

The Group Loss for Deep Metric Learning

Ismail Elezi¹, Sebastiano Vascon¹, Alessandro Torcinovich¹, Marcello Pelillo¹,
and Laura Leal-Taixé²

¹ Ca' Foscari University of Venice

² Technical University of Munich

Abstract. Deep metric learning has yielded impressive results in tasks such as clustering and image retrieval by leveraging neural networks to obtain highly discriminative feature embeddings, which can be used to group samples into different classes. Much research has been devoted to the design of smart loss functions or data mining strategies for training such networks. Most methods consider only pairs or triplets of samples within a mini-batch to compute the loss function, which is commonly based on the distance between embeddings. We propose *Group Loss*, a loss function based on a differentiable label-propagation method that enforces embedding similarity across *all* samples of a group while promoting, at the same time, low-density regions amongst data points belonging to different groups. Guided by the smoothness assumption that “similar objects should belong to the same group”, the proposed loss trains the neural network for a classification task, enforcing a consistent labelling amongst samples within a class. We show state-of-the-art results on clustering and image retrieval on several datasets, and show the potential of our method when combined with other techniques such as ensembles. To facilitate further research, we make available the code and the models at https://github.com/dvl-tum/group_loss.

Keywords: Deep Metric Learning, Image Retrieval, Image Clustering

1 Introduction

Measuring object similarity is at the core of many important machine learning problems like clustering and object retrieval. For visual tasks, this means learning a distance function over images. With the rise of deep neural networks, the focus has rather shifted towards learning a feature embedding that is easily separable using a simple distance function, such as the Euclidean distance. In essence, objects of the same class (similar) should be close by in the learned manifold, while objects of a different class (dissimilar) should be far away.

Historically, the best performing approaches get deep feature embeddings from the so-called siamese networks [3], which are typically trained using the contrastive loss [3] or the triplet loss [36,47]. A clear drawback of these losses is that they only consider pairs or triplets of data points, missing key information about the relationships between all members of the mini-batch. On a

mini-batch of size n , despite that the number of pairwise relations between samples is $\mathcal{O}(n^2)$, contrastive loss uses only $\mathcal{O}(n/2)$ pairwise relations, while triplet loss uses $\mathcal{O}(2n/3)$ relations. Additionally, these methods consider only the relations between objects of the same class (positives) and objects of other classes (negatives), without making any distinction that negatives belong to different classes. This leads to not taking into consideration the global structure of the embedding space, and consequently results in lower clustering and retrieval performance. To compensate for that, researchers rely on other tricks to train neural networks for deep metric learning: intelligent sampling [21], multi-task learning [53] or hard-negative mining [35]. Recently, researchers have been increasingly working towards exploiting in a principled way the global structure of the embedding space [31,4,10,44], typically by designing ranking loss functions instead of following the classic triplet formulations.

In a similar spirit, we propose *Group Loss*, a novel loss function for deep metric learning that considers the similarity between all samples in a mini-batch. To create the mini-batch, we sample from a fixed number of classes, with samples coming from a class forming a *group*. Thus, each mini-batch consists of several randomly chosen groups, and each group has a fixed number of samples. An iterative, fully-differentiable label propagation algorithm is then used to build feature embeddings which are similar for samples belonging to the same group, and dissimilar otherwise.

At the core of our method lies an iterative process called replicator dynamics [46,8], that refines the local information, given by the softmax layer of a neural network, with the global information of the mini-batch given by the similarity between embeddings. The driving rationale is that the more similar two samples are, the more they affect each other in choosing their final label and tend to be grouped together in the same group, while dissimilar samples do not affect each other on their choices. Neural networks optimized with the Group Loss learn to provide similar features for samples belonging to the same class, making clustering and image retrieval easier.

Our **contribution** in this work is four-fold:

- We propose a novel loss function to train neural networks for deep metric embedding that takes into account the local information of the samples, as well as their similarity.
- We propose a differentiable label-propagation iterative model to embed the similarity computation within backpropagation, allowing end-to-end training with our new loss function.
- We perform a comprehensive robustness analysis showing the stability of our module with respect to the choice of hyperparameters.
- We show state-of-the-art qualitative and quantitative results in several standard clustering and retrieval datasets.

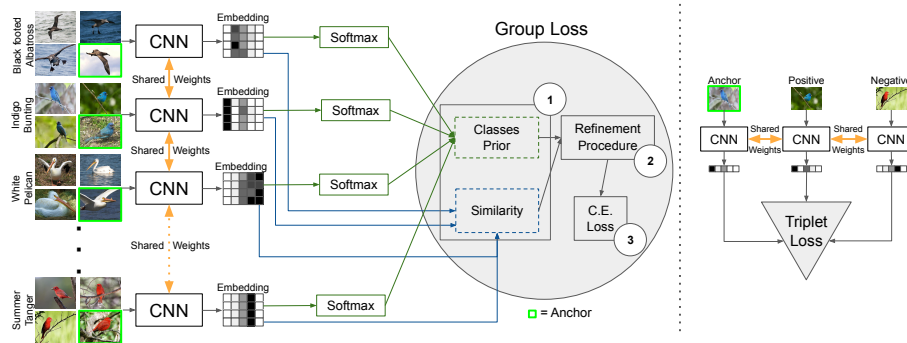


Fig. 1: A comparison between a neural model trained with the Group Loss (left) and the triplet loss (right). Given a mini-batch of images belonging to different classes, their embeddings are computed through a convolutional neural network. Such embeddings are then used to generate a similarity matrix that is fed to the Group Loss along with prior distributions of the images on the possible classes. The green contours around some mini-batch images refer to *anchors*. It is worth noting that, differently from the triplet loss, the Group Loss considers multiple classes and the pairwise relations between all the samples. Numbers from ① to ③ refer to the Group Loss steps, see Sec 3.1 for the details.

2 Related Work

Classical metric learning losses. The first attempt at using a neural network for feature embedding was done in the seminal work of Siamese Networks [3]. A cost function called *contrastive loss* was designed in such a way as to minimize the distance between pairs of images belonging to the same cluster, and maximize the distance between pairs of images coming from different clusters. In [5], researchers used the principle to successfully address the problem of face verification. Another line of research on convex approaches for metric learning led to the triplet loss [36,47], which was later combined with the expressive power of neural networks [35]. The main difference from the original Siamese network is that the loss is computed using triplets (an anchor, a positive and a negative data point). The loss is defined to make the distance between features of the anchor and the positive sample smaller than the distance between the anchor and the negative sample. The approach was so successful in the field of face recognition and clustering, that soon many works followed. The majority of works on the Siamese architecture consist of finding better cost functions, resulting in better performances on clustering and retrieval. In [37], the authors generalized the concept of triplet by allowing a joint comparison among $N - 1$ negative examples instead of just one. [39] designed an algorithm for taking advantage of the mini-batches during the training process by lifting the vector of pairwise distances within the batch to the matrix of pairwise distances, thus enabling the algorithm to learn feature embedding by optimizing a novel structured pre-

diction objective on the lifted problem. The work was later extended in [38], proposing a new metric learning scheme based on structured prediction that is designed to optimize a clustering quality metric, i.e., the normalized mutual information [22]. Better results were achieved on [43], where the authors proposed a novel angular loss, which takes angle relationship into account. A very different problem formulation was given by [17], where the authors used a spectral clustering-inspired approach to achieve deep embedding. A recent work presents several extensions of the triplet loss that reduce the bias in triplet selection by adaptively correcting the distribution shift on the selected triplets [50].

Sampling and ensemble methods. Knowing that the number of possible triplets is extremely large even for moderately-sized datasets, and having found that the majority of triplets are not informative [35], researchers also investigated sampling. In the original triplet loss paper [35], it was found that using semi-hard negative mining, the network can be trained to a good performance, but the training is computationally inefficient. The work of [21] found out that while the majority of research is focused on designing new loss functions, selecting training examples plays an equally important role. The authors proposed a distance-weighted sampling procedure, which selects more informative and stable examples than traditional approaches, achieving excellent results in the process. A similar work was that of [9] where the authors proposed a hierarchical version of triplet loss that learns the sampling all-together with the feature embedding. The majority of recent works has been focused on complementary research directions such as intelligent sampling [21,9,6,45,48] or ensemble methods [49,34,15,24,51]. As we will show in the experimental section, these can be combined with our novel loss.

Other related problems. In order to have a focused and concise paper, we mostly discuss methods which tackle image ranking/clustering in standard datasets. Nevertheless, we acknowledge related research on specific applications such as person re-identification or landmark recognition, where researchers are also gravitating towards considering the global structure of the mini-batch. In [10] the authors propose a new hashing method for learning binary embeddings of data by optimizing Average Precision metric. In [31,11] authors study novel metric learning functions for local descriptor matching on landmark datasets. [4] designs a novel ranking loss function for the purpose of few-shot learning. Similar works that focus on the global structure have shown impressive results in the field of person re-identification [54,1].

Classification-based losses. The authors of [23] proposed to optimize the triplet loss on a different space of triplets than the original samples, consisting of an anchor data point and similar and dissimilar learned proxy data points. These proxies approximate the original data points so that a triplet loss over the proxies is a tight upper bound of the original loss. The final formulation of the loss is shown to be similar to that of softmax cross-entropy loss, challenging the long-held belief that classification losses are not suitable for the task of metric learning. Recently, the work of [52] showed that a carefully tuned normalized softmax cross-entropy loss function combined with a balanced sampling strategy

can achieve competitive results. A similar line of research is that of [55], where the authors use a combination of normalized-scale layers and Gram-Schmidt optimization to achieve efficient usage of the softmax cross-entropy loss for metric learning. The work of [30] goes a step further by taking into consideration the similarity between classes. Furthermore, the authors use multiple centers for class, allowing them to reach state-of-the-art results, at a cost of significantly increasing the number of parameters of the model. In contrast, we propose a novel loss that achieves state-of-the-art results without increasing the number of parameters of the model.

3 Group Loss

Most loss functions used for deep metric learning [35,39,37,38,43,45,44,17,9,21] do not use a classification loss function, e.g., cross-entropy, but rather a loss function based on embedding distances. The rationale behind it, is that what matters for a classification network is that the output is correct, which does not necessarily mean that the embeddings of samples belonging to the same class are similar. Since each sample is classified independently, it is entirely possible that two images of the same class have two distant embeddings that both allow for a correct classification. We argue that a classification loss can still be used for deep metric learning if the decisions do not happen independently for each sample, but rather jointly for a whole *group*, i.e., the set of images of the same class in a mini-batch. In this way, the method pushes for images belonging to the same class to have similar embeddings.

Towards this end, we propose *Group Loss*, an iterative procedure that uses the global information of the mini-batch to refine the local information provided by the softmax layer of a neural network. This iterative procedure categorizes samples into different *groups*, and enforces consistent labelling among the samples of a group. While softmax cross-entropy loss judges each sample in isolation, the Group Loss allows us to judge the overall class separation for *all* samples. In section 3.3, we show the differences between the softmax cross-entropy loss and Group Loss, and highlight the mathematical properties of our new loss.

3.1 Overview of Group Loss

Given a mini-batch \mathcal{B} consisting of n images, consider the problem of assigning a class label $\lambda \in A = \{1, \dots, m\}$ to each image in \mathcal{B} . In the remainder of the manuscript, $X = (x_{i\lambda})$ represents a $n \times m$ (non-negative) matrix of image-label soft assignments. In other words, each row of X represents a probability distribution over the label set A ($\sum_{\lambda} x_{i\lambda} = 1$ for all $i = 1 \dots n$).

Our model consists of the following steps (see also Fig. 1 and Algorithm 1):

- ① **Initialization:** Initialize X , the image-label assignment using the softmax outputs of the neural network. Compute the $n \times n$ pairwise similarity matrix W using the neural network embedding.

- ② **Refinement:** Iteratively, refine X considering the similarities between all the mini-batch images, as encoded in W , as well as their labeling preferences.
- ③ **Loss computation:** Compute the cross-entropy loss of the refined probabilities and update the weights of the neural network using backpropagation.

We now provide a more detailed description of the three steps of our method.

3.2 Initialization

Image-label assignment matrix. The initial assignment matrix denoted $X(0)$, comes from the softmax output of the neural network. We can replace some of the initial assignments in matrix X with one-hot labelings of those samples. We call these randomly chosen samples *anchors*, as their assignments do not change during the iterative refine process and consequently do not directly affect the loss function. However, by using their correct label instead of the predicted label (coming from the softmax output of the NN), they guide the remaining samples towards their correct label.

Similarity matrix. A measure of similarity is computed among all pairs of embeddings (computed via a CNN) in \mathcal{B} to generate a similarity matrix $W \in \mathbb{R}^{n \times n}$. In this work, we compute the similarity measure using the Pearson’s correlation coefficient [28]:

$$\omega(i, j) = \frac{\text{Cov}[\phi(I_i), \phi(I_j)]}{\sqrt{\text{Var}[\phi(I_i)]\text{Var}[\phi(I_j)]}} \quad (1)$$

for $i \neq j$, and set $\omega(i, i)$ to 0. The choice of this measure over other options such as cosine layer, Gaussian kernels, or learned similarities, is motivated by the observation that the correlation coefficient uses data standardization, thus providing invariance to scaling and translation – unlike the cosine similarity, which is invariant to scaling only – and it does not require additional hyperparameters, unlike Gaussian kernels [7]. The fact that a measure of the linear relationship among features provides a good similarity measure can be explained by the fact that the computed features are actually a highly non-linear function of the inputs. Thus, the linear correlation among the embeddings actually captures a non-linear relationship among the original images.

3.3 Refinement

In this core step of the proposed algorithm, the initial assignment matrix $X(0)$ is refined in an iterative manner, taking into account the similarity information provided by matrix W . X is updated in accordance with the *smoothness assumption*, which prescribes that similar objects should share the same label.

To this end, let us define the *support* matrix $\Pi = (\pi_{i\lambda}) \in \mathbb{R}^{n \times m}$ as

$$\Pi = WX \quad (2)$$

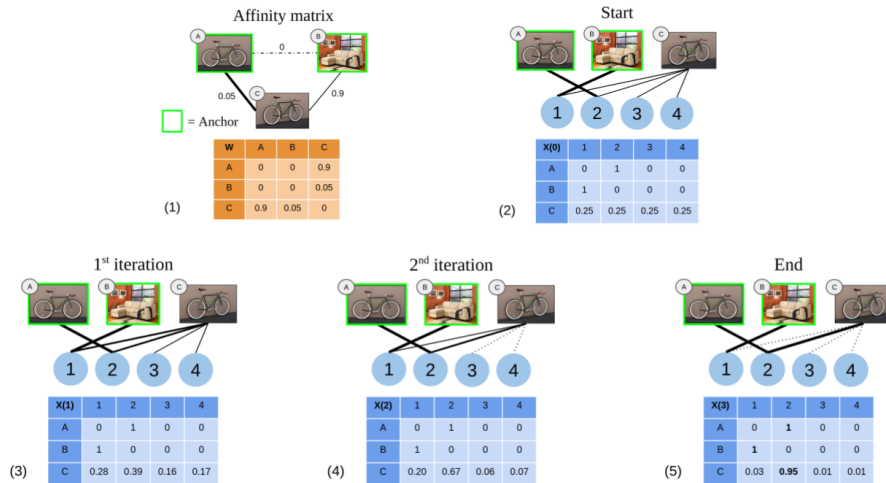


Fig. 2: A toy example of the refinement procedure, where the goal is to classify sample C based on the similarity with samples A and B. (1) The Affinity matrix used to update the soft assignments. (2) The initial labeling of the matrix. (3-4) The process iteratively refines the soft assignment of the unlabeled sample C. (5) At the end of the process, sample C gets the same label of A, (A, C) being more similar than (B, C).

whose (i, λ) -component

$$\pi_{i\lambda} = \sum_{j=1}^n w_{ij} x_{j\lambda} \quad (3)$$

represents the *support* that the current mini-batch gives to the hypothesis that the i -th image in \mathcal{B} belongs to class λ . Intuitively, in obedience to the smoothness principle, $\pi_{i\lambda}$ is expected to be high if images similar to i are likely to belong to class λ .

Given the initial assignment matrix $X(0)$, our algorithm refines it using the following update rule:

$$x_{i\lambda}(t+1) = \frac{x_{i\lambda}(t)\pi_{i\lambda}(t)}{\sum_{\mu=1}^m x_{i\mu}(t)\pi_{i\mu}(t)} \quad (4)$$

where the denominator represents a normalization factor which guarantees that the rows of the updated matrix sum up to one. This is known as multi-population replicator dynamics in evolutionary game theory [46] and is equivalent to non-linear relaxation labeling processes [32,29].

In matrix notation, the update rule (4) can be written as:

$$X(t+1) = Q^{-1}(t)[X(t) \odot \Pi(t)] \quad (5)$$

where

$$Q(t) = \text{diag}([X(t) \odot \Pi(t)] \mathbf{1}) \quad (6)$$

and $\mathbb{1}$ is the all-one m -dimensional vector. $\Pi(t) = WX(t)$ as defined in (2), and \odot denotes the Hadamard (element-wise) matrix product. In other words, the diagonal elements of $Q(t)$ represent the normalization factors in (4), which can also be interpreted as the average support that object i obtains from the current mini-batch at iteration t . Intuitively, the motivation behind our update rule is that at each step of the refinement process, for each image i , a label λ will increase its probability $x_{i\lambda}$ if and only if its support $\pi_{i\lambda}$ is higher than the average support among all the competing label hypothesis Q_{ii} .

Thanks to the Baum-Eagon inequality [29], it is easy to show that the dynamical system defined by (4) has very nice convergence properties. In particular, it strictly increases at each step the following functional:

$$F(X) = \sum_{i=1}^n \sum_{j=1}^n \sum_{\lambda=1}^m w_{ij} x_{i\lambda} x_{j\lambda} \quad (7)$$

which represents a measure of “consistency” of the assignment matrix X , in accordance to the smoothness assumption (F rewards assignments where highly similar objects are likely to be assigned the same label). In other words:

$$F(X(t+1)) \geq F(X(t)) \quad (8)$$

with equality if and only if $X(t)$ is a stationary point. Hence, our update rule (4) is, in fact, an algorithm for maximizing the functional F over the space of row-stochastic matrices. Note, that this contrasts with classical gradient methods, for which an increase in the objective function is guaranteed only when infinitesimal steps are taken, and determining the optimal step size entails computing higher-order derivatives. Here, instead, the step size is implicit and yet, at each step, the value of the functional increases.

3.4 Loss computation

Once the labeling assignments converge (or in practice, a maximum number of iterations is reached), we apply the cross-entropy loss to quantify the classification error and backpropagate the gradients. Recall, the refinement procedure is optimized via *replicator dynamics*, as shown in the previous section. By studying Equation (5), it is straightforward to see that it is composed of fully differentiable operations (matrix-vector and scalar products), and so it can be easily integrated within backpropagation. Although the refining procedure has no parameters to be learned, its gradients can be backpropagated to the previous layers of the neural network, producing, in turn, better embeddings for similarity computation.

3.5 Summary of the Group Loss

In this section, we proposed the Group Loss function for deep metric learning. During training, the Group Loss works by grouping together similar samples

Algorithm 1: The Group Loss

Input: input : Set of pre-processed images in the mini-batch \mathcal{B} , set of labels y ,
neural network ϕ with learnable parameters θ , similarity function ω ,
number of iterations T

- 1 Compute feature embeddings $\phi(\mathcal{B}, \theta)$ via the forward pass
- 2 Compute the similarity matrix $W = [\omega(i, j)]_{ij}$
- 3 Initialize the matrix of priors $X(0)$ from the softmax layer
- 4 **for** $t = 0, \dots, T-1$ **do**
- 5 $Q(t) = \text{diag}([X(t) \odot \Pi(t)] \mathbf{1})$
- 6 $X(t+1) = Q^{-1}(t) [X(t) \odot \Pi(t)]$
- 7 Compute the cross-entropy $J(X(T), y)$
- 8 Compute the derivatives $\partial J / \partial \theta$ via backpropagation, and update the weights θ

based on both the similarity between the samples in the mini-batch and the local information of the samples. The similarity between samples is computed by the correlation between the embeddings obtained from a CNN, while the local information is computed with a softmax layer on the same CNN embeddings. Using an iterative procedure, we combine both sources of information and effectively bring together embeddings of samples that belong to the same class.

During inference, we simply forward pass the images through the neural network to compute their embeddings, which are directly used for image retrieval within a nearest neighbor search scheme. The iterative procedure is not used during inference, thus making the feature extraction as fast as that of any other competing method.

4 Experiments

In this section, we compare the Group Loss with state-of-the-art deep metric learning models on both image retrieval and clustering tasks. Our method achieves state-of-the-art results in three public benchmark datasets.

4.1 Implementation details

We use the PyTorch [27] library for the implementation of the Group Loss. We choose GoogleNet [40] with batch-normalization [12] as the backbone feature extraction network. We pretrain the network on *ILSVRC 2012-CLS* dataset [33]. For pre-processing, in order to get a fair comparison, we follow the implementation details of [38]. The inputs are resized to 256×256 pixels, and then randomly cropped to 227×227 . Like other methods except for [37], we use only a center crop during testing time. We train all networks in the classification task for 10 epochs. We then train the network in the Group Loss task for 60 epochs using RAdam optimizer [18]. After 30 epochs, we lower the learning rate by multiplying it by 0.1. We find the hyperparameters using random search [2]. We use small mini-batches of size 30 – 100. As sampling strategy, on each mini-batch,

we first randomly sample a fixed number of classes, and then for each of the chosen classes, we sample a fixed number of samples.

4.2 Benchmark datasets

We perform experiments on 3 publicly available datasets, evaluating our algorithm on both clustering and retrieval metrics. For training and testing, we follow the conventional splitting procedure [39].

CUB-200-2011 [42] is a dataset containing 200 species of birds with 11,788 images, where the first 100 species (5,864 images) are used for training and the remaining 100 species (5,924 images) are used for testing.

Cars 196 [16] dataset is composed of 16,185 images belonging to 196 classes. We use the first 98 classes (8,054 images) for training and the other 98 classes (8,131 images) for testing.

Stanford Online Products dataset [39], contains 22,634 classes with 120,053 product images in total, where 11,318 classes (59,551 images) are used for training and the remaining 11,316 classes (60,502 images) are used for testing.

4.3 Evaluation metrics

Based on the experimental protocol detailed above, we evaluate retrieval performance and clustering quality on data from unseen classes of the 3 aforementioned datasets. For the retrieval task, we calculate the percentage of the testing examples whose K nearest neighbors contain at least one example of the same class. This quantity is also known as Recall@K [13] and is the most used metric for image retrieval evaluation.

Similar to all other approaches, we perform clustering using K-means algorithm [20] on the embedded features. Like in other works, we evaluate the clustering quality using the Normalized Mutual Information measure (NMI) [22]. The choice of NMI measure is motivated by the fact that it is invariant to label permutation, a desirable property for cluster evaluation.

4.4 Results

We now show the results of our model and comparison to state-of-the-art methods. Our main comparison is with other loss functions, e.g., triplet loss. To compare with perpendicular research on intelligent sampling strategies or ensembles, and show the power of the Group Loss, we propose a simple ensemble version of our method. Our ensemble network is built by training l independent neural networks with the same hyperparameter configuration. During inference, their embeddings are concatenated. Note, that this type of ensemble is much simpler than the works of [51,49,15,25,34], and is given only to show that, when optimized for performance, our method can be extended to ensembles giving higher clustering and retrieval performance than other methods in the literature. Finally, in the interest of space, we only present results for Inception network [40],




Fig. 3: Retrieval results on a set of images from the *CUB-200-2011* (left), *Cars 196* (middle), and *Stanford Online Products* (right) datasets using our Group Loss model. The left column contains query images. The results are ranked by distance. The green square indicates that the retrieved image is from the same class as the query image, while the red box indicates that the retrieved image is from a different class.

as this is the most popular backbone for the metric learning task, which enables fair comparison among methods. In supplementary material, we present results for other backbones, and include a discussion about the methods that work by increasing the number of parameters (capacity of the network) [30], or use more expressive network architectures.

Quantitative results

Loss comparison. In Table 1 we present the results of our method and compare them with the results of other approaches. On the *CUB-200-2011* dataset, we outperform the other approaches by a large margin, with the second-best model (Classification [52]) having circa 6 percentage points (*pp*) lower absolute accuracy in Recall@1 metric. On the NMI metric, our method achieves a score of 69.0 which is 2.8*pp* higher than the second-best method. Similarly, on *Cars 196*, our method achieves best results on Recall@1, with Classification [52] coming second with a 4*pp* lower score. On *Stanford Online Products*, our method reaches the best results on the Recall@1 metric, around 2*pp* higher than Classification [52] and Proxy-NCA [23]. On the same dataset, when evaluated on the NMI score, our loss outperforms any other method, be those methods that exploit advanced sampling, or ensemble methods.

Loss with ensembles. In Table 2 we present the results of our ensemble, and compare them with the results of other ensemble and sampling approaches. Our ensemble method (using 5 neural networks) is the highest performing model in *CUB-200-2011*, outperforming the second-best method (Divide and Conquer [34]) by 1*pp* in Recall@1 and by 0.4*pp* in NMI. In *Cars 196* our method outperforms the second best method (ABE 8 [15]) by 2.8*pp* in Recall@1. The second best method in NMI metric is the ensemble version of RLL [44] which gets outperformed by 2.4*pp* from the Group Loss. In *Stanford Online Products*, our ensemble reaches the third-highest result on the Recall@1 metric (after RLL [44] and GPW [45]) while increasing the gap with the other methods in NMI metric.

Qualitative results

Loss	CUB-200-2011					CARS 196					Stanford Online Products			
	R@1	R@2	R@4	R@8	NMI	R@1	R@2	R@4	R@8	NMI	R@1	R@10	R@100	NMI
Triplet [35]	42.5	55	66.4	77.2	55.3	51.5	63.8	73.5	82.4	53.4	66.7	82.4	91.9	89.5
Lifted Structure [39]	43.5	56.5	68.5	79.6	56.5	53.0	65.7	76.0	84.3	56.9	62.5	80.8	91.9	88.7
Npairs [37]	51.9	64.3	74.9	83.2	60.2	68.9	78.9	85.8	90.9	62.7	66.4	82.9	92.1	87.9
Facility Location [38]	48.1	61.4	71.8	81.9	59.2	58.1	70.6	80.3	87.8	59.0	67.0	83.7	93.2	89.5
Angular Loss [43]	54.7	66.3	76	83.9	61.1	71.4	81.4	87.5	92.1	63.2	70.9	85.0	93.5	88.6
Proxy-NCA [23]	49.2	61.9	67.9	72.4	59.5	73.2	82.4	86.4	88.7	64.9	73.7	-	-	90.6
Deep Spectral [17]	53.2	66.1	76.7	85.2	59.2	73.1	82.2	89.0	93.0	64.3	67.6	83.7	93.3	89.4
Classification [52]	59.6	72	81.2	88.4	66.2	81.7	88.9	93.4	96	70.5	73.8	88.1	95	89.8
Bias Triplet [50]	46.6	58.6	70.0	-	-	79.2	86.7	91.4	-	-	63.0	79.8	90.7	-
Ours	65.5	77.0	85.0	91.3	69.0	85.6	91.2	94.9	97.0	72.7	75.7	88.2	94.8	91.1

Table 1: Retrieval and Clustering performance on *CUB-200-2011*, *CARS 196* and *Stanford Online Products* datasets. Bold indicates best results.




Fig. 4: The effect of the number of anchors and number of samples per class. Fig. 5: The effect of the number of classes per mini-batch.

Fig. 6: Recall@1 as a function of training epochs on Cars196 dataset. Figure adapted from [23].

In Fig. 3 we present qualitative results on the retrieval task in all three datasets. In all cases, the query image is given on the left, with the four nearest neighbors given on the right. Green boxes indicate the cases where the retrieved image is of the same class as the query image, and red boxes indicate a different class. As we can see, our model is able to perform well even in cases where the images suffer from occlusion and rotation. On the *Cars 196* dataset, we see a successful retrieval even when the query image is taken indoors and the retrieved image outdoors, and vice-versa. The first example of *Cars 196* dataset is of particular interest. Despite that the query image contains 2 cars, its four nearest neighbors have the same class as the query image, showing the robustness of the algorithm to uncommon input image configurations. We provide the results of t-SNE [19] projection in the supplementary material.

4.5 Robustness analysis

Number of anchors. In Fig. 4, we show the effect of the number of anchors with respect to the number of samples per class. We do the analysis on *CUB-200-2011* dataset and give a similar analysis for *CARS* dataset in the supplementary material. The results reported are the percentage point differences in terms of

Loss+Sampling	CUB-200-2011					CARS 196					Stanford Online Products			
	R@1	R@2	R@4	R@8	NMI	R@1	R@2	R@4	R@8	NMI	R@1	R@10	R@100	NMI
Samp. Matt. [21]	63.6	74.4	83.1	90.0	69.0	79.6	86.5	91.9	95.1	69.1	72.7	86.2	93.8	90.7
Hier. triplet [9]	57.1	68.8	78.7	86.5	-	81.4	88.0	92.7	95.7	-	74.8	88.3	94.8	-
DAMLRRM [48]	55.1	66.5	76.8	85.3	61.7	73.5	82.6	89.1	93.5	64.2	69.7	85.2	93.2	88.2
DE-DSP [6]	53.6	65.5	76.9	61.7	-	72.9	81.6	88.8	-	64.4	68.9	84.0	92.6	89.2
RLL 1 [44]	57.4	69.7	79.2	86.9	63.6	74	83.6	90.1	94.1	65.4	76.1	89.1	95.4	89.7
GPW [45]	65.7	77.0	86.3	91.2	-	84.1	90.4	94.0	96.5	-	78.2	90.5	96.0	-
Teacher-Student														
RKD [26]	61.4	73.0	81.9	89.0	-	82.3	89.8	94.2	96.6	-	75.1	88.3	95.2	-
Loss+Ensembles														
BIER 6 [24]	55.3	67.2	76.9	85.1	-	75.0	83.9	90.3	94.3	-	72.7	86.5	94.0	-
HDC 3 [51]	54.6	66.8	77.6	85.9	-	78.0	85.8	91.1	95.1	-	70.1	84.9	93.2	-
ABE 2 [15]	55.7	67.9	78.3	85.5	-	76.8	84.9	90.2	94.0	-	75.4	88.0	94.7	-
ABE 8 [15]	60.6	71.5	79.8	87.4	-	85.2	90.5	94.0	96.1	-	76.3	88.4	94.8	-
A-BIER 6 [25]	57.5	68.7	78.3	86.2	-	82.0	89.0	93.2	96.1	-	74.2	86.9	94.0	-
D and C 8 [34]	65.9	76.6	84.4	90.6	69.6	84.6	90.7	94.1	96.5	70.3	75.9	88.4	94.9	90.2
RLL 3 [44]	61.3	72.7	82.7	89.4	66.1	82.1	89.3	93.7	96.7	71.8	79.8	91.3	96.3	90.4
Ours 2-ensemble	65.8	76.7	85.2	91.2	68.5	86.2	91.6	95.0	97.1	72.6	75.9	88.0	94.5	91.1
Ours 5-ensemble	66.9	77.1	85.4	91.5	70.0	88.0	92.5	95.7	97.5	74.2	76.3	88.3	94.6	91.1

Table 2: Retrieval and Clustering performance of our ensemble compared with other ensemble and sampling methods. Bold indicates best results.

Recall@1 with respect to the best performing set of parameters (see $Recall@1 = 64.3$ in Tab. 1). The number of anchors ranges from 0 to 4, while the number of samples per class varies from 5 to 10. It is worth noting that our best setting considers 1 or 2 anchors over 9 samples. Moreover, even when we do not use any anchor, the difference in Recall@1 is no more than $2pp$.

Number of classes per mini-batch. In Fig. 5, we present the change in Recall@1 on the *CUB-200-2011* dataset if we increase the number of classes we sample at each iteration. The best results are reached when the number of classes is not too large. This is a welcome property, as we are able to train on small mini-batches, known to achieve better generalization performance [14].

Convergence rate. In Fig. 6, we present the convergence rate of the model on the *Cars 196* dataset. Within the first 30 epochs, our model achieves state-of-the-art results, making our model significantly faster than other approaches. The other models except Proxy-NCA [23], need hundreds of epochs to converge.

Implicit regularization and less overfitting. In Figures 7 and 8, we compare the results of training vs. testing on *Cars 196* [16] and *Stanford Online Products* [39] datasets. We see that the difference between Recall@1 at train and test time is small, especially on *Stanford Online Products* dataset. On *Cars 196* the best results we get for the training set are circa 93% in the Recall@1 measure, only 7.5 percentage points (pp) better than what we reach in the testing set. From the works we compared the results with, the only one which reports the results on the training set is [17]. They reported results of over 90% in all three datasets (for the training sets), much above the test set accuracy which lies at 73.1% on *Cars 196* and 67.6% on *Stanford Online Products* dataset. [41] also provides results, but it uses a different network.




Fig. 7: Training vs testing Recall@1 curves on *Cars 196* dataset.




Fig. 8: Training vs testing Recall@1 curves on *Stanford Online Products* dataset.

We further implement the P-NCA [23] loss function and perform a similar experiment, in order to be able to compare training and test accuracies directly with our method. In Figure 7, we show the training and testing curves of P-NCA on the *Cars 196* [16] dataset. We see that while in the training set, P-NCA reaches results of $3pp$ higher than our method, in the testing set, our method outperforms P-NCA by around $10pp$. Unfortunately, we were unable to reproduce the results of the paper [23] on *Stanford Online Products* dataset. Furthermore, even when we turn off $L2$ -regularization, the generalization performance of our method does not drop at all. Our intuition is that by taking into account the structure of the entire manifold of the dataset, our method introduces a form of regularization. We can clearly see a smaller gap between training and test results when compared to competing methods, indicating less overfitting.

5 Conclusions and Future Work

In this work, we propose the Group Loss, a novel loss function for metric learning. By considering the content of a mini-batch, it promotes embedding similarity across all samples of the same class, while enforcing dissimilarity for elements of different classes. This is achieved with a differentiable layer that is used to train a convolutional network in an end-to-end fashion. Our model outperforms state-of-the-art methods on several datasets, and shows fast convergence. In our work, we did not consider any advanced sampling strategy. Instead, we randomly sample objects from a few classes at each iteration. Sampling has shown to have a very important role in feature embedding [21]. As future work, we will explore sampling techniques which can be suitable for our module.

Acknowledgements. This research was partially funded by the Humboldt Foundation through the Sofja Kovalevskaja Award. We thank Michele Fenzi, Maxim Maximov and Guillem Braso Andilla for useful discussions.

References

1. Alemu, L.T., Shah, M., Pelillo, M.: Deep constrained dominant sets for person re-identification. In: IEEE/CVF International Conference on Computer Vision, ICCV. pp. 9854–9863 (2019)
2. Bergstra, J., Bengio, Y.: Random search for hyper-parameter optimization. *Journal of Machine Learning Research* **13**, 281–305 (2012)
3. Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., Shah, R.: Signature verification using a” siamese” time delay neural network. In: Advances in Neural Information Processing Systems, NIPS. pp. 737–744 (1994)
4. Çakir, F., He, K., Xia, X., Kulis, B., Sclaroff, S.: Deep metric learning to rank. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR. pp. 1861–1870 (2019)
5. Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. In: IEEE Computer Vision and Pattern Recognition, CVPR. pp. 539–546 (2005)
6. Duan, Y., Chen, L., Lu, J., Zhou, J.: Deep embedding learning with discriminative sampling policy. In: IEEE Computer Vision and Pattern Recognition, CVPR (2019)
7. Elezi, I., Torcinovich, A., Vascon, S., Pelillo, M.: Transductive label augmentation for improved deep network learning. In: International Conference on Pattern Recognition, ICPR. pp. 1432–1437 (2018)
8. Erdem, A., Pelillo, M.: Graph transduction as a noncooperative game. *Neural Computation* **24**(3), 700–723 (2012)
9. Ge, W., Huang, W., Dong, D., Scott, M.R.: Deep metric learning with hierarchical triplet loss. In: European Conference in Computer Vision, ECCV. pp. 272–288 (2018)
10. He, K., Çakir, F., Bargal, S.A., Sclaroff, S.: Hashing as tie-aware learning to rank. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR. pp. 4023–4032 (2018)
11. He, K., Lu, Y., Sclaroff, S.: Local descriptors optimized for average precision. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR. pp. 596–605 (2018)
12. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International Conference on Machine Learning, ICML. pp. 448–456 (2015)
13. Jégou, H., Douze, M., Schmid, C.: Product quantization for nearest neighbor search. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(1), 117–128 (2011)
14. Keskar, N.S., Mudigere, D., Nocedal, J., Smelyanskiy, M., Tang, P.T.P.: On large-batch training for deep learning: Generalization gap and sharp minima. In: International Conference on Learning Representations, ICLR (2017)
15. Kim, W., Goyal, B., Chawla, K., Lee, J., Kwon, K.: Attention-based ensemble for deep metric learning. In: European Conference on Computer Vision. pp. 760–777 (2018)
16. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: International IEEE Workshop on 3D Representation and Recognition (3dRR-13). Sydney, Australia (2013)
17. Law, M.T., Urtasun, R., Zemel, R.S.: Deep spectral clustering learning. In: Proceedings of the 34th International Conference on Machine Learning, ICML. pp. 1985–1994 (2017)

18. Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., Han, J.: On the variance of the adaptive learning rate and beyond. In: International Conference on Learning Representations, ICLR (2020)
19. van der Maaten, L., Hinton, G.E.: Visualizing non-metric similarities in multiple maps. *Machine Learning* **87**(1), 33–55 (2012)
20. MacQueen, J.: Some methods for classification and analysis of multivariate observations. In: Proc. Fifth Berkeley Symp. on Math. Statist. and Prob., Vol. 1. pp. 281–297 (1967)
21. Manmatha, R., Wu, C., Smola, A.J., Krähenbühl, P.: Sampling matters in deep embedding learning. In: IEEE International Conference on Computer Vision, ICCV. pp. 2859–2867 (2017)
22. McDaid, A.F., Greene, D., Hurley, N.J.: Normalized mutual information to evaluate overlapping community finding algorithms. *CoRR* **abs/1110.2515** (2011)
23. Movshovitz-Attias, Y., Toshev, A., Leung, T.K., Ioffe, S., Singh, S.: No fuss distance metric learning using proxies. In: IEEE International Conference on Computer Vision, ICCV. pp. 360–368 (2017)
24. Opitz, M., Waltner, G., Possegger, H., Bischof, H.: BIER - boosting independent embeddings robustly. In: IEEE International Conference on Computer Vision, ICCV. pp. 5199–5208 (2017)
25. Opitz, M., Waltner, G., Possegger, H., Bischof, H.: Deep metric learning with BIER: boosting independent embeddings robustly. *IEEE Trans. Pattern Anal. Mach. Intell.* **42**(2), 276–290 (2020)
26. Park, W., Kim, D., Lu, Y., Cho, M.: Relational knowledge distillation. In: IEEE Computer Vision and Pattern Recognition, CVPR (2019)
27. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch. *NIPS Workshops* (2017)
28. Pearson, K.: Notes on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London* **58**, 240–242 (1895)
29. Pelillo, M.: The dynamics of nonlinear relaxation labeling processes. *Journal of Mathematical Imaging and Vision* **7**(4), 309–323 (1997)
30. Qian, Q., Shang, L., Sun, B., Hu, J., Tacoma, T., Li, H., Jin, R.: Softtriple loss: Deep metric learning without triplet sampling. In: IEEE/CVF International Conference on Computer Vision, ICCV. pp. 6449–6457 (2019)
31. Revaud, J., Almazán, J., Rezende, R.S., de Souza, C.R.: Learning with average precision: Training image retrieval with a listwise loss. In: IEEE/CVF International Conference on Computer Vision, ICCV. pp. 5106–5115 (2019)
32. Rosenfeld, A., Hummel, R.A., Zucker, S.W.: Scene labeling by relaxation operations. *IEEE Trans. Syst. Man Cybern.* **6**, 420–433 (1976)
33. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M.S., Berg, A.C., Li, F.: Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.* **115**(3), 211–252 (2015)
34. Sanakoyeu, A., Tschernezki, V., Büchler, U., Ommer, B.: Divide and conquer the embedding space for metric learning. In: IEEE Computer Vision and Pattern Recognition, CVPR (2019)
35. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR. pp. 815–823 (2015)
36. Schultz, M., Joachims, T.: Learning a distance metric from relative comparisons. In: Advances in Neural Information Processing Systems, NIPS. pp. 41–48 (2003)

37. Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. In: *Advances in Neural Information Processing Systems, NIPS*. pp. 1849–1857 (2016)
38. Song, H.O., Jegelka, S., Rathod, V., Murphy, K.: Deep metric learning via facility location. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. pp. 2206–2214 (2017)
39. Song, H.O., Xiang, Y., Jegelka, S., Savarese, S.: Deep metric learning via lifted structured feature embedding. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. pp. 4004–4012 (2016)
40. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S.E., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. pp. 1–9 (2015)
41. Vo, N., Hays, J.: Generalization in metric learning: Should the embedding layer be embedding layer? In: *IEEE Winter Conference on Applications of Computer Vision, WACV*. pp. 589–598 (2019)
42. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD Birds-200-2011 Dataset. Tech. Rep. CNS-TR-2011-001, California Institute of Technology (2011)
43. Wang, J., Zhou, F., Wen, S., Liu, X., Lin, Y.: Deep metric learning with angular loss. In: *IEEE International Conference on Computer Vision, ICCV*. pp. 2612–2620 (2017)
44. Wang, X., Hua, Y., Kodirov, E., Hu, G., Garnier, R., Robertson, N.M.: Ranked list loss for deep metric learning. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. pp. 5207–5216 (2019)
45. Wang, X., Han, X., Huang, W., Dong, D., Scott, M.R.: Multi-similarity loss with general pair weighting for deep metric learning. In: *IEEE Computer Vision and Pattern Recognition, CVPR* (2019)
46. Weibull, J.: *Evolutionary Game Theory*. MIT Press (1997)
47. Weinberger, K.Q., Saul, L.K.: Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research* **10**, 207–244 (2009)
48. Xu, X., Yang, Y., Deng, C., Zheng, F.: Deep asymmetric metric learning via rich relationship mining. In: *IEEE Computer Vision and Pattern Recognition, CVPR*
49. Xuan, H., Souvenir, R., Pless, R.: Deep randomized ensembles for metric learning. In: *European Conference Computer Vision, ECCV*. pp. 751–762 (2018)
50. Yu, B., Liu, T., Gong, M., Ding, C., Tao, D.: Correcting the triplet selection bias for triplet loss. In: *European Conference in Computer Vision, ECCV*. pp. 71–86 (2018)
51. Yuan, Y., Yang, K., Zhang, C.: Hard-aware deeply cascaded embedding. In: *IEEE International Conference on Computer Vision, CVPR*. pp. 814–823 (2017)
52. Zhai, A., Wu, H.: Classification is a strong baseline for deep metric learning. In: *British Machine Vision Conference BMVC*. p. 91 (2019)
53. Zhang, X., Zhou, F., Lin, Y., Zhang, S.: Embedding label structures for fine-grained feature representation. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. pp. 1114–1123 (2016)
54. Zhao, K., Xu, J., Cheng, M.: Regularface: Deep face recognition via exclusive regularization. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. pp. 1136–1144 (2019)
55. Zheng, X., Ji, R., Sun, X., Zhang, B., Wu, Y., Huang, F.: Towards optimal fine grained retrieval via decorrelated centralized loss with normalize-scale layer. In: *Conference on Artificial Intelligence, AAAI*. pp. 9291–9298 (2019)