

LoOp: Looking for Optimal Hard Negative Embeddings for Deep Metric Learning

Bhavya Vasudeva^{1*} Puneesh Deora^{1*} Saumik Bhattacharya² Umapada Pal¹ Sukalpa Chanda³
¹Indian Statistical Institute, Kolkata, India ²Indian Institute of Technology, Kharagpur, India
³Østfold University College, Halden, Norway

Abstract

Deep metric learning has been effectively used to learn distance metrics for different visual tasks like image retrieval, clustering, etc. In order to aid the training process, existing methods either use a hard mining strategy to extract the most informative samples or seek to generate hard synthetics using an additional network. Such approaches face different challenges and can lead to biased embeddings in the former case, and (i) harder optimization (ii) slower training speed (iii) higher model complexity in the latter case. In order to overcome these challenges, we propose a novel approach that looks for optimal hard negatives (LoOp) in the embedding space, taking full advantage of each tuple by calculating the minimum distance between a pair of positives and a pair of negatives. Unlike mining-based methods, our approach considers the entire space between pairs of embeddings to calculate the optimal hard negatives. Extensive experiments combining our approach and representative metric learning losses reveal a significant boost in performance on three benchmark datasets¹.

1. Introduction

Deep metric learning tries to learn an embedding space, where closeness between embeddings encodes the level of semantic similarities between the data samples. This is done by leveraging deep neural networks to learn the mapping between the data samples and the embedding space and enforcing the embeddings belonging to the same class to lie close while pushing those belonging to different classes further apart. Deep metric learning-based methods have achieved state-of-the-art (SOTA) results for several tasks, like face recognition [2, 24, 20, 8], re-identification [35, 19, 41, 30], image retrieval [31, 7], etc.

To train the deep networks, several loss functions with desirable properties have been formulated, which enable the

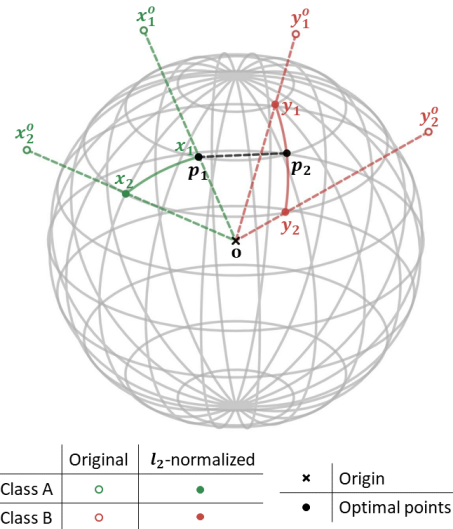


Figure 1: Illustration of the proposed approach. Given two pairs of points in the embedding space, x_1, x_2 from class A, and y_1, y_2 from class B, our method finds the optimal hard negatives by calculating the minimum distance between the curves joining these points.

neural network to learn the mapping from the data space to the embedding space. Conventional methods, such as contrastive loss [2, 9] which takes tuples with two samples and triplet loss [24, 33] which takes three sample tuples, consider the similarity between few samples. In contrast, methods like lifted structure loss [25] aim to exploit all the samples present in a batch to learn more informative representations. Other methods [28, 21] that utilize several samples in the loss have also been proposed.

However, even when considering the whole batch, not all the samples are able to contribute to the loss term. This is because many of them already satisfy the constraints present in the loss. As a result, these samples are not sufficiently informative and lead to low gradient values. To overcome this problem, the idea of using hard samples has

*Equal contribution

¹Code available at <https://github.com/puneesh00/LoOp>

been considered, *i.e.* positive samples which lie far away and negative samples which lie closer. Several works on hard negative mining [21, 11, 13, 37, 24] have been proposed. Mining-based strategies generally look for the most informative samples in the dataset and are prone to learning biased mappings, which do not generalize well to the entire dataset. On the other hand, hard negative generation-based methods [5, 40, 39] utilize an additional sub-network, either an autoencoder or a generative adversarial network [6] as the generator. This can result in harder optimization [1] and an increase in training time and computation.

To address the limitations of existing hard negative mining and generation methods, we propose a novel approach, which looks for optimal hard negative samples (LoOp) in the embedding space. This is done by finding points that minimize the distance between the curves joining a pair of positives and a pair of negatives, as shown in Fig 1. The curve joining a pair of points belonging to the same class lies in the region belonging to that class. Hence, finding the minimum distance with another curve, which is representative of some other class, allows us to consider the most informative pair of samples for computing the loss, as described in Section 3.1. Unlike mining-based methods, our approach does not neglect any samples, and unlike generation-based methods, it does not increase the training complexity or optimization difficulty. It can be easily integrated with various metric learning-based losses. We include several experimental results in Section 4 to demonstrate the efficacy of our approach.

Contributions We propose a novel approach, LoOp, which finds optimal hard negatives in the embedding space, and maximizes the contribution of each tuple in computing the pair-based metric learning loss. Our approach presents the solution to a general problem, namely finding the minimum distance between two bounded curves, which can be useful in other applications as well. Generating hard negatives in the embedding space, our approach utilizes all training samples and does not rely on a subset of samples like mining-based methods. It also avoids the computational load and training complexity introduced by generation-based methods. We also explore the optimality of LoOp with a gradient-based theoretical analysis of the loss function. Our approach generalizes well to various metric learning losses, and it can be easily combined without any increase in optimization difficulty or additional parameters. It outperforms state-of-the-art mining and generation-based methods on three benchmark datasets, Cars196 [17], CUB-200-2011 [29], and Stanford Online Products [25].

2. Related Work

Hard Negative Mining The samples which satisfy the criteria used for training do not contribute to the loss and lead to smaller gradients and slower convergence. To

combat this issue, several sampling and mining strategies [21, 11, 13, 37, 24] have been developed, which use hard positives and negatives for training. Schroff *et al.* [24] propose ‘semi-hard’ triplet mining, where the negatives, which lie close to the anchor but farther than the positives, are selected. Harwood *et al.* [11] present a smart mining-based method that adaptively selects the most informative samples. The drawback of such methods is that the trained models overfit on a subset of hard samples and underfit on the ‘easy’ samples. As they rely on a subset of training samples, they fail to exploit the information other samples provide.

Hard Negative Generation In contrast to mining-based methods, generation-based methods [5, 40, 39] try to exploit all the training samples by extracting the semantic information required to generate synthetic samples which act as hard negatives. Deep adversarial metric learning (DAML) [5] trains the embedding network and the hard negative generator network in an adversarial manner to generate hard negatives. In hardness-aware deep metric learning (HDML) [40], synthetic samples are created via interpolation in the embedding space. This is followed by finding the corresponding label-preserving mappings in the feature space using an autoencoder. They also control the level of hardness while training. Although these methods utilize all the training samples, they require an additional network as the generator. This increases the training time and computational load and can also lead to optimization difficulty [1].

Embedding Space Augmentation Some methods [36, 16] augment the embedding space directly to obtain useful synthetic samples. Yin *et al.* [36] assume that all classes follow a Gaussian distribution and translate samples from different classes about their means to generate new ones. Embedding expansion (EE) [16] is a linear interpolation-based method to generate synthetic samples and mine for hard negatives by considering the pair-wise distances between real and synthetic points of different classes. It exhibits a trade-off between the quality of hard negatives and the computational load of calculating the pair-wise distances.

Other works try to supplement the metric learning losses by introducing a regularization to optimize the direction of displacement of samples [23] or maximize the spread-out of feature descriptors [38]. Another work introduced second-order similarity [27] inspired by graph matching and clustering as a regularizer for learning local feature descriptors.

3. Methodology

Let \mathcal{Z} denote the set of data points, \mathcal{X} denote the set of embeddings, and $h(\cdot; \theta) : \mathcal{Z} \rightarrow \mathcal{X}$ denote the mapping learnt by the neural network h with parameters θ . \mathcal{C} denotes the set of classes to which the data points belong. Let \mathbf{Z} denote the array of data points sampled for training. For any index i , $\mathbf{Z}[i]$ denotes the data sample, $\mathbf{X}[i]$ denotes the corresponding embedding, and $\mathbf{c}[i]$ denotes its class.

3.1. Problem Formulation

Consider two pairs of points in the embedding space: $\mathbf{x}_1, \mathbf{x}_2$ and $\mathbf{y}_1, \mathbf{y}_2$, belonging to two different classes. They are l_2 -normalized, and lie on a hypersphere of unit radius. Most pair-wise metric learning-based losses try to consider the gap between $d(\mathbf{x}_1, \mathbf{x}_2)$ and $d(\mathbf{x}_1, \mathbf{y}_1)$ in some way, where d denotes the Euclidean distance. To consider informative samples which allow for larger values of the loss function, we can try to increase this gap. As $d(\mathbf{x}_1, \mathbf{x}_2)$ represents the distance between two positives, it utilizes both the samples from that class available in the batch and although we can try to generate \mathbf{x}_3 such that $d(\mathbf{x}_1, \mathbf{x}_3) > d(\mathbf{x}_1, \mathbf{x}_2)$, it would be mathematically difficult to ensure its class belongingness without more information (or assumptions). On the other hand, we can minimize the distance between the negatives represented by $d(\mathbf{x}_1, \mathbf{y}_1)$ by utilizing the remaining samples ($\mathbf{x}_2, \mathbf{y}_2$). In other words, we can find two points, \mathbf{p}_1 and \mathbf{p}_2 , which lie on the geodesic curves joining \mathbf{x}_1 and \mathbf{x}_2 , and \mathbf{y}_1 and \mathbf{y}_2 , respectively, such that

$$d(\mathbf{p}_1, \mathbf{p}_2) = \|\mathbf{p}_1 - \mathbf{p}_2\|_2 = \sqrt{2(1 - \mathbf{p}_1 \cdot \mathbf{p}_2)}, \quad (1)$$

is minimized. This is visualized in Fig. 1. A toy example demonstrating the behavior of triplet loss using different sets of $\mathbf{p}_1, \mathbf{p}_2$ and a gradient-based analysis is included in the supplementary.

Assumption on Data Distribution We assume that data distributions of different classes lying on the hypersphere are spherical-homoscedastic [10], defined as follows.

Definition 1. *Two distributions are said to be spherical-homoscedastic if their covariances have identical eigenvalues, i.e. the distributions are identically shaped.*

Proposition 1. *If the pairs of points $\mathbf{x}_1, \mathbf{x}_2$ and $\mathbf{y}_1, \mathbf{y}_2$ (from two different classes) belong to spherical-homoscedastic distributions, then the points on the curves $\widehat{\mathbf{x}_1\mathbf{x}_2}$ and $\widehat{\mathbf{y}_1\mathbf{y}_2}$ have a higher probability of belonging to the same classes as $\mathbf{x}_1, \mathbf{x}_2$ and $\mathbf{y}_1, \mathbf{y}_2$, respectively, than the other classes.*

Experimental validation for the assumption on data distribution (Definition 1) in our pipeline and the proof of Proposition 1 can be found in the supplementary material.

\mathbf{p}_1 is obtained by rotating \mathbf{x}_1 towards \mathbf{x}_2 along the curve $\widehat{\mathbf{x}_1\mathbf{x}_2}$ by angle α , which lies between 0 and $\alpha_0 = \cos^{-1}(\mathbf{x}_1 \cdot \mathbf{x}_2)$. Similarly, \mathbf{p}_2 is obtained by rotating \mathbf{y}_1 towards \mathbf{y}_2 along $\widehat{\mathbf{y}_1\mathbf{y}_2}$ by angle β , which lies between 0 and $\beta_0 = \cos^{-1}(\mathbf{y}_1 \cdot \mathbf{y}_2)$. This is done in three steps. First, the basis vectors for the space spanned by \mathbf{x}_1 and \mathbf{x}_2 are calculated using Gram-Schmidt orthogonalization, i.e.

$$\mathbf{n}_1 = \mathbf{x}_1; \mathbf{n}_2 = \frac{\mathbf{x}_2 - (\mathbf{x}_1 \cdot \mathbf{x}_2)\mathbf{x}_1}{\|\mathbf{x}_2 - (\mathbf{x}_1 \cdot \mathbf{x}_2)\mathbf{x}_1\|_2}. \quad (2)$$

Second, the rotation matrix \mathbf{R} is calculated using Ro-

drigues theorem, as follows:

$$\mathbf{R} = \mathbf{I} + \sin \alpha (\mathbf{n}_2 \mathbf{n}_1^T - \mathbf{n}_1 \mathbf{n}_2^T) - (1 - \cos \alpha) (\mathbf{n}_1 \mathbf{n}_1^T + \mathbf{n}_2 \mathbf{n}_2^T),$$

where \mathbf{I} is the identity matrix and $(\cdot)^T$ is the transpose operator. Finally, we find $\mathbf{p}_1 = \mathbf{R}\mathbf{x}_1$. Simplifying, we get:

$$\mathbf{p}_1 = \mathbf{n}_1 \cos \alpha + \mathbf{n}_2 \sin \alpha. \quad (3)$$

Similarly, \mathbf{p}_2 is obtained as:

$$\mathbf{p}_2 = \mathbf{n}_3 \cos \beta + \mathbf{n}_4 \sin \beta, \quad (4)$$

where \mathbf{n}_3 and \mathbf{n}_4 are the basis vectors obtained by Gram-Schmidt orthogonalization for \mathbf{y}_1 and \mathbf{y}_2 .

Incorporating these expressions for \mathbf{p}_1 and \mathbf{p}_2 in (1) and simplifying, the objective f to be minimized, is given by:

$$f(\alpha, \beta) = -\mathbf{p}_1 \cdot \mathbf{p}_2 = a \sin \alpha \sin \beta + b \cos \alpha \sin \beta + c \sin \alpha \cos \beta + d \cos \alpha \cos \beta, \quad (5)$$

where $a = -\mathbf{n}_2 \cdot \mathbf{n}_4$, $b = -\mathbf{n}_1 \cdot \mathbf{n}_4$, $c = -\mathbf{n}_2 \cdot \mathbf{n}_3$, and $d = -\mathbf{n}_1 \cdot \mathbf{n}_3$.

Further, we need to consider two constraints on α , so that \mathbf{p}_1 remains between \mathbf{x}_1 and \mathbf{x}_2 . They are given as:

$$g_1 = -\alpha \leq 0; g_2 = \alpha - \alpha_0 \leq 0. \quad (6)$$

Similarly, the constraints on β are given as:

$$g_3 = -\beta \leq 0; g_4 = \beta - \beta_0 \leq 0. \quad (7)$$

The Lagrangian function for the constrained optimization problem is given by:

$$L(\alpha, \beta, \lambda_1, \lambda_2, \lambda_3, \lambda_4) = f(\alpha, \beta) - \sum_{i=1}^4 \lambda_i g_i, \quad (8)$$

where $\lambda_i, i = 1, 2, 3, 4$ are Karush-Kuhn-Tucker (KKT) [14, 18] multipliers.

3.2. Finding Optimal Distance

As the constraints are differentiable without any critical points, we consider the KKT conditions [14, 18] to obtain the solution for this problem. They are listed as follows:

$$\begin{aligned} \frac{\partial L}{\partial \alpha} &= a \cos \alpha \sin \beta - b \sin \alpha \sin \beta + c \cos \alpha \cos \beta \\ &\quad - d \sin \alpha \cos \beta + \lambda_1 - \lambda_2 = 0, \end{aligned} \quad (9)$$

$$\begin{aligned} \frac{\partial L}{\partial \beta} &= a \sin \alpha \cos \beta + b \cos \alpha \cos \beta - c \sin \alpha \sin \beta \\ &\quad - d \cos \alpha \sin \beta + \lambda_3 - \lambda_4 = 0, \end{aligned} \quad (10)$$

$$\lambda_i g_i = 0; i = 1, 2, 3, 4, \quad (11)$$

$$\lambda_i \leq 0; i = 1, 2, 3, 4, \quad (12)$$

$$g_i \leq 0; i = 1, 2, 3, 4. \quad (13)$$

Case	α	β	λ_1	λ_2	λ_3	λ_4
1	0	$\tan^{-1}\left(\frac{b}{d}\right)$	$-\frac{\partial f}{\partial \alpha}\bigg _{\alpha=0, \beta=\hat{\beta}}$	0	0	0
2	α_0	$\tan^{-1}\left(\frac{a \sin \alpha_0 + b \cos \alpha_0}{c \sin \alpha_0 + d \cos \alpha_0}\right)$	0	$\frac{\partial f}{\partial \alpha}\bigg _{\alpha=\alpha_0, \beta=\hat{\beta}}$	0	0
3	$\tan^{-1}\left(\frac{c}{d}\right)$	0	0	0	$-\frac{\partial f}{\partial \beta}\bigg _{\alpha=\hat{\alpha}, \beta=0}$	0
4	$\tan^{-1}\left(\frac{a \sin \beta_0 + c \cos \beta_0}{b \sin \beta_0 + d \cos \beta_0}\right)$	β_0	0	0	0	$\frac{\partial f}{\partial \beta}\bigg _{\alpha=\hat{\alpha}, \beta=\beta_0}$
5	0	0	$-\frac{\partial f}{\partial \alpha}\bigg _{\alpha=0, \beta=0}$	0	$-\frac{\partial f}{\partial \beta}\bigg _{\alpha=0, \beta=0}$	0
6	0	β_0	$-\frac{\partial f}{\partial \alpha}\bigg _{\alpha=0, \beta=\beta_0}$	0	0	$\frac{\partial f}{\partial \beta}\bigg _{\alpha=0, \beta=\beta_0}$
7	α_0	0	0	$\frac{\partial f}{\partial \alpha}\bigg _{\alpha=\alpha_0, \beta=0}$	$-\frac{\partial f}{\partial \beta}\bigg _{\alpha=\alpha_0, \beta=0}$	0
8	α_0	β_0	0	$\frac{\partial f}{\partial \alpha}\bigg _{\alpha=\alpha_0, \beta=\beta_0}$	0	$\frac{\partial f}{\partial \beta}\bigg _{\alpha=\alpha_0, \beta=\beta_0}$

Table 1: Expressions for α , β , and the KKT multipliers for 8 cases of complementary slackness.

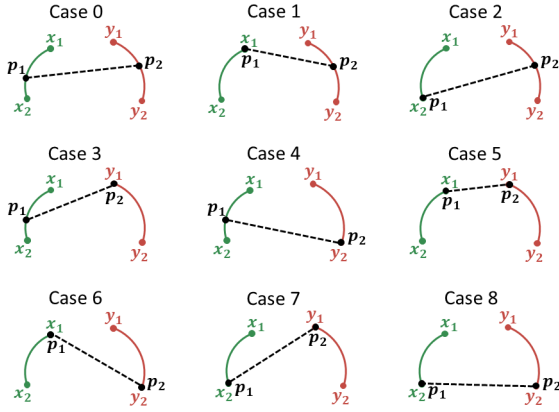


Figure 2: Illustrations of the 9 cases in which the KKT conditions for minimizing the distance between $\widehat{\mathbf{x}}_1 \widehat{\mathbf{x}}_2$ and $\widehat{\mathbf{y}}_1 \widehat{\mathbf{y}}_2$ can be satisfied. In case 0, \mathbf{p}_1 and \mathbf{p}_2 lie between the end points of the respective curves. In cases 1-4, only one of them is an end point of the respective curve. In cases 5-8, both of them are end points of the respective curves.

There are 16 possible ways to satisfy (11). However, some constraints cannot be binding simultaneously (e.g. if $g_1 = 0, g_2 \neq 0$), which leaves 9 cases to consider, which are shown in Fig. 2. The solutions are discussed below in brief. The proofs can be found in the supplementary material.

We first consider the case where none of the constraints is binding (Case 0). In order to satisfy (11), $\lambda_i = 0$ for $i = 1, 2, 3, 4$. In this case, (12) is also satisfied. Leveraging (9) and (10), we get a quadratic equation in $\tan \alpha$. On solving it and finding the corresponding value of β using (9) or (10), we get two sets of solutions:

$$\hat{\alpha} = \tan^{-1} \left(\frac{A \pm \sqrt{A^2 + 4(ab + cd)^2}}{2(ab + cd)} \right),$$

$$\hat{\beta} = \tan^{-1} \left(\frac{B \mp \sqrt{B^2 + 4(ac + bd)^2}}{2(ac + bd)} \right),$$

where $A = a^2 - b^2 + c^2 - d^2$ and $B = a^2 + b^2 - c^2 - d^2$. The obtained $\hat{\alpha}$ and $\hat{\beta}$ are used to verify the validity of (13).

Similarly, the KKT conditions are checked for the remaining 8 cases as well. The expressions for α , β , and the KKT multipliers for various cases are listed in Table 1. The first 4 rows of the table cover the cases where one of the optimal points is an end point of a curve, while the other can be any point on the other curve. The next 4 rows of the table cover the cases where both the optimal points are one of the two end points. The values obtained using these expressions are used to verify the conditions in (12) and (13). In case any condition is violated, the solution is discarded.

Once $\hat{\alpha}$ and $\hat{\beta}$ are obtained, they are substituted in (3) and (4) to obtain \mathbf{p}_1 and \mathbf{p}_2 . Finally, d is calculated using (1), which is used in different metric learning losses, as discussed below.

3.3. Optimal Hard Negative Embeddings for Deep Metric Learning

Let $d_{i,j}$ denote the distance between $\mathbf{X}[i]$ and $\mathbf{X}[j]$ and $s_{i,j}$ denote the similarity, given by $s_{i,j} = 1 - \frac{d_{i,j}^2}{2}$. Let $d_{i,j,k,l}$ denote the optimal distance between the curves joining $\mathbf{X}[i]$ and $\mathbf{X}[j]$, and $\mathbf{X}[k]$ and $\mathbf{X}[l]$, such that $\mathbf{c}[i] = \mathbf{c}[j] \neq \mathbf{c}[k] = \mathbf{c}[l]$. We evaluate our method for the following metric learning losses:

3.3.1 Triplet Loss

One of the earliest metric learning losses, triplet loss [24, 33] tries to pull a positive sample belonging to the same class as the anchor, close to the anchor, and push a negative sample belonging to a different class away from the anchor, by a margin m . It is formulated as follows:

$$\mathcal{L}_{Tri} = \frac{1}{|\mathcal{P}|} \sum_{\substack{(i,j) \in \mathcal{P} \\ k: \mathbf{c}[i] \neq \mathbf{c}[k]}} [d_{i,j} - d_{i,k} + m]_+, \quad (14)$$

where $[\cdot]_+$ denotes the hinge function, and \mathcal{P} denotes the set of classes in the batch. The modified triplet loss, obtained

by incorporating LoOp, is given by:

$$\mathcal{L}'_{Tri} = \frac{1}{|\mathcal{P}|} \sum_{\substack{(i,j) \in \mathcal{P} \\ k,l: \mathbf{c}[i] \neq \mathbf{c}[k] = \mathbf{c}[l]}} [d_{i,j} - d_{i,j,k,l} + m]_+. \quad (15)$$

3.3.2 HPHN Triplet Loss

In order to better utilize the samples present in a batch, hard positive and hard negative (HPHN) mining [34, 12] has been proposed. The expression for HPHN-triplet loss is given by:

$$\mathcal{L}_{HPHNtri} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left[\max \left(\max_{\mathbf{c}[i]=\mathbf{c}[k]} d_{i,k}, \max_{\mathbf{c}[j]=\mathbf{c}[l]} d_{j,l} \right) + m - \min \left(\min_{\mathbf{c}[i] \neq \mathbf{c}[k]} d_{i,k}, \min_{\mathbf{c}[j] \neq \mathbf{c}[l]} d_{j,l} \right) \right]_+.$$

The modified HPHN-triplet loss, obtained by incorporating LoOp, is given by:

$$\mathcal{L}'_{HPHNtri} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left[\max \left(\max_{\mathbf{c}[i]=\mathbf{c}[k]} d_{i,k}, \max_{\mathbf{c}[j]=\mathbf{c}[l]} d_{j,l} \right) + m - \min_{\mathbf{c}[i] \neq \mathbf{c}[k] = \mathbf{c}[l]} d_{i,j,k,l} \right]_+.$$

3.3.3 Lifted Structure Loss

The lifted structure (LS) loss [25] tries to push a pair of samples belonging to one class away from all the samples belonging to some different class. It is formulated as:

$$\mathcal{L}_{LS} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left[d_{i,j} + m - \min \left(\min_{\mathbf{c}[i] \neq \mathbf{c}[k]} d_{i,k}, \min_{\mathbf{c}[j] \neq \mathbf{c}[l]} d_{j,l} \right) \right]_+. \quad (16)$$

When only 2 samples of each class are present in the batch, the HPHN-triplet loss becomes the same as LS loss. The modified LS loss, obtained by incorporating LoOp, is given by:

$$\mathcal{L}'_{LS} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left[d_{i,j} + m - \min_{\mathbf{c}[i] \neq \mathbf{c}[k] = \mathbf{c}[l]} d_{i,j,k,l} \right]_+.$$

3.3.4 Multi-Similarity Loss

The multi-similarity (MS) loss [32] considers self-similarity and relative similarities between positive as well as negative pairs to learn the embedding space. It involves mining for

hard negative and positive pairs, where a pair of negatives ($\mathbf{X}[i], \mathbf{X}[j]$) (for $\mathbf{c}[i] \neq \mathbf{c}[j]$) is selected when:

$$s_{i,j}^- > \min_{\mathbf{c}[i]=\mathbf{c}[k]} s_{i,k} - \epsilon.$$

Similarly, a pair of positives is selected, when:

$$s_{i,j}^+ < \max_{\mathbf{c}[i] \neq \mathbf{c}[k]} s_{i,k} + \epsilon.$$

In both the conditions, ϵ is a hyperparameter. The second step is general pair weighting, and the final expression for MS loss is given by:

$$\mathcal{L}_{MS} = \frac{1}{|\mathcal{P}|} \sum_{i \in \mathcal{P}} \left[\frac{1}{\alpha_m} \log \left(1 + \sum_{\mathbf{c}[i]=\mathbf{c}[j]} e^{-\alpha_m (s_{i,j} - \lambda)} \right) + \frac{1}{\beta_m} \log \left(1 + \sum_{\mathbf{c}[i] \neq \mathbf{c}[j]} e^{\beta_m (s_{i,j} - \lambda)} \right) \right],$$

where α_m and β_m are hyperparameters, and λ denotes the margin. For the modified version of MS loss, obtained by incorporating LoOp, the criteria for selecting hard negatives is given as follows:

$$s_{i,k,j,l}^- > \min_{\mathbf{c}[i]=\mathbf{c}[k]} s_{i,k} - \epsilon.$$

The expressions for selecting hard positives as well as for general pair weighting remain unchanged.

3.4. Implementation

Algorithm 1 presents the steps involved in the proposed approach. In order to make the implementation efficient, the calculations for all possible combinations are carried out in parallel. Given a batch size B_S and number of samples per class N , the number of such combinations C is given as:

$$C = \frac{\frac{B_S}{2} \left(\frac{B_S}{2} - 1 \right)}{2} - \frac{B_S}{N} \times \frac{\frac{N}{2} \left(\frac{N}{2} - 1 \right)}{2} = \frac{B_S(B_S - N)}{8}.$$

The runtime of our method is documented in Section 4.3.

4. Experiments

We conduct numerous experiments to evaluate the effectiveness of LoOp for both image clustering and retrieval tasks [25]. We report the standardized metrics F_1 score and normalized mutual information (NMI) for image clustering [22], and Recall@K values for image retrieval.

Datasets We evaluate our method on three widely used benchmark datasets described below. We use the conventional zero-shot setting [25, 5], *i.e.* no intersection in train and test classes. Results for train-validate-test split are included in the supplementary.

- The CUB-200-2011 dataset [29] comprises 11,788 images of 200 bird species. We use the first 100 species (5864 images) for training and the remaining 100 (5924 images) for testing.

Algorithm 1: Training with optimal hard negatives

input : Set of training images and labels, ensuring number of samples of each class is even; hyperparameters depending on the loss function; number of iterations T

output: Trained network parameters θ

```
1 Initialize  $\theta$ 
2 for iteration= $1, 2, \dots, T$  do
3   Extract embeddings of the images in the training
   batch and perform  $l_2$ -normalization
4   Create pairs of embeddings (by choosing
   alternate elements from  $\mathbf{X}$ ) to get the sets  $\mathbf{X}_1$ 
   and  $\mathbf{X}_2$ , with the corresponding classes
    $c_{\mathbf{X}_1} = c_{\mathbf{X}_2}$ 
5   Create 4 arrays:  $\mathbf{X}_1, \mathbf{X}_2, \mathbf{Y}_1, \mathbf{Y}_2$ , containing
   the possible combinations of the
   aforementioned pairs, ensuring that
    $c_{\mathbf{X}_1} \neq c_{\mathbf{Y}_1}$ 
6   Calculate the optimal distance between each
   combination of pairs, i.e. between the
   corresponding elements of the arrays  $\mathbf{X}_1, \mathbf{X}_2,$ 
    $\mathbf{Y}_1, \mathbf{Y}_2$  in parallel
7   Find the minimum distance for each unique pair
   in  $\mathbf{X}_1, \mathbf{X}_2$ 
8   Calculate the metric learning-based loss using
   the set of optimal distances
9   Update network parameters
10 return  $\theta$ 
```

- The Cars196 dataset [17] consists of 16,185 images of 196 classes of cars. We use the first 98 car classes (8054 images) as the training data and the remaining 98 classes (8131 images) as testing data.
- The Stanford Online Products (SOP) dataset [25] contains 120,053 images of 22,634 online products listed on eBay.com. The train set consists of 59,551 images of first 11,318 product classes, while the test set uses the remaining 60,502 images of 11,316 classes.

Settings We use the MXNet package for implementing our proposed approach. All the experiments were conducted using one NVIDIA GeForce RTX 2080 Ti, 11 GB memory. For pre-processing input images, we resize them to 256×256 , randomly crop them to size 227×227 , horizontally flip them left or right. We use an ImageNet ILSVRC [3] pre-trained GoogLeNet [26] with a fully connected randomly initialized layer as our feature extracting network. For a fair comparison, we compare with methods that use the same network. Results and comparisons on other network architectures have been included in the supplementary. Throughout the experiments, we use 512-dimensional

feature embedding vectors. We use the Adam optimizer [15] for training our models, with $\beta_1 = 0.9, \beta_2 = 0.999$, learning rate of 10^{-4} . The hyperparameters α_m and β_m are set as 2 and 50, respectively. For Cars196, CUB-200-2011 datasets, we use a batch size of 32, and for the SOP dataset, we set the batch size to 128. The SOTA methods use a larger batch size of 120 or 128 for the Cars196, CUB-200-2011 dataset as well. As shown in Section 4.1, larger batch size provides a boost in performance metrics. t-SNE visualizations at different epochs demonstrating the training process can be found in the supplementary.

4.1. Effect of Hyperparameters

In this section, we study the effects of changing the values of various hyperparameters on the Recall@1 (%) performance.

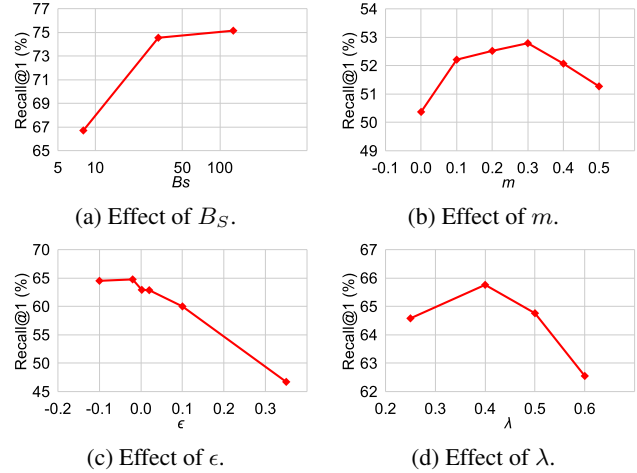


Figure 3: Variation in Recall@1 (%) values with various hyperparameters.

Effect of Batch Size Fig. 3a shows the effect of varying B_S (log scale is used for x -axis) for Cars196 dataset using HPHN-triplet loss. The number of unique classes in the batch is 4. It is observed that as B_S increases, there is an increase in the Recall@1 value. However, the rate of increase is lower for higher B_S .

Effect of Margin Fig. 3b shows the effect of varying m , which is used in triplet, HPHN-triplet, and LS losses. These values are for CUB-200-2011 dataset for LS loss, with $B_S = 32$. It is observed that as m increases, the performance initially increases, and then drops. This shows that for smaller values of m , the embeddings of dissimilar classes are not sufficiently separated, whereas for higher values of m , they are forced to form more compact and distant clusters. This might lead to overfitting resulting in subsequent drop in performance.

Effect of ϵ and λ Next, we consider the effect of varying two hyperparameters of the MS loss. In both cases, the

results are reported for the Cars196 dataset, with $B_S = 8$ and $N = 4$. Fig. 3c presents the results of varying ϵ , for $\lambda = 0.5$. From the range of the y -axis, it can be observed that the performance of MS loss is quite sensitive to this parameter. It is noteworthy that since our approach finds the hardest negatives for each pair in a batch, the role of ϵ (to make the criteria for selection of hard positives more strict) is diminished, and we get performance gains for lower values of ϵ . Fig. 3d presents the results of varying λ , for $\epsilon = -0.02$. As the value of similarity lies between -1 and 1 , embeddings belonging to the same class are constrained to have similarity more than λ , while those belonging to dissimilar classes should have a smaller similarity than λ . Varying λ for MS loss has a similar effect as varying m for LS loss.

4.2. Effect of l_2 -normalization

Most of the metric learning losses considered in this work perform l_2 -normalization on the embeddings before computing the loss. In order to observe the effect of l_2 -normalization, we consider the case where the optimal points lie on the line segments $\overline{\mathbf{x}_1\mathbf{x}_2}$ and $\overline{\mathbf{y}_1\mathbf{y}_2}$ rather than the curves joining these points. In other words, the optimal points are not constrained to lie on the hypersphere. The optimization problem is then formulated differently. A brief discussion is presented here, and the details can be found in the supplementary material.

The optimal points \mathbf{p}_1 and \mathbf{p}_2 can be obtained as:

$$\begin{aligned}\mathbf{p}_1 &= (1 - k_1)\mathbf{x}_1 + k_1\mathbf{x}_2, \\ \mathbf{p}_2 &= (1 - k_2)\mathbf{y}_1 + k_2\mathbf{y}_2,\end{aligned}\quad (17)$$

where k_1 and k_2 are the parameters to be optimized. f can be obtained by incorporating these expressions in (1). The constraints in this case are given as:

$$\begin{aligned}g_1 &= -k_1 \leq 0; g_2 = k_1 - 1 \leq 0, \\ g_3 &= -k_2 \leq 0; g_4 = k_2 - 1 \leq 0.\end{aligned}\quad (18)$$

The Lagrangian function can be obtained using (8). The partial derivatives are given as:

$$\begin{aligned}\frac{\partial L}{\partial k_1} &= (\mathbf{x}_2 - \mathbf{x}_1) \cdot (k_1(\mathbf{x}_2 - \mathbf{x}_1) - k_2(\mathbf{y}_2 - \mathbf{y}_1) \\ &\quad + \mathbf{x}_1 - \mathbf{y}_1) + \lambda_1 - \lambda_2 = 0,\end{aligned}\quad (19)$$

$$\begin{aligned}\frac{\partial L}{\partial k_2} &= (\mathbf{y}_2 - \mathbf{y}_1) \cdot (-k_1(\mathbf{x}_2 - \mathbf{x}_1) + k_2(\mathbf{y}_2 - \mathbf{y}_1) \\ &\quad - \mathbf{x}_1 + \mathbf{y}_1) + \lambda_3 - \lambda_4 = 0.\end{aligned}\quad (20)$$

Considering KKT conditions, the solution when all the multipliers are 0 is obtained as follows:

$$\begin{aligned}\hat{k}_1 &= \frac{\mathbf{u} \cdot \mathbf{v} \mathbf{v} \cdot \mathbf{w} - \|\mathbf{v}\|_2^2 \mathbf{u} \cdot \mathbf{w}}{(\mathbf{u} \cdot \mathbf{v})^2 - \|\mathbf{u}\|_2^2 \|\mathbf{v}\|_2^2}, \\ \hat{k}_2 &= \frac{-\mathbf{u} \cdot \mathbf{v} \mathbf{u} \cdot \mathbf{w} + \|\mathbf{u}\|_2^2 \mathbf{v} \cdot \mathbf{w}}{(\mathbf{u} \cdot \mathbf{v})^2 - \|\mathbf{u}\|_2^2 \|\mathbf{v}\|_2^2},\end{aligned}\quad (21)$$

where $\mathbf{u} = \mathbf{x}_1 - \mathbf{x}_2$, $\mathbf{v} = \mathbf{y}_1 - \mathbf{y}_2$, and $\mathbf{w} = \mathbf{x}_1 - \mathbf{y}_1$. The solutions to the remaining cases can be obtained by following a similar approach as presented in Section 3.2.




Figure 4: Performance (%) comparison of triplet, Ours (Triplet) without l_2 -normalization and Ours (Triplet) with l_2 -normalization for clustering and retrieval tasks.

The results of our approach for triplet loss with and without l_2 -normalization as well as comparison with triplet loss (baseline) for Cars196 dataset are shown in Fig. 4. It is observed that our approach performs better when the optimal distance is calculated between the curves rather than the line segments. This is because the points are l_2 -normalized in the former case. Nevertheless, the latter case without l_2 -normalization performs better than the baseline. This shows that the proposed approach to consider the minimum distance between two pairs of samples from different classes, greatly benefits the deep metric learning framework.

4.3. Time and Memory Requirement

Incorporating LoOp into metric learning losses leads to a minimal overhead in computation and time. For MS loss, the time per iteration changes from 0.071s to 0.121s. For HPHN-triplet loss, the time per iteration for LoOp is 0.058s. Thus, the time taken by LoOp remains in milliseconds for various metric learning losses.

In terms of computation and memory requirement, HPHN-triplet loss requires $O(B_S^2)$ combinations to compute the loss, which is same as the requirement of our approach. In contrast, the memory requirement of EE [16] depends on the number of interpolated points n . It needs to consider $O(n^2 B_S^2)$ combinations to compute the loss. Thus, our method is both time and memory efficient.

4.4. Comparison with State-of-the-art

The results of our proposed approach and comparison with the baseline metric learning losses as well as SOTA methods for hard negative generation and sampling are shown in Tables 2 and 3, respectively. In Table 2, it can be observed that for CUB-200-2011, LoOp outperforms the baseline metric learning losses in all cases. With the exception of EE for LS loss, our approach also shows better performance than the SOTA hard negative generation methods. The improvements with respect to previous best values range from 0.1% to 6% in Recall@1 and 0.7% to 2.2% in

Method	CUB-200-2011						Cars196						SOP				
	NMI	F ₁	R@1	R@2	R@4	R@8	NMI	F ₁	R@1	R@2	R@4	R@8	NMI	F ₁	R@1	R@10	R@100
Triplet	49.8	15.0	35.9	47.7	59.1	70.0	52.9	17.9	45.1	57.4	69.7	79.2	86.3	20.2	53.9	72.1	85.7
DAML (Triplet)	51.3	17.6	37.6	49.3	61.3	74.4	56.5	22.9	60.6	72.5	82.5	89.9	87.1	22.3	58.1	75.0	88.0
HDML (Triplet)	55.1	21.9	43.6	55.8	67.7	78.3	59.4	27.2	61.0	72.6	80.7	88.5	87.2	22.5	58.5	75.5	88.3
EE (Triplet)	55.7	22.4	44.3	57.0	68.1	78.9	60.3	25.1	57.2	70.5	81.3	88.2	87.4	24.8	62.4	79.0	91.0
LoOp (Triplet)	59.9	26.5	50.3	63.1	74.4	82.7	60.2	28.0	61.7	72.9	82.1	88.7	88.7	30.6	69.6	84.9	93.7
HPHNtri	58.1	24.2	48.3	61.9	73.0	82.3	57.4	22.6	60.3	73.4	83.5	90.5	91.4	43.3	75.5	88.8	95.4
EE (HPHNtri)	60.5	27.0	51.7	63.5	74.5	82.5	63.1	32.0	71.6	80.7	87.5	92.2	91.5	43.6	77.2	89.6	95.5
LoOp (HPHNtri)	61.2	27.5	51.8	63.6	74.9	83.9	64.8	34.1	74.5	83.3	89.4	93.6	90.8	36.9	77.6	90.2	96.1
LS	56.4	22.6	46.9	59.8	71.2	81.5	57.8	25.1	59.9	70.4	79.6	87.0	87.2	25.3	62.6	80.9	91.2
DAML (LS)	59.5	26.6	49.0	62.2	73.7	83.3	63.1	31.9	72.5	82.1	88.5	92.9	89.1	31.7	66.3	82.8	92.5
EE (LS)	61.2	28.2	54.2	66.6	76.7	85.2	59.1	27.2	65.2	76.4	85.6	89.5	89.6	35.3	70.6	85.5	93.6
LoOp (LS)	60.3	27.4	52.8	65.0	76.2	84.5	57.4	25.9	66.2	76.8	84.7	89.8	89.8	35.9	77.1	89.7	95.6
MS	59.3	26.0	50.9	63.0	74.1	83.3	63.3	31.7	71.0	80.8	87.5	92.6	89.3	33.7	75.0	88.7	95.7
LoOp (MS)	61.1	28.4	52.0	64.3	75.0	84.1	63.0	30.6	72.6	81.5	88.4	92.8	89.4	34.2	76.6	89.8	95.8

Table 2: Clustering and retrieval performance of generation-based methods for CUB-200-2011, Cars196 and SOP datasets. **Bold** numbers indicate the best values within each metric learning loss. **Blue** numbers indicate cases where the combination of the proposed approach with a metric learning loss outperforms the baseline metric learning loss.

Method	NMI	F ₁	R@1	R@2	R@4	R@8
Rand-disjoint	49.8	15.0	35.9	47.7	59.1	70.0
Semi-hard	53.4	17.9	40.6	52.3	64.2	75.0
DE-DSP	53.7	19.8	41.0	53.2	64.8	-
Dis-weighted	56.3	25.4	44.1	57.5	70.1	80.5
Smart mining	58.1	-	45.9	57.7	69.6	79.8
LoOp	59.9	26.5	50.3	63.1	74.4	82.7

(a) CUB-200-2011 dataset

Method	NMI	F ₁	R@1	R@2	R@4	R@8
Rand-disjoint	52.9	17.9	45.1	57.4	69.7	79.2
Semi-hard	55.7	22.4	53.2	65.4	74.3	83.6
DE-DSP	55.0	22.3	59.3	71.3	81.3	-
Dis-weighted	58.3	25.4	59.4	72.3	81.6	87.2
Smart mining	58.2	-	56.1	68.3	78.0	85.9
LoOp	60.2	28.0	61.7	72.9	82.1	88.7

(b) Cars196 dataset

Method	NMI	F ₁	R@1	R@10	R@100
Rand-disjoint	86.3	20.2	53.9	72.1	85.7
Semi-hard	86.7	22.1	57.8	75.3	88.1
DE-DSP	87.4	22.7	58.2	75.8	88.4
Dis-weighted	87.9	23.4	58.9	77.2	89.6
LoOp	88.7	30.6	69.6	84.9	93.7

(c) SOP dataset

Table 3: Comparison with clustering and retrieval performance of sampling-based methods for (a) CUB-200-2011, (b) Cars196 and (c) SOP datasets using triplet loss. **Bold** numbers indicate the best values. **Blue** numbers indicate cases where the combination of LoOp and triplet loss outperforms the baseline. - indicates not reported.

NMI. For Cars196 dataset, our approach provides a boost in performance for most of the baseline losses. The improvements with respect to previous best values range from 0.7% to 2.9% in Recall@1. For SOP dataset, it can be seen that

for all cases, LoOp improves over the baseline losses. In terms of comparison with SOTA methods, there is an increase of upto 7.2% in Recall@1, while NMI shows an increase of 0.1% to 1.3%.

It is noteworthy that although HPHN-triplet and MS loss look for hard positive and negative samples as part of their formulation, LoOp further improves the results obtained by these methods. The maximum increase in Recall@1 value is 14.2% for HPHN-triplet and 1.6% for MS loss for the Cars196 dataset, as compared to the respective baselines. MS loss has a strict criterion that selects only the most informative samples, and hence the increase in performance is not very significant.

Table 3 presents the comparison with sampling methods, like random sampling of disjoint tuples (rand-disjoint), semi-hard negative mining [24], DE-DSP [4], distance-weighted sampling [21], and smart mining [11]. It can be seen that LoOp outperforms all methods for all three datasets for both image clustering and retrieval tasks.

5. Conclusion

We propose an approach to find optimal hard negative embeddings for deep metric learning. This is done by calculating the points which minimize the distance between two curves, each representative of a different class. Our approach can be easily integrated with representative metric learning losses, avoids the computational cost of generating hard data samples, and utilizes the training set effectively unlike hard negative mining-based methods. We demonstrated that our approach boosts the performance of existing metric learning losses and outperforms state-of-the-art hard negative sampling and generation-based methods.

References

- [1] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning (ICML)*, volume 70, pages 214–223, 2017. 2
- [2] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 539–546, 2005. 1
- [3] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009. 6
- [4] Y. Duan, L. Chen, J. Lu, and J. Zhou. Deep embedding learning with discriminative sampling policy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 8
- [5] Y. Duan, W. Zheng, X. Lin, J. Lu, and J. Zhou. Deep adversarial metric learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2780–2789, 2018. 2, 5
- [6] I. Goodfellow et al. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2672–2680. 2014. 2
- [7] A. Gordo et al. Deep image retrieval: Learning global representations for image search. In *European Conference on Computer Vision (ECCV)*, pages 241–257, Cham, 2016. Springer International Publishing. 1
- [8] M. Guillaumin, J. Verbeek, and C. Schmid. Is that you? metric learning approaches for face identification. In *IEEE International Conference on Computer Vision (ICCV)*, pages 498–505, 2009. 1
- [9] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 1735–1742, 2006. 1
- [10] O. Hamsici and A. Martínez. Spherical-homoscedastic distributions: The equivalency of spherical and normal distributions in classification. *Journal of Machine Learning Research*, 8:1583–1623, 07 2007. 3
- [11] B. Harwood, V. Kumar B.G., G. Carneiro, I. Reid, and T. Drummond. Smart mining for deep metric learning. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2840–2848, 2017. 2, 8
- [12] A. Hermans*, L. Beyer*, and B. Leibe. In defense of the triplet loss for person re-identification. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2017. 5
- [13] C. Huang, C. C. Loy, and X. Tang. Local similarity-aware deep feature embedding. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1262–1270, 10 2016. 2
- [14] W. Karush. Minima of functions of several variables with inequalities as side conditions. Master’s thesis, Department of Mathematics, University of Chicago, USA, 1939. 3
- [15] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2017. 6
- [16] B. Ko and G. Gu. Embedding expansion: Augmentation in embedding space for deep metric learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7255–7264, 2020. 2, 7
- [17] J. Krause, M. Stark, Jia Deng, and L. Fei-Fei. 3d object representations for fine-grained categorization. In *IEEE International Conference on Computer Vision Workshops (ICCVW)*, page 554–561, 2013. 2, 6
- [18] H. W. Kuhn and A. W. Tucker. Nonlinear programming. In *Second Berkeley Symposium on Mathematical Statistics and Probability*, 1951. 3
- [19] W. Li, R. Zhao, T. Xiao, and X. Wang. DeepReID: Deep filter pairing neural network for person re-identification. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 152–159, 2014. 1
- [20] J. Lu, J. Hu, and Y. Tan. Discriminative deep metric learning for face and kinship verification. *IEEE Transactions on Image Processing*, 26(9):4269–4282, 2017. 1
- [21] R. Manmatha, C. Wu, A. Smola, and P. Krahenbuhl. Sampling matters in deep embedding learning. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2859–2867, 10 2017. 1, 2, 8
- [22] C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, USA, 2008. 5
- [23] D. D. Mohan, N. Sankaran, D. Fedorishin, S. Setlur, and V. Govindaraju. Moving in the right direction: A regularization for deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [24] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823, 2015. 1, 2, 4, 8
- [25] H. Song, Y. Xiang, S. Jegelka, and S. Savarese. Deep metric learning via lifted structured feature embedding. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4004–4012, 2016. 1, 2, 5, 6
- [26] C. Szegedy et al. Going deeper with convolutions. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, 2015. 6
- [27] Y. Tian, X. Yu, B. Fan, F. Wu, H. Heijnen, and V. Balntas. SOSNet: Second order similarity regularization for local descriptor learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [28] E. Ustinova and V. Lempitsky. Learning deep embeddings with histogram loss. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 29, pages 4170–4178. Curran Associates, Inc., 2016. 1
- [29] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical report, 2011. 2, 5
- [30] F. Wang, W. Zuo, L. Lin, D. Zhang, and L. Zhang. Joint learning of single-image and cross-image representations for

- person re-identification. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1288–1296, 2016. [1](#)
- [31] J. Wang et al. Learning fine-grained image similarity with deep ranking. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1386–1393, 2014. [1](#)
- [32] X. Wang, X. Han, W. Huang, D. Dong, and M. R. Scott. Multi-similarity loss with general pair weighting for deep metric learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5022–5030, 2019. [5](#)
- [33] K. Q. Weinberger and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research*, 10:207–244, 2009. [1](#), [4](#)
- [34] H. Xuan, A. Stylianou, and R. Pless. Improved embeddings with easy positive triplet mining. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2463–2471, 2020. [5](#)
- [35] D. Yi, Z. Lei, S. Liao, and S. Z. Li. Deep metric learning for person re-identification. In *International Conference on Pattern Recognition*, pages 34–39, 2014. [1](#)
- [36] X. Yin, X. Yu, K. Sohn, X. Liu, and M. Chandraker. Feature transfer learning for face recognition with under-represented data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [2](#)
- [37] Y. Yuan, K. Yang, and C. Zhang. Hard-aware deeply cascaded embedding. In *IEEE International Conference on Computer Vision (ICCV)*, pages 814–823, 11 2017. [2](#)
- [38] X. Zhang, F. X. Yu, S. Kumar, and S. Chang. Learning spread-out local feature descriptors. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. [2](#)
- [39] Y. Zhao, Z. Jin, G. Qi, H. Lu, and X. Hua. An adversarial approach to hard triplet generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. [2](#)
- [40] W. Zheng, Z. Chen, J. Lu, and J. Zhou. Hardness-aware deep metric learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 72–81, 2019. [2](#)
- [41] J. Zhou, P. Yu, W. Tang, and Y. Wu. Efficient online local metric adaptation via negative samples for person re-identification. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2439–2447, 2017. [1](#)