we elaborate on the observation that the training set labels, predicted after training the model with a noise-robust loss function (i.e. the pseudo labels), are more accurate than the ground truth. This observation is supported by the results in Figure 3, depicting the accuracy of pseudo labels predicted on CIFAR100, contaminated with various levels of asymmetric (panel a) and symmetric (panel b) noise. The pseudo labels are more accurate than the corrupted ground truth and bring a higher gain in performance as the noise ratio increases.




Figure 3: Accuracy of pseudo labels on all simulated settings with asymmetric (a) and symmetric (b) noise, evaluated on CIFAR100. The correctness of the ground truth is represented on the x axis, while the accuracy of predicted pseudo labels on the y axis. In all experiments, the pseudo labels have a higher accuracy than the corrupted ground truth and this gain increases with the noise ratio.




Figure 4: Accuracy of the entire training set (in blue) compared to the clean train subset (in red); The clean subset’s percentual size is depicted in green. The example is performed on CIFAR100, with 40% symmetric noise.

As proposed by other approaches [1], the loss value on train samples can be used to discriminate between clean and

misclassified samples. The sample correctness probability is computed by fitting a 2 components GMM on the distribution of losses [19]. The underlying probability is used as a sample weight $w_i = p(k = 0 | l_i)$, where l_i is the loss for sample i and $k = 0$ is the GMM component associated to the clean samples (lowest loss). **Figure 4** depicts the evolution of the clean training set identified by GMM on an example: its accuracy grows from 0.6 to 0.93 while the size stabilizes at 60% of the training set.

4.2 Weighted supervised contrastive learning

A modification to the contrastive loss defined in **Equation 6** has been proposed to leverage label information [16]:

$$L_i = -\log \frac{1}{|P(i)|} \sum_{p \in P(i)} \frac{\exp(\mathbf{z}_i^T \mathbf{z}_p / \tau)}{\sum_{a \in A(i)} \exp(\mathbf{z}_i^T \mathbf{z}_a / \tau)}, \quad (7)$$

where $P(i) = \{j \in I \setminus \{i\}, y_j = \tilde{y}_i\}$ with \tilde{y}_i the prediction of the model for input \mathbf{x}_i .

As explained in the previous section, the loss value for the training set samples is used to fit a GMM with 2 components, corresponding to correctly and incorrectly labeled samples. We adapted the supervised representation loss to employ w , a weighting factor representing the sample probability of membership to the correctly labeled component. Thus, likely misclassified samples having large loss values would contribute only marginally to the supervised representations:

$$L_i = -\log \frac{1}{|P(i)|} \sum_{p \in P(i)} \frac{\tilde{w}_{p,i} \exp(\mathbf{z}_i^T \mathbf{z}_p / \tau)}{\sum_{a \in A(i)} \exp(\mathbf{z}_i^T \mathbf{z}_a / \tau)}, \quad (8)$$

where $\tilde{w}_{p,i}$ is a modified version of w_p such as $\tilde{w}_{p,i} = 1$ if $p = j(i)$ else $\tilde{w}_{p,i} = w_i$. If all samples are considered as noisy, **Equation 8** is simplified into the classical unsupervised contrastive loss in **Equation 6**.

5 Experiments

The framework is assessed on three benchmarks and the contribution of each block identified in **Figure 2** is analyzed.

5.1 Datasets

CIFAR10 and **CIFAR100** [17]. These experiments assess the accuracy of the method against synthetic label noise. The two datasets are contaminated with simulated symmetric or asymmetric label noise reproducing the heuristic in Ma et al. [23]. The symmetric noise consists in corrupting an equal arbitrary ratio of labels for each class. The noise level varies from 0.2 to 0.8. For asymmetric noise [22, 30], sample

labels have been flipped within a specific set of classes, thus providing confusion between predetermined pairs of labels. For CIFAR100, 20 groups of super-classes have been created, each consisting of 5 sub-classes. The label flipping is performed only within each super-class circularly. The asymmetric noise ratio is explored between 0.2 and 0.4.

Webvision [21]. This is a real-world dataset with noisy labels. It contains 2.4 million images crawled from the web (Google and Flickr) that share the same 1,000 classes from the ImageNet dataset. The noise ratio varies from 0.5% to 88%, depending on the class. In order to speed-up the training time, we used mini Webvision [14], consisting of only top 50 classes in the Google subset (66,000 images).

Clothing1M [39]. Clothing 1M is a large real-world dataset consisting of 1 million images on 14 classes of clothing articles. Being gathered from e-commerce websites, Clothing1M embeds an unknown ratio of label noise. Additional validation and test sets, consisting of 14k and 10k clean labeled samples have been made available. In order to speed-up the training time, we selected a subset of 56,000 images keeping the initial class distribution.

Both Webvision and Clothing1M images were resized to 128×128 . Therefore, the reported results may differ from other papers cropping the images to a 224×224 resolution.

5.2 Settings

We use the contrastive SimCLR framework [4] with a ResNet18 [10] (without ImageNet pre-training) as encoder. A projection head is added after the encoder for the contrastive learning with the following architecture: a multi-layer perceptron with one hidden layer and a ReLU non-linearity. The classifier following the contrastive step has a simple multilayer architecture: a single hidden layer with batch normalization and a ReLU activation function.

For all supervised classification, we use SGD optimizer with momentum 0.9 and cosine learning rate annealing. The NFL hyperparameter γ is set to 0.5. Unlike the original paper, the ELR hyperparameters do not depend on the noise type: the regularization coefficient λ_{elr} and the momentum β are set to 3.0 and 0.7.

All codes are implemented in the PyTorch framework [29]. The experiments for CIFAR are performed with a single Nvidia TITAN V-12GB and the experiments for Webvision and Clothing1M are performed with a single Nvidia Tesla V100-32GB, demonstrating the accessibility of the method. All experiments presented in the next section evaluate our method’s performance with the top-1 accuracy score.

6 Results

6.1 Impact of contrastive pre-training

To evaluate the impact of the contrastive pre-training on the classification model, the proposed method (pre-training phase) is compared with a baseline classifier, trained for 200 epochs without contrastive learning. For each simulated dataset, we compare robust losses (e.g. NLF+RCE and ELR) and cross entropy. Results for CIFAR10 and CIFAR100 are depicted in [Table 1](#) for different levels of symmetric and asymmetric noise. The pre-training improves the accuracy of the three different baselines for both datasets with different types and ratios of label noise. The largest differences are observed for the noisiest case with 80% noise. The pre-training outperforms the baselines by large margins between 10 and 75 for CIFAR10 and between 5 and 30 for CIFAR100.

Table 1: Accuracy for CIFAR10 and CIFAR100 using symmetric (0.2 - 0.8) and asymmetric noise (0.2 - 0.4). Training from scratch or from pre-trained representations. Best scores are in bold for each noise scenario and each loss.

		CIFAR10		CIFAR100		
Type	η	Loss	Base	Pre-t.	Base	Pre-t.
Sym	0.2	ce	77.2	87.7	55.6	56.5
		elr	90.3	93.0	64.1	67.4
		nfl+rce	91.0	92.7	66.6	68.8
	0.4	ce	58.2	78.0	39.9	41.9
		elr	82.3	92.0	56.9	62.0
		nfl+rce	87.0	91.4	60.2	66.3
	0.6	ce	35.2	59.2	21.8	26.8
		elr	64.2	90.4	40.6	55.7
		nfl+rce	80.2	88.1	47.0	61.8
	0.8	ce	17.0	27.3	7.80	12.4
		elr	18.3	84.8	16.2	45.3
		nfl+rce	42.8	59.9	20.1	50.2
Asym	0.2	ce	84.0	87.9	59.0	57.8
		elr	91.8	92.4	70.3	70.2
		nfl+rce	90.2	91.5	63.9	68.4
	0.3	ce	79.2	83.9	50.6	50.4
		elr	89.6	91.7	69.8	69.3
		nfl+rce	86.7	89.9	53.5	63.5
	0.4	ce	75.3	77.8	41.8	42.4
		elr	72.3	89.5	67.6	67.6
		nfl+rce	80.0	82.4	40.6	47.8

In addition to the comparisons with ELR and NLF+RCE, performed using our implementations (column Base in [Table 1](#)), we present the results reported by other recent competing methods. As shown in the introduction, numerous contributions have been made to the field in the last years. Six

Table 2: Accuracy scores compared with 6 methods on CIFAR10 (C10) and CIFAR100 (C100). The cases most affected by dropout are presented, with symmetric (S) and asymmetric (A) noise. Top-2 scores are in bold

	C10 80% S	C10 40% A	C100 80% S	C100 40% A
Ours (ELR)	84.8	89.5	45.3	67.6
ELR [22]	73.9	91.1	29.7	73.2
Taks [31]	40.2	73.4	16.0	35.2
Co-teach+ [40]	23.5	68.5	14.0	34.3
DivideMix [19]	92.9	93.4	59.6	72.1
SELF [27]	69.9	89.1	42.1	53.8
JoCoR [36]	25.5	76.1	12.9	32.3

recent representative methods are selected for comparison: Taks [31], Co-teaching+ [40], ELR [22], DivideMix [19], SELF [27], and JoCoR [36]. The results are presented in [Table 2](#). The difference between the scores reported by ELR and those obtained with our run (using the same implementation, but slightly different hyper-parameters and a ResNet34 versus ResNet18) suggests that the method is less stable on data contaminated with asymmetric noise and sensitive to small changes hyperparameters. Moreover, ELR proposes hyperparameters having different values depending on the type of dataset (i.e. CIFAR10/CIFAR100) and underlying noise (i.e. symmetric/asymmetric), identified after a hyperparameter search exercise. The best scores are reported by DivideMix and they surpass all other techniques. One can note DivideMix uses a PreAct ResNet18 while we use a classical ResNet18. Moreover, a recent study [28] attempted to replicate these values and reported significantly lower results on CIFAR100 (i.e. 49.5% instead of 59.6% on symmetric data and 50.9% instead of 72.1% on asymmetric data). Our framework compares favourably with the other competing methods, both on symmetric and asymmetric noise.

Webvision and Clothing1M results are presented in [Table 3](#). The contrastive framework outperforms the respective baselines for the three loss functions. Because the images have a reduced size, and for Clothing1M, we use a smaller training set, the direct comparison with competing methods is less relevant. However, the observed gap in performance is significant and promising for training images with higher resolution. Moreover, a ResNet50 model has been trained with our framework on the Webvision dataset with a higher resolution (224×224). The accuracy reaches respectively 75.7% and 76.2% for CE and ELR. These results are very close to the values reported with DivideMix (77.3%) and

ELR+ (77.8%) using a larger model, Inception-ResNet-v2 (the difference is 4% on the ImageNet benchmark [3]).

Table 3: Top-1 accuracy for mini-Webvision and Clothing1M. Best scores are in bold for each dataset and each loss. Pre-t represents the pre-training phase while Fine-tune refers to the results after the fine-tuning step.

Loss	Webvision			Clothing1M		
	Base.	Pre-t.	Fine-tune	Base.	Pre-t.	Fine-tune
ce	51.8	57.1	58.4	54.8	59.1	61.5
elr	53.0	58.1	59.0	57.4	60.8	60.4
nfl+rce	49.9	54.8	58.2	57.4	59.4	60.1

Supported by this first set of experiments, the preliminary pre-training with contrastive learning shows great performances. The accuracy of both traditional and robust-loss classification models is significantly improved.

6.2 Sensitivity to the hyperparameters

Estimating the best hyperparameters is complex for datasets with noisy labels as clean validation sets are not available. For instance, Ortego et al. [28] show that two efficient methods (eg. ELR and DivideMix) could be sensitive to specific hyperparameters. Therefore a hyperparameter sensitivity study has been carried out to estimate the stability of the framework for the learning rate. Figure 5 depicts the sensitivity on CIFAR100 with 80% noise. CE and NFL+RCE seem to have opposite behaviors. The CE reaches competitive results with small learning rates but is prompt to overfitting for higher learning rates. The NFL+RCE loss tends to underfitting for the lowest learning rates but is quite robust for higher values. The ELR loss has the smallest sensitivity to the learning for the investigated range but does not reach the best values obtained with CE or NFL+RCE. We can assume that the regularization term coupled with pre-training is very efficient. It prevents memorization of the false labels as observed with CE.

This sensitivity analysis is limited to the learning rate. Investigating the impact of other hyperparameters, such as the momentum β or the regularization factor λ_{elr} , could be interesting. In their original papers, ELR and NFL+RCE reach respectively 25.2% and 30.3% with other hyperparameters. These values are still far from the improvements brought by the contrastive pre-training but it suggests that the results could be improved with different hyperparameters.

Our empirical results indicate that the analyzed methods may be sensitive to hyperparameters. Despite the promised robustness to label noise, the analyzed robust losses are




Figure 5: Learning rate sensitivity for CIFAR100 with 80% noise. The explored learning rate values are $\{0.001, 0.01, 0.1, 1.0\}$. The baseline (dashed line) is compared with our framework (solid line).

also affected by overfitting or underfitting. Our experiments have been built upon the parameters recommended in each issuing paper (e.g. ELR, SIMCLR) but, since the individual building blocks can be affected by small variations in input parameters, the performance of our method may also be impacted. Finding a relevant method to estimate proper hyperparameters in NLL remains a challenge. In the absence of a clean validation set, identifying when overfitting starts also remains an open challenge.

6.3 Impact of the fine-tuning phase

Experimental results on synthetic label noise, depicted in Figure 6, show that continuing the presented pre-training block (Figure 2) with the fine-tuning phase increases the accuracy in over 65% of cases on CIFAR10 and over 80% of cases on CIFAR100. For both datasets, asymmetric noise data benefit more from this approach than symmetric noise. All experiments only use the input parameters proposed in the loss-issuing papers.

The sample selection has also got a positive impact on the two real-world datasets, as shown in Table 3 by the "Fine-tune" columns. The average accuracy improvement is about 1.8%. Only the ELR loss function slightly decreases the performance on Clothing1M.

Enriching pretrained models with sample weighting and selection, pseudo labels instead of corrupted targets, and supervised contrastive pre-training can improve the classification accuracy. However, such an approach raises the question of a trade-off between complexity, accuracy improvement, and computation time.

7 Discussion

In addition to the presented fine-tuning phase, we evaluated the performance of other promising techniques, such

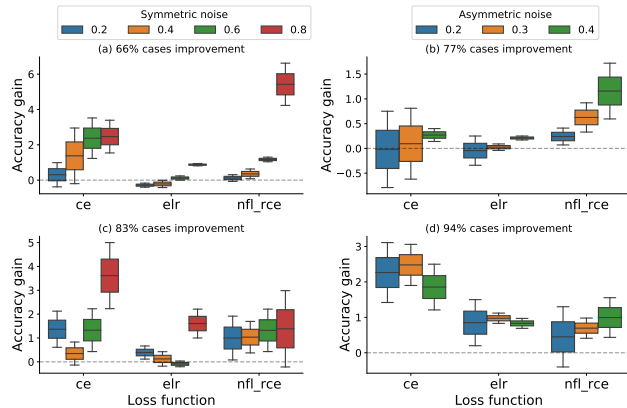


Figure 6: Accuracy gain when performing the fine-tuning phase after the pre-training block (computed as the difference between fine-tuning accuracy and pre-training accuracy). The plot gathers the results for all noise ratios on CIFAR10 (panels a, b) and CIFAR100 (c, d) with symmetric (first column) and asymmetric (second column) noise.

as the dynamic bootstrapping with mixup [1]. This strategy has been developed to help convergence under extreme label noise conditions. The improvement that dynamic bootstrapping can bring when used after pre-training is depicted in Figure 7. In most of the cases, this technique improves the accuracy, as indicated by the positive accuracy gain scores, measuring the difference between the accuracy after dynamic bootstrapping and the accuracy of the pre-training phase. ELR and CE benefit most from this addition for CIFAR100. The impact of the dynamic bootstrapping should also be analyzed for the fine-tuning phase and for larger datasets, such as Webvision or Clothing1M.




Figure 7: Gain for the dynamic bootstrapping (CIFAR100) with asymmetric (a) and symmetric noise (b). Bootstrapping is an alternative to the proposed fine-tuning phase. Each color is associated to a noise ratio.

One of the major drawbacks of our method is the extra computational time needed to learn representations with contrastive learning. The pre-training phase doubles the execution time of a reference baseline, consisting of performing only a single classification step, while the entire

framework increases the execution time 3 to 4 times the baseline value. However, the contrastive learning does not increase the need for GPU memory if the batch size is limited for the contrastive learning [11, 26]. The computational time could be reduced by initializing the contrastive step with the pretrained weights from ImageNet.

Most state-of-the-art approaches also leverage computationally expensive settings, consisting of larger models (e.g. ResNet50), dual model training, or data augmentation such as mixup. In this work, we explored the limits of a restricted computational setting, consisting of a single GPU and 8GB RAM. All experiments use a ResNet18 model, batch sizes of 256, and for real-world datasets, the images have been rescaled (e.g. 128×128 instead of 224×224). We also foresee that the contrastive learning step could be improved by images with higher resolutions as smaller details could be identified in the representation embedding.

There remain multiple open problems for future research, such as: i) identifying the start of the memorization phase in the absence of a clean dataset, ii) studying the impact of contrastive learning on other models for noisy labels such as DivideMix, iii) comparing SimCLR approach in the context of noisy labels with other contrastive frameworks and other self-supervised approaches, and iv) having a better theoretical understanding of the interaction between the initial state precomputed with contrastive learning and the classifier in presence of noisy labels. Moreover, the analysis carried out in this work should be validated on larger settings, in particular on Clothing1M with a ResNet50, higher resolutions, and the full dataset.

8 Conclusions

In this work, we presented a contrastive learning framework optimized with several adaptations for noisy label classification. Supported by an extensive range of experiments, we conclude that a preliminary representation pre-training improves the performance of both traditional and robust-loss classification models. Additionally, multiple techniques can be used to fine-tune and further optimize these results; however, no approach provides a significant improvement systematically on all types of datasets and label noise. The cross-entropy penalized by Early-Learning Regularization (ELR) shows the best overall results for synthetic noise but also real-world datasets.

However, the training phases remain sensitive to input configuration. Overfitting is the common weakness of all studied models. When trained with tuned parameters, even traditional (cross-entropy) models provide competitive results, while robust-losses are less sensitive. The typical noisy label adaptations, such as sample selection or weight-

ing, the usage of pseudo labels, or supervised contrastive losses, improve the performance to a lesser extent but increase the framework’s complexity. We hope that this work will promote the use of contrastive learning to improve the robustness of the classification process with noisy labels.

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. 2, 4, 8
- [2] Devansh Arpit, Stanisław Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *International Conference on Machine Learning*, pages 233–242. PMLR, 2017. 2
- [3] Simone Bianco, Remi Cadene, Luigi Celona, and Paolo Napolitano. Benchmark analysis of representative deep neural network architectures. *IEEE Access*, 6:64270–64277, 2018. 7
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. 1, 2, 3, 5
- [5] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 3
- [6] William Falcon and Kyunghyun Cho. A framework for contrastive self-supervised learning and designing a new approach. *arXiv preprint arXiv:2009.00104*, 2020. 3
- [7] Aritra Ghosh, Himanshu Kumar, and PS Sastry. Robust loss functions under label noise for deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017. 3
- [8] J. Goldberger and E. Ben-Reuven. Training deep neural networks using a noise adaptation layer. In *ICLR*, 2017. 1
- [9] 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*, pages 8527–8537, 2018. 1, 2
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3, 5
- [11] 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. 1, 2, 8
- [12] Olivier Henaff. Data-efficient image recognition with contrastive predictive coding. In *International Conference on Machine Learning*, pages 4182–4192. PMLR, 2020. 2
- [13] 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*, pages 10456–10465, 2018. 1
- [14] 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. 1, 2, 5
- [15] Yannis Kalantidis, Mert Bulent Sariyildiz, Noe Pion, Philippe Weinzaepfel, and Diane Larlus. Hard negative mixing for contrastive learning. *arXiv preprint arXiv:2010.01028*, 2020. 2
- [16] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33, 2020. 5
- [17] A Krizhevsky. Learning multiple layers of features from tiny images. *Master’s thesis, University of Tront*, 2009. 5
- [18] Phuc H Le-Khac, Graham Healy, and Alan F Smeaton. Contrastive representation learning: A framework and review. *IEEE Access*, 2020. 2
- [19] Junnan Li, Richard Socher, and Steven C.H. Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *International Conference on Learning Representations*, 2020. 1, 2, 5, 6
- [20] Junnan Li, Caiming Xiong, and Steven CH Hoi. Mopro: Webly supervised learning with momentum prototypes. *arXiv preprint arXiv:2009.07995*, 2020. 2
- [21] 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
- [22] Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning regularization prevents memorization of noisy labels. *Advances in Neural Information Processing Systems*, 33, 2020. 1, 2, 3, 5, 6
- [23] 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. 1, 2, 3, 5
- [24] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Barambe, and Laurens van der Maaten. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 181–196, 2018. 1
- [25] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *Proceedings*

-
- of the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6707–6717, 2020. 2
- [26] Jovana Mitrovic, Brian McWilliams, and Melanie Rey. Less can be more in contrastive learning. In *"I Can't Believe It's Not Better!" NeurIPS 2020 workshop*, 2020. 8
- [27] 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. In *International Conference on Learning Representations*, 2019. 2, 6
- [28] Diego Ortego, Eric Arazo, Paul Albert, Noel E. O'Connor, and Kevin McGuinness. Multi-objective interpolation training for robustness to label noise, 2020. 2, 6, 7
- [29] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. 5
- [30] 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. 1, 5
- [31] Heon Song, Nariaki Mitsuo, Seiichi Uchida, and Daiki Suehiro. No regret sample selection with noisy labels. *arXiv preprint arXiv:2003.03179*, 2020. 6
- [32] 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
- [33] Hwanjun Song, Minseok Kim, Dongmin Park, and Jae-Gil Lee. Learning from noisy labels with deep neural networks: A survey. *arXiv preprint arXiv:2007.08199*, 2020. 2
- [34] Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 322–330, 2019. 1, 2
- [35] Zhuowei Wang, Jing Jiang, Bo Han, Lei Feng, Bo An, Gang Niu, and Guodong Long. Seminll: A framework of noisy-label learning by semi-supervised learning. *CoRR*, abs/2012.00925, 2020. URL <https://arxiv.org/abs/2012.00925>. 2
- [36] Hongxin Wei, Lei Feng, Xiangyu Chen, and Bo An. Combating noisy labels by agreement: A joint training method with co-regularization, 2020. 6
- [37] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3733–3742, 2018. 3
- [38] 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*, pages 6838–6849, 2019. 1
- [39] Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2691–2699, 2015. 5
- [40] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *International Conference on Machine Learning*, pages 7164–7173. PMLR, 2019. 6
- [41] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016. 1
- [42] Hui Zhang and Quanming Yao. Decoupling representation and classifier for noisy label learning, 2020. 2
- [43] Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems*, 31:8778–8788, 2018. 1