

Embedding contrastive unsupervised features to cluster in- and out-of-distribution noise in corrupted image datasets

Paul Albert, Eric Arazo, Noel E. O'Connor, Kevin McGuinness
paul.albert@insight-centre.org

School of Electronic Engineering,
Insight SFI Centre for Data Analytics, Dublin City University (DCU)

Abstract. Using search engines for web image retrieval is a tempting alternative to manual curation when creating an image dataset, but their main drawback remains the proportion of incorrect (noisy) samples retrieved. These noisy samples have been evidenced by previous works to be a mixture of in-distribution (ID) samples, assigned to the incorrect category but presenting similar visual semantics to other classes in the dataset, and out-of-distribution (OOD) images, which share no semantic correlation with any category from the dataset. The latter are, in practice, the dominant type of noisy images retrieved. To tackle this noise duality, we propose a two stage algorithm starting with a detection step where we use unsupervised contrastive feature learning to represent images in a feature space. We find that the alignment and uniformity principles of contrastive learning allow OOD samples to be linearly separated from ID samples on the unit hypersphere. We then spectrally embed the unsupervised representations using a fixed neighborhood size and apply an outlier sensitive clustering at the class level to detect the clean and OOD clusters as well as ID noisy outliers. We finally train a noise robust neural network that corrects ID noise to the correct category and utilizes OOD samples in a guided contrastive objective, clustering them to improve low-level features. Our algorithm improves the state-of-the-art results on synthetic noise image datasets as well as real-world web-crawled data. Our work is fully reproducible github.com/PaulAlbert31/SNCF.

Keywords: Computer vision, image classification, label noise, out-of-distribution noise

1 Introduction

Convolutional neural networks (CNN) show remarkable classification accuracies on large image datasets carefully curated by large numbers of annotators [18, 26], which entails that the data gathering and labeling process has become a significant part of the effort required when designing a deep learning image classification algorithm. Recent research proposes to reduce the cost of the data gathering




Fig. 1. Visualization of the linear separation between OOD and ID unsupervised contrastive representations on the 2D hypersphere. CIFAR-10 corrupted with $r_{in} = r_{out} = 0.2$, OOD from ImageNet32. Linear separability in 2D at the dataset level is 92.49% but increases to 98%+ for 128D

effort in multiple ways, including semi-supervised learning [9] where only a small part of the images have been labeled by humans, which is often the most time consuming part of the dataset building process; unsupervised learning [5], where visual features are learned without the need for labels; and web-crawled datasets [24] constructed from search engine queries, although this results in the inclusion of incorrect samples assigned to the incorrect class or not belonging to the distribution and corrupting the CNN’s convergence. This paper tackles this latter challenge. Designing algorithms capable of training highly accurate CNNs even when trained on imperfect web-crawled data is an important step towards the widespread deployment and take up of computer vision algorithms in practice. CNNs have been shown to completely overfit noisy samples in a dataset without proper regularization [48], which degrades performance. More specifically, the noise observed in web-crawled datasets has been categorized as both in-distribution and out-of-distribution, the latter being the dominant type [1]. While in-distribution (ID) noisy images can be directly used to train the network after correcting their assigned label, out-of-distribution (OOD) images cannot be assigned to any category. Since a trusted in-distribution dataset is unavailable and the identity of clean and noisy samples is unknown, out-of-distribution detection algorithms [47, 11, 13], which use a classifier trained on clean data to be able to detect OOD samples post training, cannot be used. This further complicates the noise detection process. Once noisy images have been identified, simply ignoring out-of-distribution images has been shown to be sub-optimal as these samples still contain meaningful information for learning low-level features that can be leveraged to improve the representations learned [14, 46]. This paper proposes to tackle the in-distribution and out-of-distribution duality of the noise present

in web-crawled datasets specifically to improve the final classification accuracy. To detect the noise, we observe that unsupervised contrastive representations for OOD samples are linearly separated from ID ones on the hypersphere (see Figure 1) and train a robust network that will use current representations to correct ID noisy samples and use OOD data to improve low-level representations using contrastive learning. Our contributions are:

- A dual noise detection approach utilizing the alignment and uniformity principles of contrastive learning to detect noisy samples in a spectral embedding of unsupervised representations.
- A noise robust algorithm capable of training a CNN on a dataset corrupted with in-distribution and out-of-distribution noise, correcting the label of in-distribution whilst using out-of-distribution noise to improve low-level features.
- Experiments on controlled and real world noisy datasets demonstrating the state-of-the-art performance of our algorithm.

2 Related work

2.1 Web-crawled datasets

Web-crawled datasets are a low supervision alternative to dataset building that helps to democratize deep learning approaches. Here we will give some examples. Webvision [24] is a dataset constituted of 2.4 million images gathered using search queries on the same 1k classes as the ILSVRC12 [20] challenge. Albert *et al.* [1] have shown that the noise present in Webvision is predominantly out-of-distribution. Clothing1M [44] is a 1M images clothes classification dataset, popular dataset in the label noise community which, as shown by the authors, only contains ID noise. More recently, Sun *et al.* [37] released additional web datasets for fine grained classification tasks.

2.2 Tackling ID and OOD noise

Most of the noise robust algorithms that showed improvements on web-noise assume that all the noise is ID. The strategies used to combat ID noise include: robust loss functions [31,30], meta-learning [45,41], sample weighing using a mentor network [16,17], network regularization [49,25], semi-supervised learning [22,51,28], contrastive learning to detect the noise in MOIT [29] or improve the label correction using nearest neighbor clustering together with a semi-supervised correction in ScanMix [33]. Another class of noise robust algorithms have emerged recently to tackle noisy datasets presenting both in- and out-of-distribution noise. EvidentialMix [32] and DSOS [1] differentiates between ID noisy and OOD data using a custom noise landscape and JoSRC [46] evidences OOD samples as having a low agreement between two consistent views of the same image. The methods we compare against are described in more details in 4.1 and we direct the interested reader to a recent label noise survey by Song *et al.* [15] for an in-depth overview of state-of-the-art label noise robust algorithms.

3 Algorithm description

This paper studies image classification in the presence of label noise, where part of the available image dataset $\mathcal{X} = \{x_i\}_{i=1}^N$ and its associated classification labels $\mathcal{Y} = \{y_i\}_{i=1}^N$, with the class distribution $\{c\}_{c=1}^C$, is corrupted by N_o out-of-distribution samples and N_n in-distribution noisy samples, where $N_c = N - N_o - N_n$ is the number of in-distribution clean examples. N_o , N_n as well as the identity of the ID noisy and OOD samples are unknown. Examples of such datasets are web-crawled datasets: Webvision [24], Clothing1M [44], and the Weby Supervised Fine-Grained Recognition datasets [37]. We propose here an algorithm capable of training a convolutional neural network (CNN) Ψ on the corrupted dataset \mathcal{X} without over-fitting to the noise and capable to accurately classify examples belonging to the class distribution.

3.1 Unsupervised feature learning

First, our algorithm learns unsupervised representations from the images themselves, independently of their label. We aim here to relate images to each other in order to capture clusters of similar images. To do so, we train the N -pairs unsupervised contrastive learning algorithm which has been successfully used in metric learning [36] and unsupervised learning on images and text [21]. Given two mini-batches of size B formed from two strongly data augmented views x'_i and x''_i of $x_i \in \mathcal{X}$, we enforce u'_i and u''_i , their associated contrastive representations through Ψ , to be similar to each other and dissimilar to every other image in the batch. We compute the unsupervised contrastive loss

$$l_{unsup} = -\frac{1}{B} \sum_{i=1}^B \log \left(\frac{\exp(ip(u''_i, u'_i)/\tau_2)}{\sum_{k=1}^B \exp(ip(u''_k, u'_i)/\tau_2)} \right), \quad (1)$$

where $ip(u_1, u_2) = \frac{u_1^T u_2}{\|u_1\|_2 \|u_2\|_2}$ is the inner product operation, measuring the similarity between contrastive representations, and τ_2 a temperature hyper-parameter, fixed to 0.2 for every experiment. Mixup [49] can be optionally used to further augment x'_i by linearly interpolating it with other augmented samples from the mini-batch with a parameter μ drawn from a beta distribution with parameter 1 to produce $x'_{mix} = \mu x'_i + (1 - \mu)x'_j$ with x'_j a random sample from the mini-batch (different for every x_i) and u'_{mix} the associated representation of x'_{mix} . We then use

$$l_{mix} = -\frac{1}{B} \sum_{i=1}^B \log \left(\mu \frac{\exp(ip(u''_i, u'_{mix})/\tau_2)}{\sum_{k=1}^B \exp(ip(u''_k, u'_{mix})/\tau_2)} + (1 - \mu) \frac{\exp(ip(u''_j, u'_{mix})/\tau_2)}{\sum_{k=1}^B \exp(ip(u''_k, u'_{mix})/\tau_2)} \right), \quad (2)$$

the N -pairs loss paired with mixup as a data augmentation. This unsupervised contrastive objective has been proposed as part of the iMix [21] algorithm.

3.2 Embedding of unsupervised features

We propose not to use the learned unsupervised features directly but instead to perform a non-linear spectral dimensionality reduction on an affinity matrix (embedding). The aim of the embedding is to capture the affinities between samples and their neighbors where ID clean samples will be very similar to other samples from the same class, ID noisy samples will be similar to other ID samples from a different class and OOD samples dissimilar to any ID sample. This motivates computing the embedding at the dataset level to ensure that the similarity of ID noise to other classes is captured. We first compute the sparse similarity matrix S of size $N \times N$ where for each sample in the dataset, we compute the affinity to a fixed neighborhood size of 50 neighbors.

$$S_{ij} = (u_i^T u_j / \|u_i\|_2 \|u_j\|_2)^\gamma, \quad (3)$$

with u_i the unsupervised representation for sample x_i (not augmented) and $\gamma = 3$ a hyper-parameter regulating the importance of distant neighbors. With I_N the identity matrix of size N and D the diagonal normalization matrix where $D_{ii} = \sum_{j=1}^N S_{ij}$, we compute the normalized Laplacian

$$L = I_N - D^{-1/2} S D^{-1/2}. \quad (4)$$

We finally compute the first k eigenvectors of the normalized Laplacian L by solving

$$(L - \lambda)V = 0 \quad (5)$$

and concatenating the first k eigenvectors V of L (by increasing order of the eigenvalues λ , omitting the smallest), providing us with k features per sample to form the embedding E . In practice we use $k = 20$ for every dataset. This embedding process is commonly referred to as spectral embedding [27,35].

3.3 Unsupervised clustering of noise

Using the embedding E , we cluster embedded unsupervised features to identify three kinds of samples: clean ID, noisy ID, and OOD. In the generic case where the three types of noise are expected in the dataset, we apply the clustering at the class level and aim to discover three clusters for each class: a high density cluster of ID clean samples, a low density cluster of OOD samples and ID noisy outliers. In the case where no ID noise is present, we observe the cluster separation at the dataset level and use a two mode Gaussian mixture to retrieve each cluster.

Why does OOD noise cluster? Contrary to previous research where OOD noise is considered an outlier to the distribution [43], we observe in this paper that unsupervised contrastive learning can be effectively used to cluster noise in the feature space. We expand here on our intuition as to why this works using the alignment and uniformity principles for contrastive learning formalized by Wang *et al.* [42]. Unsupervised contrastive learning pulls together augmented representations of a same image while pushing apart representations of any other

sample in the mini-batch. Since images from a same class will be similar to each other’s augmentations, they will cluster together in the feature space to create one (or more) mode for the class (alignment principle). On the other hand, by considering OOD samples as being uniformly sampled from the set of all images, meaning much more varied in appearance than the ID set, we would expect that no compact mode would appear and that these samples would remain uniformly distributed in the feature space yet separated from the ID examples, pulled together into their respective class modes (uniformity principle). Since the features are L^2 normalized during training they exist on the surface of a unit hypersphere and one side of the sphere will contain well represented ID classes, clustered into their respective modes, while OOD noise will remain uniformly distributed be on the other side of the hypersphere and linearly separable from ID samples. Section 4.2 proposes experiments to evidence the linear separability of ID and OOD samples in the unsupervised contrastive feature space. The spectral embedding we propose has a key role to play in the clustering of the OOD noise which, although separable from the ID samples, is much more spread-out than the compact class modes of ID images. We remedy this problem by computing the affinities in S not over a fixed distance threshold but using a fixed number of neighbors.

OPTICS [2] is an algorithm which allows us to detect clusters as well as outliers: each feature point is ordered to create a chain where neighboring points are ordered next to each other. Each feature point is then labeled with a reachability cost to neighbors in a neighborhood of size V . The higher the cost, the more likely a sample is to be an outlier. Finally, clusters are identified in the ordering where “valleys” of low reachability cost evidence a cluster, themselves separated with high cost outliers.

Discovering clean and noisy clusters. Because the difficulty of learning similar unsupervised features varies from class to class in an image dataset, we propose to modify the OPTICS algorithm to become more flexible to our problem. We aim here to be able to detect varying valley sizes in the ordered reachability plot where different classes in the image dataset will have more compact classes (fine grained classes) than others (classes with highly diverse examples). In practice, we compute three different reachability orderings for three different neighborhood sizes V (75, 50, 25 neighbors), which allows us to account for cluster compactness variations across classes and noise levels. The algorithm chooses the optimal cluster assignment at the class level as being the cluster with the lower amount of outliers given at least two clusters are identified (clean and OOD). This allows us to reduce the amount of hyper-parameters to tune for the clustering to the ξ parameter of OPTICS which controls the decision boundary between clusters and outliers. Higher values for ξ imply a higher tolerance threshold meaning a lower amount of outliers.

No ID noise. In the case where we expect no ID noisy samples in the dataset, we only aim to discover a clean and an OOD cluster without outliers. In this case, the OOD cluster can be retrieved at the dataset level and we choose instead

to fit a 2 component Gaussian mixture on the embedded features to retrieve the clusters.

Clean or OOD. Once the clusters are evidenced in the ordering, the final step for the detection is to classify the clusters into clean or OOD. Although the average reachability score within the cluster could at first glance be considered a good indication of the OOD nature of a cluster, by computing the affinity matrix over a fixed neighborhood size, distances are not accurately preserved. We propose instead to compute the density of the cluster in the original unsupervised feature space, where for each sample in the dataset we compute the average distance to all other points in the cluster. We then select the cluster with the lowest density as the OOD cluster.

3.4 Spectral Noise clustering from Contrastive Features (SNCF)

Clustering the embedded unsupervised feature space provides three subsets of \mathcal{X} : \mathcal{X}_c , \mathcal{X}_n and \mathcal{X}_o , respectively the clean, ID noisy and OOD subsets. We aim to use all the available samples to train our CNN and do so by correcting ID noisy samples to their true label and using OOD samples to learn more robust low-level features. We train from scratch on each type of noise separately without using the unsupervised features to initialize the classification network.

Correcting in-distribution noise Although the unsupervised features allow the detection of incorrectly assigned samples, we find that this is not sufficient to accurately assign ID noise to the right class, especially since they might be close to other ID noisy samples themselves assigned to another incorrect class. We propose instead to correct the ID noise during the supervised training phase, using knowledge learned on clean ID samples during a warm-up pre-training. We then estimate the true labels of the detected ID noise using a consistency regularization approach. For every ID noisy sample in \mathcal{X}_n two weakly augmented versions are produced. The network then predicts on both samples ($p_{i,1}$ and $p_{i,2}$) and returns an average prediction, which, after temperature sharpening τ_1 and normalization, is used as the corrected class assignment: $y_i = \left(\frac{p_{i,1} + p_{i,2}}{2} \right)^{\tau_1}$ with $\tau_1 = 2$ in every experiment. We find temperature sharpening to be necessary to reduce the entropy of the guessed label and to encourage the network to produce more confident predictions.

Out-of-distribution samples OOD samples cannot be corrected to any label in the distribution but we propose to include them in an additional guided contrastive loss minimization objective to learn low level features. Once the noise detection algorithm has run, we re-embed the unsupervised features of detected OOD noise and use OPTICS to discover clusters of the most similar samples in the OOD data. At training time, we augment each sample in the dataset into one weakly and one strongly augmented view, producing two mini-batches of the same images augmented differently. We then enforce ID samples belonging to the

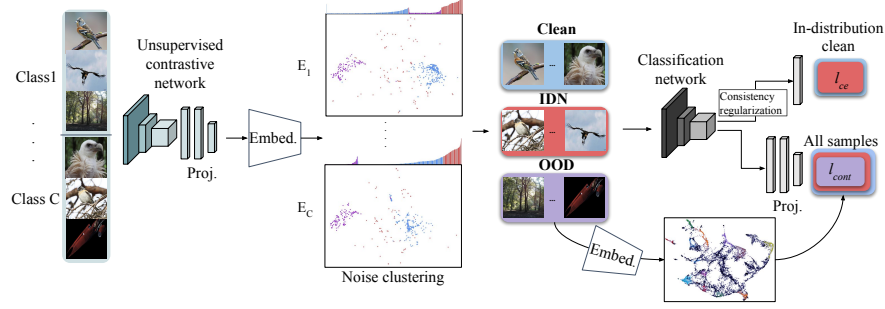


Fig. 2. Visualization of the algorithm. The unsupervised features are embedded to create E and evaluated at the class level (E_1, \dots, E_C) to cluster clean and OOD samples. The detected OOD samples are re-embedded from their unsupervised features to detect clusters of similar images. We correct the ID noise using consistency regularization and the OOD sample’s cluster assignments are used together with the classes of all in-distribution samples in a guided contrastive objective

same class (corrected for the ID noise) as well as OOD samples from the same unsupervised cluster to be similar while being dissimilar to every other example in the mini-batch. OOD samples not assigned to any unsupervised cluster are considered similar to their augmented view only. This guided contrastive objective is described in equation 7. Note here that the similarities are enforced between the two mini-batches of augmented views alone. We attempted to enforce similarities inside the same augmented batch but noticed no accuracy improvements.

3.5 Loss objectives

We consider here that p_i is the current softmax prediction of Ψ on sample $x_i \in \mathcal{X}$. Our algorithm aims to optimize over two objectives during training. The first is the classification objective on the detected clean samples \mathcal{X}_c and the ID samples from \mathcal{X}_n whose label has been corrected. We use the cross entropy loss:

$$l_{ce} = \sum_{i=1}^{N_c + N_n} y_i^T \log(p_i). \quad (6)$$

Secondly, we minimize the guided contrastive learning objective, grouping samples of the same class and OOD samples from the same OOD cluster together using their respectively weakly and strongly augmented projected representations r_i and r'_i , projected from the classification space to the contrastive space

$$l_{cont} = -\frac{1}{N} \sum_{i=1}^N \frac{1}{B} \sum_{b=1}^B e_{i,b} \log \left(\frac{\exp(ip(r_b, r'_i)/\tau_2)}{\sum_{k=1}^B \exp(ip(r_k, r'_i)/\tau_2)} \right), \quad (7)$$

with $e_{i,b} = 1$ if sample i is considered similar to sample b , $e_{i,b} = 0$ otherwise. Note that this objective can be paired with mixup as in the unsupervised objective in

equation 2. The final loss minimized by our algorithm is:

$$l = l_{ce} + \beta l_{cont}, \quad (8)$$

where β is an hyper-parameter (typically 1). Figure 2 illustrates the algorithm.

4 Experiments

4.1 Implementation details

We form each mini-batch by aggregating an equal number of clean ID, noisy ID and OOD samples. Since the OOD is ignored in the ID objective (eq 6) and in order to have the same batch size for the ID and the contrastive forward pass, we form the ID mini-batch by aggregating two weakly augmented views of the clean data with one weakly augmented view for the ID noisy data (ID clean, ID noisy, OOD for the contrastive mini-batch). For the weak data augmentations we use cropping with padding and random horizontal flip and the strong SimCLR augmentations [6]. We warm-up the network on the detected clean data from scratch for 15 epochs in every experiment (except 5 for WebVision) and start both the ID noise correction and the guided contrastive objective after this. For a fair comparison with other approaches, the unsupervised features are not used to initialize the network in the robust classification phase. More experiments are proposed in the supplementary material. Since our algorithm minimizes a contrastive loss, we find that adding a non-linear projection head [6,7] to project features from the classification space to the contrastive space is beneficial in reconciling the training objectives. The final number of projected contrastive features is 128 and the projection head is not used at test time. We use stochastic gradient descent (SGD) with a weight decay of 5×10^{-4} and mixup [49] augmentation with $\alpha = 1$ for all experiments.

Training the unsupervised algorithm. We train the unsupervised algorithm using the same network as the robust classification phase. In cases where the resolution is 227×227 or above, we train and evaluate the unsupervised features at resolution 84×84 as this helps to keep training time and memory consumption reasonable yet still separates the OOD and ID clusters. The algorithm is trained for 2000 epochs, with a batch size of 256, starting with a learning rate of 0.01 and reducing it by a factor of 10 at epochs 1000, 1500. We use the mixup version on the unsupervised objective (iMix [21]).

Synthetically corrupted datasets. We conduct a first series of experiments on synthetically corrupted versions of CIFAR-100 [19] where we control the ID noise and OOD noise. We use the same configuration as in Albert *et al.* [1] and note r_{in} and r_{out} the corruption ratios for ID noisy and OOD noise respectively with $r_{in} + r_{out}$ the total noise level. Our focus here is on the OOD noise rate more than ID noise, which is less present in web-crawled datasets [1]. We introduce OOD noise by replacing original images with images from another dataset, either ImageNet32 [8] or Places365 [50]. For the ID noise, we randomly flip the labels of a portion of the dataset to a random label (uniform noise). The dataset size remains 50K images after noise injection. We train on CIFAR-100 using

a PreActResNet18 [18] trained with a batch size of 256 for 100 epochs with a learning rate of 0.1, reducing it by a factor of 10 at epochs 50 and 80.

Web noise corruption. We conduct experiments on miniImageNet [40] corrupted by web noise from Jiang *et al.* [17] (Controlled Noisy Web Labels, CNWL) where the severity of the web noise corruption is controlled. This dataset is an example where ID noise is very limited and where we find that using the 2 components GMM is sufficient to detect the noise at the dataset level (see Section 3.3). We train on this dataset at two different resolutions, first 299×299 , which is the original configuration proposed by Jiang *et al.* [17] and second the 32×32 resolution adopted in recent works [10,45,33]. For the 299×299 configuration, we train an InceptionResNetV2 [38] with a batch size of 64 for 200 epochs with a learning rate of 0.01, reducing it by a factor of 10 at epochs 100, 160. For the 32×32 configuration, we use the same configuration as CIFAR-100.

Real-world dataset. We evaluate our model on the (mini)Webvision [24] dataset reduced to the first 50 classes (65k images). We train an InceptionResNetV2 [38] at a 227×227 resolution with a batch size of 64 for 100 epochs with a learning rate of 0.01, reducing it by 10 at epochs 50, 80.

Baselines. We introduce here the state-of-the-art approaches we compared with as well as the abbreviations used in the tables. Cross-entropy (CE), dropout (D), and mixup (M) are simple baselines obtained by training with no noise correction and dropout [4] or mixup [49] as regularization. MentorNet [16] (MN) and MentorMix [17] (MM) use teacher networks to weight noisy samples. FaMUS [45] (FaMUS) uses meta learning to learn to correct noisy samples. Bootstrapping [31] (B) corrects noisy samples using a fixed interpolation with pseudo-labels; Dynamic Bootstrapping [3] (DB) expands the idea by correcting only high loss noisy samples retrieved using a beta mixture. The S-model [12] (SM) corrects noisy samples using a noise adaptation layer optimized using an expectation maximization algorithm. DivideMix [22] (DM) uses a Gaussian mixture to detect high loss samples and correct them using a semi-supervised consistency regularization algorithm; the idea is expanded upon in PropMix [10] (PM) where self-supervised initialization is used and only the simplest of the noisy samples are corrected while the hardest are discarded. ScanMix [33] (SM) also improves on DM by correcting the label using a semi-supervised contrastive algorithm together with a semantic clustering in a self-supervised feature space, optimized using an EM algorithm. EvidentialMix [32] (EDM) refines the noisy sample detection of DM to account for OOD samples and uses the evidential loss [34] to evidence separate OOD and ID noisy modes. JoSRC [46] (JoSRC) proposes to use the Jensen-Shannon divergence between a consistency regularization guessed label and the original label to detect noisy samples and further select samples with low agreement between views as OOD. Robust Representation Learning [23] (RRL) trains a weakly supervised prototype objective to promote clean samples to be close to their class prototypes. Finally, Dynamic Softening for Out-of-distribution Samples [1] (DSOS) computes the collision entropy of the interpolation between the original label and network prediction to separate ID noisy and OOD samples.

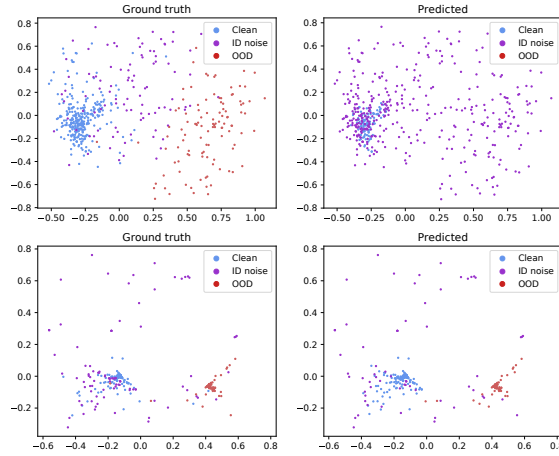


Fig. 3. Feature embedding for class 1 of CIFAR-100 corrupted with $r_i = r_o = 0.2$ (ImageNet32 OOD). The top row presents a 2D visualization obtained using Isomap [39] of the raw contrastive features and the second row presents a Isomap visualization of the embedding E . Embedding the features allows to evidence the OOD cluster

4.2 Clustering the unsupervised features

This section presents the experiments on the linear separability of ID and OOD data in an unsupervised feature space, the importance of embedding the unsupervised features when clustering the noise and the accuracy our noise retrieval algorithm. First, to validate our hypothesis over the separability of ID and OOD samples in the unsupervised feature space, we propose to train a linear classifier on the unsupervised features and evaluate its capacity to linearly separate the two distributions. We observe that the classifier can linearly separate the two distribution with error rates below 3% for synthetically corrupted CIFAR-100 with $r_{in} = r_{out} = 0.2$ and below 1% for miniImagenet corrupted with web noise (CNWL). The linear separability is less accurate when using Places365 as the OOD corruption dataset; we argue that this is because of lower image variability in the dataset, justified by the lower number of classes and the fine-grained nature of the classification task. A table and more visualizations of the separation on the 2D-hypersphere are available in the supplementary material. Second, Figure 3 provides a visualization of the importance of embedding the unsupervised features to perform the noise clustering for a class of CIFAR-100 where we compare applying the clustering algorithm on the raw unsupervised contrastive features against the spectral embedding E . The left column is the ground-truth and the right represent the detection made by the clustering algorithm. We use Isomap [39] to reduce the dimensionality to 2 to be able to visualize the features. The spectral embedding E is essential to evidence the OOD cluster, not originally present in the raw features.

Table 1. Mitigating ID noise and OOD noise on CIFAR-100 corrupted with ImageNet32 or Places365 images. We run all the algorithms using publicly available implementations by authors. We report best and last accuracy. We bold (underline) the highest best (final) accuracy

Corruption	r_{out}	r_{in}	CE	M	DB	JoSRC	ELR	EDM	DSOS	RRL	Ours
INet32	0.2	0.2	63.68/55.52	66.71/62.52	65.61/65.61	67.37/64.17	68.71/68.51	71.03/70.42	70.54/70.54	72.64/72.33	72.95 / <u>72.70</u>
	0.4	0.2	58.94/44.31	59.54/53.16	54.79/54.42	61.70/61.37	63.21/63.07	61.89/61.83	62.49/62.05	66.04/65.44	67.62 / <u>67.14</u>
	0.6	0.2	46.02/26.03	42.87/40.39	42.50/42.50	37.95/37.11	44.79/44.60	21.88/14.59	49.98/49.14	26.76/24.51	53.26 / <u>51.26</u>
	0.4	0.4	41.39/18.45	38.37/33.85	35.90/35.90	41.53/41.44	34.82/34.21	24.15/01.62	43.69/42.88	31.29/30.64	54.04 / <u>52.66</u>
Places365	0.2	0.2	59.88/53.61	66.31/59.69	65.86/65.83	67.06/66.73	68.58/68.45	70.46/70.25	69.72/69.12	72.62 / <u>72.49</u>	71.25/71.14
	0.4	0.2	53.46/42.46	59.75/48.55	55.81/55.61	60.83/60.64	62.66/62.34	61.80/61.55	59.47/59.47	65.82 / <u>65.79</u>	64.03/63.48
	0.6	0.2	39.55/21.42	39.17/33.69	40.75/40.61	39.83/39.63	37.10/36.51	23.67/14.66	35.48/35.41	49.27/49.27	49.83 / <u>49.83</u>
	0.4	0.4	32.06/13.85	34.36/27.63	35.05/34.86	33.23/32.58	34.71/33.86	20.33/11.88	29.54/29.48	26.67/24.34	50.95 / <u>47.61</u>

Table 2. Ablation study on CIFAR-100 corrupted with ImageNet32 with $r_{out} = 0.4$ and $r_{in} = 0.2$. corr = correction and rm = remove

		Embed	Contrastive	Best	Last
No noise corr	CE	✗	✗	58.94	44.31
	+ mixup	✗	✗	59.54	53.16
	+ guided contrastive	✗	✓	62.83	56.29
Noise corr	ID corr only	✗	✗	57.02	55.43
	rm OOD only	✗	✗	60.73	53.88
	ID corr and rm OOD	✗	✗	54.81	54.20
	ID corr only	✓	✗	61.40	58.90
	rm OOD only	✓	✗	60.87	54.08
	ID corr and rm OOD	✓	✗	61.83	61.45
	ID corr only	✓	✓	63.91	62.94
	ID corr and rm OOD	✓	✓	64.51	64.04
Other	OOD corr only	✓	✓	63.41	58.39
	ID + OOD corr	✓	✓	65.22	64.42
	+ equal sampling	✓	✓	67.62	67.14
Other	- mixup	✓	✓	61.66	59.40

4.3 Synthetic noise corruption

We study the capacity of our algorithm to mitigate ID noise and OOD noise on synthetically corrupted version of the CIFAR-100 dataset. Table 1 reports results when using ImageNet32 or Places365 as a OOD corruption. We notice that the OOD corruption using the Places365 dataset is more harmful than corrupting with ImageNet32, especially for high noise levels.

4.4 Ablation study

Table 2 illustrates the importance of each element of the proposed method on CIFAR-100 corrupted with OOD noise from ImageNet32. We study multiple cases including retrieving OOD and ID clusters on the un-embedded raw unsupervised contrastive features (Noise corr without embedding); correcting only the OOD or ID examples while considering the rest clean (ID/OOD corr only); joint effect of the ID and OOD correction (ID + OOD corr); studying the effect of removing the OOD samples from the training set instead of using them in the guided contrastive objective in equation 7 (ID corr and rm OOD); running the algorithm without mixup (- mixup). We point out how important mixup is (especially in the




Fig. 4. Hyper-parameter tuning for OPTICS. We report accuracy results obtained for ID/OOD clustering setting different ξ values in OPTICS

Table 3. Web-corrupted miniImageNet from the CNWL [17] (32×32). We run our algorithm; other results are from [10]. We denote with \star algorithms using an ensemble of networks to predict and with \dagger algorithms using unsupervised initialization. We report best accuracy

Noise level	CE	M	\star DM	MM	FaMUS	$\star\dagger$ SM	$\star\dagger$ PM	Ours
20	47.36	49.10	50.96	51.02	51.42	59.06	61.24	61.56
40	42.70	46.40	46.72	47.14	48.03	54.54	56.22	59.94
60	37.30	40.58	43.14	43.80	45.10	52.36	52.84	54.92
80	29.76	33.58	34.50	33.46	35.50	40.00	43.42	45.62

classification loss) to avoid overfitting to the noise. Figure 4 reports the quality of our robust classification algorithm on ID and OOD clustering for different values of ξ in OPTICS where values inferior to 0.03 lead to the best results. We choose $\xi = 0.01$ for all datasets.

4.5 Results on web-noise

We consider here the controlled noisy web labels (CNWL) dataset, where mini-ImageNet is corrupted with OOD images from web queries. Table 3 reports results when training at resolution 32×32 and Table 4 at resolution 299×299 . Finally, in Table 5 we train on the first 50 classes of the Webvision dataset (mini-Webvision) a real world web-crawled dataset and report top-1 and top-5 accuracy results on the validation set on Webvision and on the test set on the ImageNet1k (ILSVRC12) dataset. Since our algorithm uses only one network and to compare against ensemble methods, we report an additional result where we ensemble two networks trained from different random initializations. We also report results when training for 150 epochs to compare fairly against FaMUS. Our algorithm slightly outperforms the state-of-the-art for top-1 accuracy but more convincingly so for top-5 accuracy on Webvision. Because of the guided contrastive loss, the network learns more generalizable features which reduce the risk of catastrophic classification errors (when the predicted class is completely semantically different from the correct predictions). mini-Webvision in particular proposes fine grained classification on species of birds, amphibians and marine animals which reward generalizable features for top-5 evaluation.

Table 4. Web noise on web corrupted miniImageNet (Red MiniImageNet [17]) trained at a high resolution (299×299). We run our algorithm, other results are from [1]. We report best accuracy

Noise level	CE	D	SM	B	M	MN	MM	DSOS	Ours
0	70.9/68.5	71.8/65.7	71.4/68.4	71.8/68.4	72.8/72.3	71.2/68.9	74.3/73.7	74.52/74.10	74.80/74.60
30	66.1/56.5	66.6/55.0	65.2/56.3	66.6/56.7	66.8/61.8	66.2/64.0	68.3/67.2	69.84/67.86	69.96/69.64
50	60.9/51.7	62.1/50.01	61.3/51.3	62.6/52.5	63.2/58.4	61.7/58.0	63.3/61.8	66.14/65.18	66.48/66.38
80	48.8/39.8	49.5/37.6	49.0/40.6	50.1/40.1	50.7/45.5	49.3/43.4	50.2/48.4	55.26/52.24	55.54/54.96

Table 5. Classification accuracy for the proposed and other state-of-the-art methods. We denote with \star algorithms using an ensemble of networks to predict and with \dagger algorithms using unsupervised initialization.. We train the network on the mini-Webvision dataset and test on the ImageNet 1k test set (ILSVRC12). We bold the best results

		100 epochs										150 epochs	
		M	MM	\star DM	\star ELR+	RRL	\star DSOS	$\star\dagger$ PM	$\star\dagger$ SM	Ours	\star Ours	FaMUS	\star Ours
mini-WebVision	top-1	75.44	76.0	77.32	77.78	77.80	78.76	78.84	80.04	78.16	79.84	79.40	80.24
	top-5	90.12	90.2	91.64	91.68	91.30	92.32	90.56	93.04	92.60	93.64	92.80	93.44
ILSVRC12	top-1	71.44	72.9	75.20	70.29	74.40	75.88	—	75.76	74.20	76.64	77.00	77.12
	top-5	89.40	91.10	90.84	89.76	90.90	92.36	—	92.60	93.32	94.20	92.76	94.32

5 Conclusion

This paper proposes to use the alignment and uniformity principles of unsupervised contrastive learning to detect clean and OOD label noise clusters in an embedded feature space. We show that the unsupervised contrastive features for OOD and ID samples are, to a large extent, linearly separated on the unit hypersphere and compute a fixed neighborhood spectral embedding to reduce differences in cluster densities. We adapt the OPTICS algorithm, ordering samples in a neighbor chain and computing the reachability cost to neighbors. Clusters are evidenced by valleys in the reachability plot and a voting system automatically selects the best cluster assignment at the class level given multiple neighborhood sizes. Once the noise has been identified, we train a robust classifier that corrects the labels of known ID noisy samples using a consistency regularization estimation and uses ID and OOD samples together in an auxiliary guided contrastive objective. We report state-of-the-art results on a variety of noisy datasets including synthetically corrupted versions of CIFAR-100, controlled web noise in miniImageNet, and mini-Webvision as a real-world web-crawled dataset.

Acknowledgments

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under grant number 16/RC/3835 - Vistamilk and 12/RC/2289.P2 - Insight as well as the support of the Irish Centre for High End Computing (ICHEC).

References

1. Albert, P., Ortego, D., Arazo, E., O'Connor, N., McGuinness, K.: Addressing out-of-distribution label noise in webly-labelled data. In: Winter Conference on Applications of Computer Vision (WACV) (2022) [2](#), [3](#), [9](#), [10](#), [14](#)
2. Ankerst, M., Breunig, M.M., Kriegel, H.P., Sander, J.: Optics: Ordering points to identify the clustering structure. *ACM Sigmod record* **28**(2), 49–60 (1999) [6](#)
3. Arazo, E., Ortego, D., Albert, P., O'Connor, N., McGuinness, K.: Unsupervised Label Noise Modeling and Loss Correction. In: International Conference on Machine Learning (ICML) (2019) [10](#)
4. Arpit, D., Jastrzebski, S., Ballas, N., Krueger, D., Bengio, E., Kanwal, M.S., Maharaaj, T., Fischer, A., Courville, A., Bengio, Y., et al.: A closer look at memorization in deep networks. In: International Conference on Machine Learning (ICML) (2017) [10](#)
5. Borgli, H., Thambawita, V., Smedsrud, P.H., Hicks, S., Jha, D., Eskeland, S.L., Randel, K.R., Pogorelov, K., Lux, M., Nguyen, D.T.D., et al.: HyperKvasir, a comprehensive multi-class image and video dataset for gastrointestinal endoscopy. *Scientific data* (2020) [2](#)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A Simple Framework for Contrastive Learning of Visual Representations. In: International Conference on Machine Learning (ICML) (2020) [9](#)
7. Chen, T., Kornblith, S., Swersky, K., Norouzi, M., Hinton, G.: Big self-supervised models are strong semi-supervised learners. *arXiv: 2006.10029* (2020) [9](#)
8. Chrabaszcz, P., Loshchilov, I., Hutter, F.: A downsampled variant of imagenet as an alternative to the cifar datasets. *arXiv: 1707.08819* (2017) [9](#)
9. Coates, A., Ng, A., Lee, H.: An analysis of single-layer networks in unsupervised feature learning. In: International conference on artificial intelligence and statistics (AISTATS) (2011) [2](#)
10. Cordeiro, F.R., Belagiannis, V., Reid, I., Carneiro, G.: PropMix: Hard Sample Filtering and Proportional MixUp for Learning with Noisy Labels. *arXiv: 2110.11809* (2021) [10](#), [13](#)
11. Fort, S., Ren, J., Lakshminarayanan, B.: Exploring the Limits of Out-of-Distribution Detection. In: Advances in neural information processing systems (NeurIPS) (2021) [2](#)
12. Goldberger, J., Ben-Reuven, E.: Training deep neural-networks using a noise adaptation layer. In: International Conference on Learning Representations (ICLR) (2017) [10](#)
13. Hendrycks, D., Mazeika, M., Dietterich, T.: Deep anomaly detection with outlier exposure. In: International Conference on Learning Representations (ICLR) (2019) [2](#)
14. Huang, J., Fang, C., Chen, W., Chai, Z., Wei, X., Wei, P., Lin, L., Li, G.: Trash to Treasure: Harvesting OOD Data with Cross-Modal Matching for Open-Set Semi-Supervised Learning. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2021) [2](#)
15. Hwanjun, S., Minseok, K., Dongmin, P., Jae-Gil, L.: Learning from noisy labels with deep neural networks: A survey. *arXiv: 2007.08199* (2020) [3](#)
16. Jiang, L., Zhou, Z., Leung, T., Li, L., Fei-Fei, L.: MentorNet: Learning Data-Driven Curriculum for Very Deep Neural Networks on Corrupted Labels. In: International Conference on Machine Learning (ICML) (2018) [3](#), [10](#)

17. Jiang, L., Huang, D., Liu, M., Yang, W.: Beyond Synthetic Noise: Deep Learning on Controlled Noisy Labels. In: International Conference on Machine Learning (ICML) (2020) [3](#), [10](#), [13](#), [14](#)
18. Kaiming, H., Xiangyu, Z., Shaoqing, R., Jian, S.: Deep Residual Learning for Image Recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) [1](#), [10](#)
19. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Tech. rep., University of Toronto (2009) [9](#)
20. Krizhevsky, A., Sutskever, I., Hinton, G.: Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems (NeurIPS) (2012) [3](#)
21. Lee, K., Zhu, Y., Sohn, K., Li, C.L., Shin, J., Lee, H.: i-Mix: A Strategy for Regularizing Contrastive Representation Learning. In: International Conference on Learning Representations (ICLR) (2021) [4](#), [9](#)
22. Li, J., Socher, R., Hoi, S.: DivideMix: Learning with Noisy Labels as Semi-supervised Learning. In: International Conference on Learning Representations (ICLR) (2020) [3](#), [10](#)
23. Li, J., Xiong, C., Hoi, S.C.: Learning from noisy data with robust representation learning. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2021) [10](#)
24. Li, W., Wang, L., Li, W., Agustsson, E., Van Gool, L.: WebVision Database: Visual Learning and Understanding from Web Data. arXiv: 1708.02862 (2017) [2](#), [3](#), [4](#), [10](#)
25. Liu, S., Niles-Weed, J., Razavian, N., Fernandez-Granda, C.: Early-Learning Regularization Prevents Memorization of Noisy Labels. In: Advances in Neural Information Processing Systems (NeurIPS) (2020) [3](#)
26. Mingxing, T., Quoc, L.: EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In: International Conference on Machine Learning (ICML) (2019) [1](#)
27. Ng, A.Y., Jordan, M.I., Weiss, Y.: On spectral clustering: Analysis and an algorithm. In: Advances in neural information processing systems (NeurIPS) (2002) [5](#)
28. Ortego, D., Arazo, E., Albert, P., O'Connor, N., McGuinness, K.: Towards Robust Learning with Different Label Noise Distributions. In: International Conference on Pattern Recognition (ICPR) (2020) [3](#)
29. Ortego, D., Arazo, E., Albert, P., O'Connor, N.E., McGuinness, K.: Multi-Objective Interpolation Training for Robustness to Label Noise. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [3](#)
30. Patrini, G., Rozza, A., Krishna Menon, A., Nock, R., Qu, L.: Making Deep Neural Networks Robust to Label Noise: A Loss Correction Approach. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) [3](#)
31. Reed, S., Lee, H., Anguelov, D., Szegedy, C., Erhan, D., Rabinovich, A.: Training deep neural networks on noisy labels with bootstrapping. In: International Conference on Learning Representations (ICLR) (2015) [3](#), [10](#)
32. Sachdeva, R., Cordeiro, F.R., Belagiannis, V., Reid, I., Carneiro, G.: EvidentialMix: Learning with Combined Open-set and Closed-set Noisy Labels. In: IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (2020) [3](#), [10](#)
33. Sachdeva, R., Cordeiro, F.R., Belagiannis, V., Reid, I., Carneiro, G.: ScanMix: Learning from Severe Label Noise via Semantic Clustering and Semi-Supervised Learning. arXiv: 2103.11395 (2021) [3](#), [10](#)
34. Sensoy, M., Kaplan, L., Kandemir, M.: Evidential deep learning to quantify classification uncertainty. In: Advances in Neural Information Processing Systems (NeurIPS) (2018) [10](#)

35. Shi, J., Malik, J.: Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence (TPAMI)* (2000) 5
36. Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. In: *Advances in neural information processing systems (NeurIPS)* (2016) 4
37. Sun, Z., Yao, Y., Wei, X.S., Zhang, Y., Shen, F., Wu, J., Zhang, J., Shen, H.T.: Webly Supervised Fine-Grained Recognition: Benchmark Datasets and An Approach. In: *IEEE/CVF International Conference on Computer Vision (ICCV)* (2021) 3, 4
38. Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A.: Inception-v4, inception-resnet and the impact of residual connections on learning. In: *Association for the Advancement of Artificial Intelligence (AAAI)* (2016) 10
39. Tenenbaum, J.B., Silva, V.d., Langford, J.C.: A global geometric framework for nonlinear dimensionality reduction. *Science* (2000) 11
40. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., Wierstra, D.: Matching Networks for One Shot Learning. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2016) 10
41. Vyas, N., Saxena, S., Voice, T.: Learning Soft Labels via Meta Learning. *arXiv:2009.09496* (2020) 3
42. Wang, T., Isola, P.: Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In: *International Conference on Machine Learning (ICML)* (2020) 5
43. Wang, Y., Liu, W., Ma, X., Bailey, J., Zha, H., Song, L., Xia, S.T.: Iterative Learning With Open-Set Noisy Labels. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018) 5
44. Xiao, T., Xia, T., Yang, Y., Huang, C., Wang, X.: Learning from massive noisy labeled data for image classification. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015) 3, 4
45. Xu, Y., Zhu, L., Jiang, L., Yang, Y.: Faster meta update strategy for noise-robust deep learning. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2021) 3, 10
46. Yao, Y., Sun, Z., Zhang, C., Shen, F., Wu, Q., Zhang, J., Tang, Z.: Jo-SRC: A Contrastive Approach for Combating Noisy Labels. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2021) 2, 3, 10
47. Yu, Q., Aizawa, K.: Unsupervised out-of-distribution detection by maximum classifier discrepancy. In: *IEEE International Conference on Computer Vision (ICCV)* (2019) 2
48. Zhang, C., Bengio, S., Hardt, M., Recht, B., Vinyals, O.: Understanding deep learning requires re-thinking generalization. In: *International Conference on Learning Representations (ICLR)* (2017) 2
49. Zhang, H., Cisse, M., Dauphin, Y., Lopez-Paz, D.: mixup: Beyond Empirical Risk Minimization. In: *International Conference on Learning Representations (ICLR)* (2018) 3, 4, 9, 10
50. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million Image Database for Scene Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2017) 9
51. Zhou, T., Wang, S., Bilmes, J.: Robust Curriculum Learning: From clean label detection to noisy label self-correction. In: *International Conference on Learning Representations (ICLR)* (2020) 3