

Smart Mining for Deep Metric Learning

*Ben Harwood², *Vijay Kumar B G¹, Gustavo Carneiro¹, Ian Reid¹, Tom Drummond²

¹The University of Adelaide, ²Monash University

{vijay.kumar, gustavo.carneiro, ian.reid}@adelaide.edu.au, {ben.harwood, tom.drummond}@monash.edu

Abstract

To solve deep metric learning problems and producing feature embeddings, current methodologies will commonly use a triplet model to minimise the relative distance between samples from the same class and maximise the relative distance between samples from different classes. Though successful, the training convergence of this triplet model can be compromised by the fact that the vast majority of the training samples will produce gradients with magnitudes that are close to zero. This issue has motivated the development of methods that explore the global structure of the embedding and other methods that explore hard negative/positive mining. The effectiveness of such mining methods is often associated with intractable computational requirements. In this paper, we propose a novel deep metric learning method that combines the triplet model and the global structure of the embedding space. We rely on a smart mining procedure that produces effective training samples for a low computational cost. In addition, we propose an adaptive controller that automatically adjusts the smart mining hyper-parameters and speeds up the convergence of the training process. We show empirically that our proposed method allows for fast and more accurate training of triplet ConvNets than other competing mining methods. Additionally, we show that our method achieves new state-of-the-art embedding results for CUB-200-2011 and Cars196 datasets.

1. Introduction

The development of deep metric learning models for the estimation of effective feature embedding [2, 4, 9, 11, 15, 16, 17, 13, 22, 25, 27, 26] is at the core of many recently proposed computer vision methods [3, 14, 19, 24, 28]. The main advantage of such models lies in their ability to automatically learn metric spaces, where samples from similar classes tend to be close together, while samples from different classes are more likely to be far from each other. The main scenario envisaged for such an approach involves

the presence of an extremely large number of classes (more than 10^5 classes) and low number of samples per class (in $[10^1, 10^2]$), where the implementation of traditional classifiers becomes challenging [19, 13].

Arguably, the most explored deep learning model that can estimate feature embedding is based on triplet networks [6, 24], which are an extension of the siamese network [1]. Triplet networks are composed of three identical convolutional neural networks (ConvNets) that are trained using triplets of samples: an anchor sample, a positive sample of the same class as the anchor, and a negative sample of a different class. The training procedure is based on a loss function that penalises large relative distances between the anchor and positive samples and small relative distances between the anchor and negative samples. Therefore, this training procedure relies on triplets that contain hard positive cases (anchor and positive samples that are far apart) and hard negative cases (anchor and negative samples that are close together). In other words, these hard samples will form triplets that produce gradients with sufficiently large magnitude. Assuming that a training set has N samples, then the set of triplets has complexity size $O(N^3)$, which means that its formation is infeasible even for datasets of modest sizes (e.g., $N = 10^5$). This issue has led to the implementation of importance sampling techniques [14, 16, 24] that stochastically under-samples the set of triplets. Here, their success relies on using enough samples to guarantee that a certain fraction of the hard positives and negatives are available for training¹. Given the high complexity involved in finding hard positive and negative samples, another training procedure has been developed in order to guarantee training samples with large gradient magnitudes: the incorporation of loss functions that take into account the global structure of the embedding space [9, 13, 22].

In this paper, we propose a novel deep metric learning approach that combines a global [9] and a triplet loss [6, 24] computed using training samples acquired from a smart sampling method that has low computational complexity [5]

*Vijay Kumar B G and Ben Harwood contributed equally to this work

¹We have not found a formal study that describes the number of samples used for training versus the fraction of hard positive/negatives.

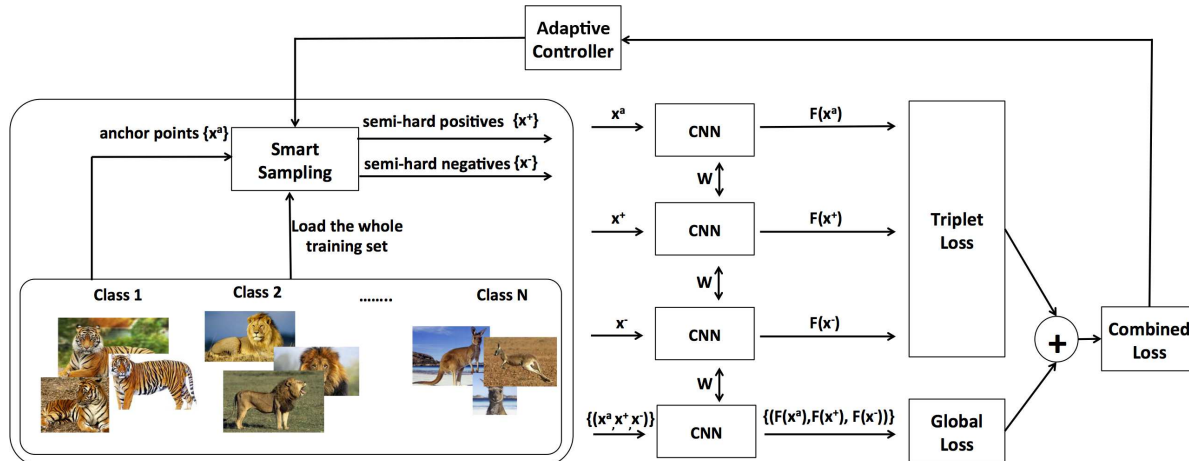


Figure 1. Our proposed deep metric learning model that combines a triplet and a global loss and uses a smart sampling procedure that is capable of quickly searching the entire training set to select effective training samples. The hyper-parameters of the smart sampling are automatically estimated by the proposed adaptive controller with the goal of accelerating the training process.

and can find effective training samples that produce gradients with large magnitude (see Fig.1). Essentially, our smart sampling method circumvents the importance sampling issue mentioned above, enabling our model to be robustly trained with more effective hard negatives and positives, and without the need for a stochastic under-sampling of the training set. Additionally, we propose a novel adaptive controller that accelerates learning by monitoring training performance, estimating its own internal parameters and then automatically adjusting the smart sampling hyper-parameters. We show empirically that our proposed method allows for fast and more accurate training of triplet ConvNets than other competing mining methods. Additionally, we show that our method achieves new state-of-the-art embedding results for CUB-200-2011 and Cars196 datasets.

2. Related Work

In this section, we review recent approaches for selecting hard positives and negatives for training triplet and siamese networks, methods that explore the global structure of the embedding space, and the approximate nearest neighbour search that forms the basis of our proposed method. As pointed out by Shrivastava et al. [15], hard negative and positive mining is a relabeling of the problem of *bootstrapping* [20], where the idea is to start the training of the embedding model with triplets containing positives and negatives that appear to be well separated, and gradually introduce more challenging positive and negative samples as we progress with the training. One of the major issues associated with this approach is on how to introduce such challenging samples - in particular: 1) how to effectively and efficiently sample the training set to select effective training samples, particularly considering that there are N^3 triplets from a training set containing N samples, and 2) what is the

definition of challenging positive and negative samples.

Wang et al. [24] described a way to build triplets based on a manual annotation of sample relevance. Using such relevance, the idea is to use importance sampling to build triplets, but this approach is limited by the fact that it needs those manual annotations. More recently proposed approaches rely on image label, such as the siamese network that gradually introduces hardest possible positive and negative samples [16]. This is achieved by randomly sampling the training set for pairs of anchor and positive samples, and sorting these pairs in descending order with respect to the distance between the two samples in the embedding space. A similar approach is applied for pairs of anchor and negative samples, but the sorting is in ascending order. Then, the training pairs are formed by the top pairs in both lists. We use this sampling scheme as the hard mining baseline. Han et al. [4] introduced an efficient reservoir sampling method to select positive and negative samples, but they do not apply any type of importance sampling to select challenging samples. In FaceNet, Schroff et al. [14] introduced a triplet training approach, where pairs of anchor and positive samples are randomly selected, and pairs of anchor and negative samples are selected from a subset of the training set (i.e., the mini-batch in regular deep learning model training) using a criterion that selects "semi-hard" negatives: pairs of anchor and negative samples are selected if they are close, but at least farther than the anchor-positive pair. This semi-hard negative sampling improves the robustness of training by avoiding overfitting outliers in the training set. An efficient computation of the full matrix of pairwise distances of a subset of the training set (i.e., the mini-batch) allows Song et al. [19] to design of a new loss function that integrates all positive and negative samples to form a lifted structured embedding. However, differently from our work, the lifted

structured embedding only works for the mini-batch instead of the full training set.

The aforementioned issues present in the training of triplet models has motivated the development of approaches that explore the global structure of the embedding space. Kumar et al. [9] proposed a global loss function that uses first and second order statistics to allow for robust training of triplet networks in an approach that improves the training robustness, but still relies on stochastic sampling of positive and negative samples. Ustinova and Lempitsky [22] presented a loss function that minimises the integral of the product between the distribution of negative similarities and the cumulative density function of the positive similarities. Similarly, Song et al. [13] introduced a loss function that optimises a global clustering quality metric (NMI). As shown by Kumar et al. [9], it appears that a combination of local and global losses can produce the most effective embedding spaces, so we believe that the last two approaches mentioned above [22, 13] still have room for improvement, but that improvement depends on more effective hard negative and positive sampling approaches.

In seeking a more effective approach to find hard triplets, we observe that hard negative mining (and to a lesser extent hard positive mining) can be framed as an instance of the well studied approximate nearest neighbour (ANN) search problem. In particular, when mining for negatives we are primarily interested in avoiding the computational cost of exhaustively searching through the entire training set. Fortunately, ANN search methods are able to trade-off a small decrease in nearest neighbour recall for large gains in computational efficiency. In the context of hard negative mining, a small set of nearest neighbours in the current embedding can be guaranteed to contain samples from at least two difference classes (due to training with very few samples per class). A FANNG (Fast Approximate Nearest Neighbour Graph) [5] is a graph based index that can find these neighbourhoods quickly and with a very high rate of recall. Additionally, FANNGs are built in the full embedding space which allows the triplet selection to reuse exact distances that have been computed during the ANN search. FANNG provides state-of-the-art performance at high recall rates while adding only a single tuning parameter for the indexing quality and another for the ANN search quality.

3. Proposed Method

We first describe the architecture of a triplet network [24, 6, 14, 25] and the loss function used in its training. Then, we describe the sampling method applied in the training process. Assume that the training set is represented by $\mathcal{T} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$, with $\mathbf{x}_i \in \mathbb{R}^{n \times n}$ and $y_i \in \{1, \dots, C\}$. The feature embedding is denoted by $f(\mathbf{x}, \theta_f)$, where $f : \mathbb{R}^n \times \mathbb{R}^k \rightarrow \mathbb{R}^m$, with $\theta_f \in \mathbb{R}^k$ representing the network parameters (weight matrices, bias vectors, and normalisation

parameters). The triplet network comprises three identical deep convolutional neural networks (ConvNet) containing L layers, each defined by:

$$f(\mathbf{x}, \theta_f) = f_{\text{out}} \circ r_L \circ h_L \circ f_L \circ \dots \circ r_1 \circ h_1 \circ f_1(\mathbf{x}), \quad (1)$$

where θ_f is defined above, $f_l(\cdot)$ denotes the linear transforms, $h_l(\cdot)$ represents a normalisation function, and $r_l(\cdot)$ denotes a non-linear activation function (e.g., ReLU [12]). Also in (1), note that $\mathbf{f}_l = [\mathbf{f}_{l,1}, \dots, \mathbf{f}_{l,n_l}]$ represents an array of n_l pre-activation functions.

3.1. Triplet Networks

The triplet network [24, 6, 14, 25] is trained with an input composed of an anchor point \mathbf{x}_i (from class y_i), another point from the same class $\mathbf{x}_i^+ = \mathbf{x}_j$ (with $i \neq j$ and $y_i = y_j$), and a point from a different class $\mathbf{x}_i^- = \mathbf{x}_k$ (with $k \neq i$ and $y_i \neq y_k$). The loss function for each triplet is defined by:

$$J^t(\mathbf{x}_i, \mathbf{x}_i^+, \mathbf{x}_i^-, \theta_f) = \max \left(0, 1 - \frac{\|f^{(1)}(\mathbf{x}_i, \theta_f) - f^{(3)}(\mathbf{x}_i^-, \theta_f)\|_2}{\|f^{(1)}(\mathbf{x}_i, \theta_f) - f^{(2)}(\mathbf{x}_i^+, \theta_f)\|_2 + m} \right), \quad (2)$$

where m is the margin, \mathbf{x}_i^+ and \mathbf{x}_i belong to the same class, \mathbf{x}_i^- and \mathbf{x}_i are from different classes, and $f^{(1)}(\cdot)$, $f^{(2)}(\cdot)$ and $f^{(3)}(\cdot)$ are constrained to be the same network parameterized by θ_f .

The training of the triplet network can be made more robust with the introduction of a loss that explores the global structure of the embedding [9]. In particular, the triplet loss in (2) can be extended with a global loss that assumes that the distribution of distances between anchor and positive samples and anchor and negative samples follow a Gaussian distribution. This global loss aims to: 1) minimise the variance of the two distributions, 2) minimise the mean value of the distances between anchor and positive samples, and 3) maximise the mean value of the distances between anchor and negative samples, as follows:

$$J^g(\{\mathbf{x}_i\}_{i=1}^N, \{\mathbf{x}_i^+\}_{i=1}^N, \{\mathbf{x}_i^-\}_{i=1}^N, \theta_f) = (\sigma^{2+} + \sigma^{2-}) + \lambda \max(0, \mu^+ - \mu^- + t), \quad (3)$$

where $\mu^+ = \sum_{i=1}^N d_i^+ / N$, $\mu^- = \sum_{i=1}^N d_i^- / N$, $\sigma^{2+} = \sum_{i=1}^N (d_i^+ - \mu^+)^2 / N$, $\sigma^{2-} = \sum_{i=1}^N (d_i^- - \mu^-)^2 / N$, with μ^+ and σ^{2+} denoting the mean and variance of the matching pair distance distribution, μ^- and σ^{2-} representing the mean and variance of the non-matching pair distance distribution, $d_i^+ = \frac{\|f^{(1)}(\mathbf{x}_i, \theta_f) - f^{(2)}(\mathbf{x}_i^+, \theta_f)\|_2^2}{4}$, $d_i^- = \frac{\|f^{(1)}(\mathbf{x}_i, \theta_f) - f^{(3)}(\mathbf{x}_i^-, \theta_f)\|_2^2}{4}$, λ is a term that balances the importance of each term, t is the margin between the mean of

the matching and non-matching distance distributions and N is the size of the training set. Note in (3), that we assume a triplet network (i.e., $f^{(1)}(\cdot)$, $f^{(2)}(\cdot)$ and $f^{(3)}(\cdot)$ are the same network), where the squared Euclidean distances of the matching and non-matching pairs of the i^{th} triplet are constrained to be $0 \leq d_i^+, d_i^- \leq 1$ because of the division by 4, and the normalisation layer enforces that the norm of the embedding is 1.

3.2. Smart Mining

As discussed in Sec. 2, semi-hard mining has proved an effective method for training triplet networks [14] with the primary aim of finding sets of triplets that will continue to progress the training of the network. Naively, this can be achieved by selecting triplets that provide the greatest violation of the triplet constraint. For instance, given an anchor \mathbf{x}_i , the hardest positive is defined as

$$\mathbf{x}_i^+ = \arg \max_{(\mathbf{x}_j, y_j) \in \mathcal{T}, \mathbf{x}_j \neq \mathbf{x}_i, y_j = y_i} \|f^{(1)}(\mathbf{x}_i, \theta_f) - f^{(2)}(\mathbf{x}_j, \theta_f)\|_2^2, \quad (4)$$

and the hardest negative as

$$\mathbf{x}_i^- = \arg \min_{(\mathbf{x}_j, y_j) \in \mathcal{T}, \mathbf{x}_j \neq \mathbf{x}_i, y_j \neq y_i} \|f^{(1)}(\mathbf{x}_i, \theta_f) - f^{(3)}(\mathbf{x}_j, \theta_f)\|_2^2. \quad (5)$$

In order to avoid the costly $\arg \max$ over the entire training set, semi-hard mining is instead commonly performed over the stochastic subset of samples used in each mini-batch [16, 14]. This method has the additional advantage of avoiding repeated attempts at learning from the hardest triplets that may never improve from epoch to epoch.

We define a novel off-line mining strategy that consists of first finding a set of approximate nearest neighbours $\mathcal{S} \subset \mathcal{T}$. Then, for all triplets with anchor \mathbf{x}_i the set of neighbours \mathcal{S}_i is used to determine appropriate positive and negative samples. To avoid mining poorly structured regions of the embedding, we limit our selection of negative samples to only include negatives where there is at least one positive sample that is closer to the anchor than the negative is. Positive samples are then chosen to guarantee a non-zero response from the loss function (2).

More formally, we define a smart negative as any negative sample $\mathbf{x}_i^- \in \mathcal{S}_i$ such that

$$\|f^{(1)}(\mathbf{x}_i, \theta_f) - f^{(3)}(\mathbf{x}_i^-, \theta_f)\|_2^2 > \kappa \cdot \|f^{(1)}(\mathbf{x}_i, \theta_f) - f^{(2)}(\mathbf{x}_i^{+NN}, \theta_f)\|_2^2, \quad (6)$$

where κ is a global tuning variable and \mathbf{x}_i^{+NN} is the closest positive to \mathbf{x}_i (note that this is not the positive used to form the triplet). The relationship between the exclusion boundary, the anchor, positive samples and negative samples can be seen in Figure 2.

Mining outside the region defined by the distance between an anchor \mathbf{x}_i and the closest positive sample \mathbf{x}_i^{+NN}




Figure 2. A simplified 2-dimensional projection of the neighbours for anchor \mathbf{x}_i , here \mathcal{S}_i contains two positive and four negative samples. Distance $d(x, y)$ is squared Euclidean. a) The current κ and clustering of class y_i specify an exclusion boundary containing all negatives samples, as such none are currently deemed suitable for training. b) On a subsequent epoch a smaller κ and tighter class clustering now yields a negative sample outside the exclusion boundary. This negative and the positive sample further outside the exclusion boundary are used to complete a triplet that is guaranteed to violate the triplet constraint.

ties the selection of negatives to how tightly the class y_i is currently clustered in the embedding space. Additionally, the global parameter κ provides a tunable scaling factor for the radius of these hyper-spherical exclusion boundaries centred on each anchor. Experimentally we have found that the best results are achieved by beginning training with a larger value for κ and then gradually relaxing this constraint throughout the training. This allows previously excluded negatives to be selected for training during later epochs either because the positive neighbours have formed a tighter neighbourhood or the global exclusion value has been sufficiently reduced. The practical details of implementing this mining scheme are discussed below.

3.2.1 Implementing Smart Mining with FANNG

At the beginning of each training epoch we perform a full forward pass of the training set \mathcal{T} to generate the current feature embedding $f(\mathbf{x}, \theta_f)$. The indexing graph used in FANNG [5] is then constructed using the traverse-add algorithm (Alg. 4 in [5]) with the embedding of each element of \mathcal{T} forming a vertex in the graph. At each vertex, a list of outbound edges connect to un-occluded neighbours in a way that approximates the local surface structures of a lower dimensional manifold. Experimental results show that the order of these edge lists remains low (between 15-30 edges) and is independent of the size of the data set and the extrinsic dimensionality of the embedding space. The newly formed traversable graph enables the computationally efficient collection of the approximate nearest neighbour set \mathcal{S} .

As described in [5], the traverse-add algorithm can be repeatedly applied until a specified percentage success rate is reached. Once our target build percentage of 98% is achieved, our approach diverges from the original building process for FANNGs. Rather than applying the backtrack

search (Alg. 3 in [5]) to further refine the graph, we instead use the same backtrack search algorithm to immediately generate the approximate nearest neighbour set \mathcal{S} . Since the graph vertices provide a complete index of the training samples, we can compute each neighbour list \mathcal{S}_i by passing the vertex $f(\mathbf{x}_i, \theta_f)$ to the backtracking search algorithm as both the query vector and starting point for the search. Because the collection of these neighbour lists does not modify the indexing graph, the searches can be performed in parallel. Each query returns a pre-specified number of nearest neighbours sorted in ascending order by distance from the query vertex, as well as the distances themselves. The size of the neighbour lists is selected to guarantee that both positive and negative samples will always be seen in the list.

3.2.2 Triplet Construction

Once \mathcal{S} is computed, the class label information y is used to separate the neighbours into several lists. We perform a single iterative pass over each neighbour list \mathcal{S}_i while maintaining a count of samples from class y_i and a count of all samples from outside that class. Once the first positive sample has been found, the exclusion boundary is computed. Then any future negative samples that satisfy (6) are added to the list of valid negatives. Each subsequent positive sample is added to the list of valid positives along with the current number of valid negatives. With this information we can ensure that a positive sample is not put into a triplet with a later negative sample that is further from the anchor. Lastly, to construct each mined triplet in the current epoch, we take the first unused negative from the list of valid negatives associated with the triplets anchor as well as the first valid positive that is also valid for the chosen negative. Random triplets are used in rare cases where there are no valid negatives. If there are no valid positive samples associated with the chosen negative, then a positive sample is uniformly selected from the set $\mathcal{T} \setminus \mathcal{N}_i$. Algorithm 1 illustrates this triplet selection process in pseudo-code. It is important to note that while each negative is used no more than once for a given anchor during any given epoch, positive samples can be used multiple times with the same anchor. However, the unique negatives will always ensure that no triplets are repeated. In general, our method will select softer negative and positive samples ahead of harder options.

3.2.3 Runtime Complexity

A naive hard mining algorithm that selects $O(N)$ triplets will have a worst case complexity of $O(N^3)$ on any given epoch. Assuming that the samples are equally distributed between C classes, then the complexity can instead be expressed as $O(N \frac{N}{C} \left(N - \frac{N}{C}\right))$. As $C \rightarrow N$, this com-

Algorithm 1: Triplet Selection

Input: Training samples X , Nearest neighbours \mathcal{S} ,
Class labels y , Scale parameter κ

Output: Mined triplets T

```

1 for each sorted neighbour list  $s_i$  do
2    $neg \leftarrow$  empty list of negatives
3    $pos \leftarrow$  empty list of positives/valid negative range
4   for each neighbour  $s_i[j]$  of sample  $x_i$  do
5     if isEmpty( $pos$ ) then
6       if class( $s_i[j]$ )  $\neq y_i$  then
7         continue
8        $bound \leftarrow \kappa \cdot \text{distance}(x_i, s_i[j])$ 
9        $pos.add(s_i[j], \emptyset)$ 
10      continue
11     if distance( $x_i, s_i[j]$ )  $< bound$  then
12       continue
13     if class( $s_i[j]$ )  $\neq y_i$  then
14        $neg.add(s_i[j])$ 
15     if class( $s_i[j]$ )  $= y_i$  then
16        $pos.add(s_i[j], \text{clone}(neg))$ 
17   for each triplet  $t[j]$  with anchor  $x_i$  do
18     if isEmpty( $neg$ ) then
19        $t[j] \leftarrow$  random triplet
20       continue
21      $t[j] \leftarrow x_i, neg[0], \text{random positive } \notin pos$ 
22     for each positive  $pos[k]$  do
23       if  $neg[0] \notin \text{validRange}(pos[k])$  then
24         continue
25        $t[j] \leftarrow x_i, neg[0], pos[k]$ 
26       break
27      $neg.remove(neg[0])$ 
28 return  $T$ 

```

plexity reduces to a best case of $O(N^2)$.

The smart mining algorithm requires the construction of a nearest neighbour index. Exhaustive index construction is $O(N^3)$ due to the sorting of all N^2 pairwise distances. However, we can guarantee a worst case complexity of $O(N^2)$ by instead building an approximation of the index. Using this index to find negatives up to the closest positive sample for each anchor can be performed with worst case complexity $O(N^2)$ regardless of class distribution. Given that $O(N^2)$ is the best case complexity for the naive hard mining approach above, we can conclude that our method is computationally more efficient.

For semi-hard mining, such as [14], algorithmic complexity is reduced by limiting triplet selection to a brute force search within each minibatch. Given an epoch with

M minibatches, the argmax for each anchor results in a total complexity of $O(M \left(\frac{N}{M}\right)^2)$, or simply $O(\frac{N^2}{M})$. For comparative purposes, we note that larger minibatches (i.e. smaller M) tend to reduce training error [16] up until performance begins to be limited by the naive use of argmax. Even so, as $M \rightarrow 1$ the semi-hard mining complexity approaches $O(N^2)$ and the information available in each minibatch also approaches that of both the naive and smart mining.

3.2.4 Automatic Parameter Selection

Up to this point, running our mining scheme requires hand tuning for the hyper-parameter κ . We propose a more robust solution that closes the loop on the triplet mining and training losses. At the beginning of each epoch, we would like to estimate what value of κ will produce triplets of a suitable difficulty for the current network. One such goal could be to ensure that the error from the training set is consistent with the current error of the validation set. We estimate κ with a simple linear model

$$\kappa = \alpha \mathbf{e} + \beta, \tag{7}$$

that finds the least-squares solution for internal parameters α and β from a vector of recent training errors \mathbf{e} and their associated κ . Once we have computed the internal parameters, we can obtain the estimated value

$$\kappa = \alpha e_t + \beta, \tag{8}$$

by providing the current target error e_t . The model is initialised at the beginning of the third training epoch with an initial estimate of the internal parameters. At the beginning of the triplet mining on each subsequent epoch, the training results from the previous epoch are used to update the model. The inclusion of as little as 2% mined triplets per batch is enough to control the training losses.

As training progresses and the embedding improves, it is expected that both the training and validation error will decrease. Targeting a low training error will guarantee that most of the next epoch is spent on triplets that will not make a significant impact on the training. So instead, we can deliberately separate the training and validation errors so that the training error is kept high, while the validation error continues to decrease. To achieve this, we replace the use of the current validation error with a constant value that represents our target training error. Experimental results have shown that a target of between 50% and 75% training error is capable of producing more accurate embeddings in far fewer epochs. To maintain a high training error, it is best to use batches that are 50% to 100% mined triplets.

A comparison of hand tuned and adaptive parameter selection can be seen in figure 3. Training error gives an indi-




Figure 3. A comparison of training performance using hand tuned and adaptive selection of κ . Training and validation error is shown for the first 20 epochs while training on CUB- 200-2011 [23].

cation of the fraction of each batch that is producing a non-zero gradient and so can continue to shape the embedding space. The validation error is produced by evaluating the embedding with a reserved set of samples not used for training and is used as an inverse measure of the current quality of the embedding. Since the adaptive method is able to select harder triplets, while avoiding triplets that are so hard that the embedding structure could be damaged, we see that it can produce a higher quality embedding. Additionally, the steeper decent of the adaptive validation indicates that these results can be reached while also using fewer training epochs. In practice when using GPU accelerated code, our triplet selection accounts for less than 1% of the total epoch runtime (the majority of the cost being the forward and backward propagation of the selected triplets). As such, the ability to produce high quality embeddings while converging in comparatively fewer epochs will greatly reduce the overall training time.

4. Experiments

For the experiments, we follow the protocol used in previous papers [18, 19, 13], which uses unseen classes from the CUB- 200-2011 [23] and Cars196 [8] datasets in order to assess the clustering quality and k nearest neighbour retrieval [7]. Our proposed method combining **triplet** and **global** losses using **FANNNG** [5] with and without automated hyper-parameter selection (i.e., the **adaptive** controller) is compared with the following state-of-the-art deep metric learning approaches: (1) **triplet** learning with **semi-hard** negative mining [14] (with and without **FANNNG** [5]), (2) **lifted structured** embedding [19], (3) **N-pairs** metric loss [18], (4) **clustering** [13], and (5) **triplet** combined with **global** loss [9]. For the approaches (1), (2), (3) and (4) above, we report the results from Song et al.’s paper [13]. For the remaining approaches (i.e. our proposed method, and (5)), we use the same training and test set split described in [19] across all datasets. Specifically, the means

CUB200-2011 [23] has 11,788 images of 200 bird species, from which we take the first 100 species for training and use the remaining 100 species for testing. Cars196 [8] has 16,185 images from 196 car models, from which we take the first 98 classes for training and use the remaining 98 for testing. For all our experiments, we initialize the network with pre-trained GoogLeNet [21] weights and randomly initialize the final fully connected layer similar to [19]. We set the embedding size to 64 [19] and the learning rate for the randomly initialized fully connected layer is multiplied by 10 to achieve faster convergence similar to [19].

For the experiments using triplet combined with global loss [9] and for our proposed approach, we let the training procedure run for a maximum of 20 epochs or until convergence (if fewer epochs were required). During the first two epochs, triplet mining was completely disabled to allow for batches comprised of only random triplets. Similar to [14, 9], we set the margin for the triplet and global loss to 0.2 and 0.01 respectively. We start experiment with an initial learning rate of 0.1 and gradually decrease it by a factor of 2 after every 3 epochs. We use a weight decay of 0.0005 for all of our experiments.

4.1. Quantitative Results

Here we report quantitative results based on the normalised mutual information (NMI) [10] score defined by the ratio of mutual information and product of entropies for two clustering assignments - this measures the label agreement between these two clustering assignments (ignoring the permutations). We also report the k nearest neighbour performance using the Recall@K metric [13].

Tables 1 and 2 show the NMI and k nearest neighbour performance with the Recall@K metric results defined above comparing our method to the state of the art for the datasets CUB-200-2011 [23] and Cars196 [8]. From these tables, we can first see that **Triplet + FANNG** significantly improves upon the **Semi-hard** [14] results with respect to all measures, and showing that the smart mining process using FANNGs is more effective than the more commonly used stochastic under-sampling of the training set. The combination of **Triplet + FANNG + Global** shows gains with respect to all measures, when compared to **Triplet + Global** and **Triplet + FANNG**, demonstrating the importance of both the smart mining process and the use of a global loss. The final model **Triplet + FANNG + Global + Adaptive** shows competitive results with respect to all measures as well as a much faster convergence rate (see Fig.4). For instance, for the CUB-200-2011 dataset [23], **Triplet + FANNG + Global + Adaptive** converges in just four epochs, while **Triplet + FANNG + Global** takes 20 epochs to converge. Similarly for Cars196 [8], **Triplet + FANNG + Global + Adaptive** converges in just four epochs, while **Triplet + FANNG + Global** takes 20 epochs

Table 1. Clustering and recall performance on CUB-200-2011 [23]. Our proposals are highlighted.

Method	NMI	R@1	R@2	R@4	R@8
Semi-hard [14]	55.38	42.59	55.03	66.44	77.23
Lifted Structure [19]	56.50	43.57	56.55	68.59	79.63
N-pairs [18]	57.24	45.37	58.41	69.51	79.49
Triplet + Global [9]	58.61	49.04	60.97	72.33	81.85
Clustering [13]	59.23	48.18	61.44	71.83	81.92
Triplet + FANNG	58.10	45.90	57.65	69.63	79.83
Triplet + FANNG + Global	60.09	49.44	61.60	73.09	82.85
Triplet + FANNG + Global + Adaptive	59.90	49.78	62.34	74.05	83.31

Table 2. Clustering and recall performance on Cars196 [8]. Our proposals are highlighted

Method	NMI	R@1	R@2	R@4	R@8
Semi-hard [14]	53.35	51.54	63.78	73.52	82.41
Lifted Structure [19]	56.88	52.98	65.70	76.01	84.27
N-pairs [18]	57.79	53.90	66.76	77.75	86.35
Triplet + Global [9]	58.20	61.41	72.51	81.75	88.39
Clustering [13]	59.04	58.11	70.64	80.27	87.81
Triplet + FANNG	58.24	56.11	68.34	77.99	85.92
Triplet + FANNG + Global	59.70	64.20	75.22	83.24	88.94
Triplet + FANNG + Global + Adaptive	59.50	64.65	76.20	84.23	90.19




Figure 4. A comparison of the convergence rate of our different methods using Recall@1 on CUB-200-2011 dataset (left) and Cars196 (right).

to converge. The accelerated rate of convergence is only achievable when the difficulty of the mined triplets is targeted at the right level for each individual epoch.

4.2. Qualitative Results

Figures 5 and 6 show triplets for visual inspection. The first column of each figure contains randomly selected anchor points from the training set. Each row then contains



Figure 5. a) Triplets mined from the CUB-200-2011 [23] training set using FANNG [5]. b) Random triplets constructed with the same anchor.

the positive and negative sample images that complete each of the triplets. For each of the mined triplets, the negative sample is guaranteed to be a shorter distance from the anchor when compared to the positive sample in the embedding being mined. As per our smart mining algorithm, each of the mined positives is the closest possible positive to the anchor, while still maintaining distance relationships. These properties can be clearly seen when the mined triplets are compared to a randomly generated triplet. The anchors of the mined triplets appear to have stronger similarities with the negative samples, while the random triplet anchors are much closer in appearance to the positive samples. While the mined positive samples are dissimilar from the anchors, in many cases it appears that they are still sharing more features with the anchor than the random positives are sharing with the same anchor. By presenting difficult (but not impossible) triplets more often, our mined triplets enable faster learning of the embedding.

5. Conclusion

From the results in Tables 1-2, we see that **Triplet + FANNG + Global + Adaptive** significantly outperforms the current state of the art methods [9, 13] in terms of clustering and recall performances. Furthermore, it is worth noting that **Triplet + FANNG** performs substantially better than its counterpart **Semi-hard** [14] with respect to the clustering and recall performances, thus highlighting the importance of the smart mining process. Comparing **Triplet**

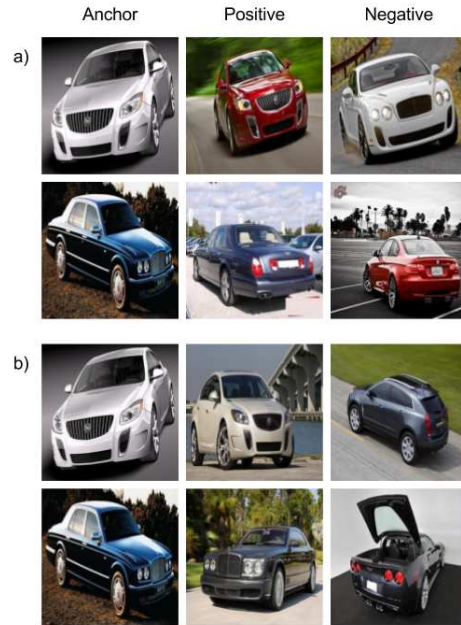


Figure 6. a) Triplets mined from the Cars196 [8] training set using FANNG [5]. b) Random triplets constructed with the same anchor.

+ **FANNG + Global** and **Triple + FANNG**, we can conclude that the global loss is indeed an important component in improving the clustering and recall performance of the embedding. Finally, **Triplet + FANNG + Global + Adaptive** and **Triplet + FANNG + Global** show almost equally strong results, but the former has a significantly faster training process.

In this paper, we proposed a novel triplet-based deep metric learning approach that combines a global structure loss with a triplet loss. We rely on a smart mining process to train our approach, which allows an effective selection of training samples at a low computational cost. Furthermore, we also extend this smart mining with an adaptive controller that automatically selects its hyper-parameters throughout training. By searching the entire training set, we pay a high upfront cost, but make good use of the extra available information to ultimately improve the convergence rate of the training process without compromising on the quality of the embedding. Using CUB-200-2011 [23] and Cars196 [8], we show that our proposed method achieves fast and more accurate training of triplet ConvNets than other competing mining methods. Our approach sets new state-of-the-art deep metric learning results for these two datasets.

Acknowledgements: This research was supported by the Australian Research Council through the Centre of Excellence in Robotic Vision, CE140100016, and through Laureate Fellowship FL130100102 to IDR. We would like thank Guosheng Lin and Chunhua Shen for their insightful discussions.

References

- [1] J. Bromley, J. W. Bentz, L. Bottou, I. Guyon, Y. LeCun, C. Moore, E. Säckinger, and R. Shah. Signature verification using a siamese time delay neural network. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(04):669–688, 1993.
- [2] A. Dosovitskiy, J. T. Springenberg, M. Riedmiller, and T. Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 766–774, 2014.
- [3] M. Guillaumin, J. Verbeek, and C. Schmid. Multiple instance metric learning from automatically labeled bags of faces. In *Computer Vision–ECCV 2010*, pages 634–647. Springer, 2010.
- [4] X. Han, T. Leung, Y. Jia, R. Sukthankar, and A. C. Berg. Matchnet: Unifying feature and metric learning for patch-based matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3279–3286, 2015.
- [5] B. Harwood and T. Drummond. Fanng: Fast approximate nearest neighbour graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5713–5722, 2016.
- [6] E. Hoffer and N. Ailon. Deep metric learning using triplet network. *arXiv preprint arXiv:1412.6622*, 2014.
- [7] H. Jegou, M. Douze, and C. Schmid. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1):117–128, 2011.
- [8] J. Krause, M. Stark, J. Deng, and L. Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 554–561, 2013.
- [9] B. Kumar, G. Carneiro, and I. Reid. Learning local image descriptors with deep siamese and triplet convolutional networks by minimising global loss functions. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [10] C. D. Manning, P. Raghavan, H. Schütze, et al. *Introduction to information retrieval*, volume 1. Cambridge university press Cambridge, 2008.
- [11] J. Masci, D. Migliore, M. M. Bronstein, and J. Schmidhuber. Descriptor learning for omnidirectional image matching. In *Registration and Recognition in Images and Videos*, pages 49–62. Springer, 2014.
- [12] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, pages 807–814, 2010.
- [13] H. Oh Song, S. Jegelka, V. Rathod, and K. Murphy. Deep metric learning via facility location. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [14] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 815–823, 2015.
- [15] A. Shrivastava, A. Gupta, and R. Girshick. Training region-based object detectors with online hard example mining. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [16] E. Simo-Serra, E. Trulls, L. Ferraz, I. Kokkinos, P. Fua, and F. Moreno-Noguer. Discriminative learning of deep convolutional feature point descriptors. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 118–126, 2015.
- [17] K. Simonyan, A. Vedaldi, and A. Zisserman. Learning local feature descriptors using convex optimisation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2014.
- [18] K. Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in Neural Information Processing Systems*, pages 1849–1857, 2016.
- [19] H. O. Song, Y. Xiang, S. Jegelka, and S. Savarese. Deep metric learning via lifted structured feature embedding. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [20] K.-K. Sung. Learning and example selection for object and pattern detection. 1996.
- [21] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [22] E. Ustinova and V. Lempitsky. Learning deep embeddings with histogram loss. In *Advances in Neural Information Processing Systems*, pages 4170–4178, 2016.
- [23] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- [24] J. Wang, Y. Song, T. Leung, C. Rosenberg, J. Wang, J. Philbin, B. Chen, and Y. Wu. Learning fine-grained image similarity with deep ranking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1386–1393, 2014.
- [25] P. Wohlhart and V. Lepetit. Learning descriptors for object recognition and 3d pose estimation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [26] Y. Yuan, K. Yang, and C. Zhang. Hard-aware deeply cascaded embedding. *CoRR*, <http://arxiv.org/abs/1611.05720>, 2016.
- [27] S. Zagoruyko and N. Komodakis. Learning to compare image patches via convolutional neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [28] B. Zhuang, G. Lin, C. Shen, and I. Reid. Fast training of triplet-based deep binary embedding networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.