

Embedding Expansion: Augmentation in Embedding Space for Deep Metric Learning

Byungsoo Ko*, Geonmo Gu*
Clova Vision, NAVER Corp.

{kobiso62, korgm403}@gmail.com

Abstract

Learning the distance metric between pairs of samples has been studied for image retrieval and clustering. With the remarkable success of pair-based metric learning losses [5, 36, 25, 21], recent works [44, 9, 42] have proposed the use of generated synthetic points on metric learning losses for augmentation and generalization. However, these methods require additional generative networks along with the main network, which can lead to a larger model size, slower training speed, and harder optimization. Meanwhile, post-processing techniques, such as query expansion [7, 6] and database augmentation [29, 3], have proposed a combination of feature points to obtain additional semantic information. In this paper, inspired by query expansion and database augmentation, we propose an augmentation method in an embedding space for metric learning losses, called **embedding expansion**. The proposed method generates synthetic points containing augmented information by a combination of feature points and performs hard negative pair mining to learn with the most informative feature representations. Because of its simplicity and flexibility, it can be used for existing metric learning losses without affecting model size, training speed, or optimization difficulty. Finally, the combination of embedding expansion and representative metric learning losses outperforms the state-of-the-art losses and previous sample generation methods in both image retrieval and clustering tasks. The implementation is publicly available¹.

1. Introduction

Deep metric learning aims to learn a distance metric for measuring similarities between given data points. It has played an important role in a variety of applications in computer vision, such as image retrieval [33, 11], re-identification [43, 17, 39], clustering [14], and face recog-

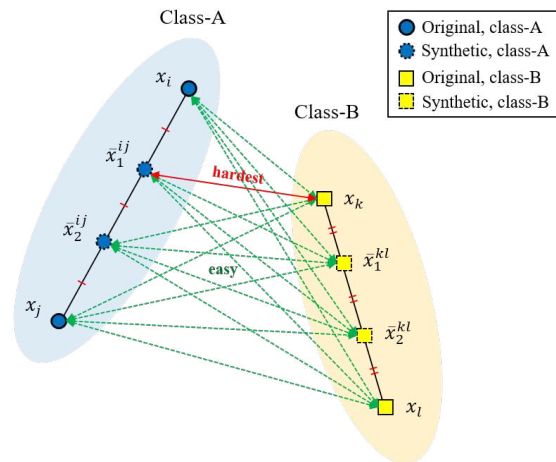


Figure 1. Illustration of our proposed embedding expansion consisting of two steps. In the first step, given a pair of embedding points from the same class, we perform linear interpolation on the line of the embedding points to generate internally dividing synthetic points into $n + 1$ equal parts, where n is the number of synthetic points ($n = 2$ in the Figure). Secondly, we select the hardest negative pair within the possible negative pairs of original and synthetic points. Rectangles and circles represent the two different classes, where the plain boundary indicates original points (x_i, x_j, x_k, x_l), and the dotted boundary indicates synthetic points ($\bar{x}_1^{ij}, \bar{x}_2^{ij}, \bar{x}_1^{kl}, \bar{x}_2^{kl}$).

niton [5, 22, 27]. The core idea of deep metric learning is to learn an embedding space by pulling the same class samples together and by pushing different class samples apart. To learn an embedding space, many of the metric learning losses take pairs of samples to optimize the loss with the desired properties. Conventional pair-based metric learning losses are contrastive loss [5, 12] and triplet loss [36, 22], which take 2-tuple and 3-tuple samples, respectively. N-pair loss [25] and lifted structured loss [21] aim to exploit a greater number of negative samples to improve the conventional metric learning losses. Recent works [34, 30, 37, 20]

* Authors contributed equally.

¹<https://github.com/clovaai/embedding-expansion>

have been proposed to consider the richer structured information among multiple samples.

Along with the importance of the loss function, the sampling strategy also plays an important role in performance. Different strategies for the same loss function can lead to extremely different results [37, 40]. Thus, there has been active research on sampling strategy and hard sample mining methods [37, 2, 13, 22]. One drawback of sampling and mining strategies is that it can lead to a biased model due to training with a minority of selected hard samples and ignoring a majority of easy samples [37, 22, 44]. To address this problem, hard sample generation methods [44, 9, 42] have been proposed to generate hard synthesis with easy samples. However, those methods require an additional sub-network as a generator, such as a generative adversarial network and an auto-encoder, which can cause a larger model size, slower training speed, and more training difficulty [4].

In this paper, we propose a novel augmentation method in the embedding space for deep metric learning, called embedding expansion (EE). Inspired by query expansion [7, 6] and database augmentation techniques [29, 3], the proposed method combines feature points to generate synthetic points with augmented image representations. As illustrated in Figure 1, it generates internally dividing points into $n + 1$ equal parts within pairs of the same classes and performs hard negative pair mining among original and synthetic points. By exploiting synthetic points with augmented information, it attains a performance boost through a more generalized model. The proposed method is simple and flexible enough that it can be combined with existing pair-based metric learning losses. Unlike the previous sample generation method, the proposed method does not suffer from the problems caused by using an additional generative network, because it performs simple linear interpolation for sample generation. We demonstrate that combining the proposed method with existing metric learning losses achieves a significant performance boost, while it also outperforms the previous sample generation methods on three famous benchmarks (CUB200-2011 [32], CARS196 [16], and Stanford Online Products [21]) in both image retrieval and clustering tasks.

2. Related Work

Sample Generation Recently, there have been attempts to generate potential hard samples for pair-based metric learning losses [44, 9, 42]. The main purpose of generating samples is to exploit a large number of easy negatives and train the network with this extra semantic information. The deep adversarial metric learning (DAML) framework [9] and the hard triplet generation (HTG) [42] use generative adversarial networks to generate synthetic samples. The hardness-aware deep metric learning (HDML) framework [44] exploits an auto-encoder to generate label-

preserving synthesis and control the hardness of synthetic negatives. Even though training with synthetic samples generated by the above methods can give a performance boost, they require additional generative networks alongside the main network. This can result in a larger model size, slower training time, and harder optimization [4]. The proposed method also generates samples to train with augmented information, while it does not require any additional generative networks and suffer from the above problems.

Query Expansion and Database Augmentation Query expansion (QE) in image retrieval has been proposed in [7, 6]. Given a query image feature, it retrieves a rank list of image features from a database that matches the query and combines the high ranked retrieved image features, along with the original query. Then, it re-queries the combined image features to retrieve an expanded set of matching images and repeats the process as necessary. Similar to query expansion, database augmentation (DBA) [29, 3] replaces every image feature in a database with a combination of itself and its neighbors, to improve the quality of image features. Our proposed embedding expansion is inspired by these concepts namely the combination of image features to augment image representations by leveraging the features of their neighbors. The key difference is that both techniques are used during post-processing, while the proposed method is used during the training phase. More specifically, the proposed method generates multiple combinations from the same class to augment semantic information for metric learning losses.

3. Preliminaries

This section introduces the mathematical formulation of the representative pair-based metric learning losses. We define a function f which projects data space \mathcal{D} to the embedding space \mathcal{X} by $f(\cdot; \theta) : \mathcal{D} \rightarrow \mathcal{X}$, where f is a neural network parameterized by θ . Feature points in the embedding space can be sampled as $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]$, where N is the number of feature points and each point \mathbf{x}_i has a label $\mathbf{y}[i] \in \{1, \dots, C\}$. \mathcal{P} is a set of positive pairs among the feature points.

Triplet loss [36, 22] considers triplet of points and pulls the anchor point closer to the positive point of the same class than to the negative point of the different class by a fixed margin m :

$$L_{\text{triplet}} = \frac{1}{|\mathcal{P}|} \sum_{\substack{(i,j) \in \mathcal{P} \\ k: \mathbf{y}[k] \neq \mathbf{y}[i]}} [d_{i,j}^2 - d_{i,k}^2 + m]_+, \quad (1)$$

where $[\cdot]_+$ is a hinge function and $d_{i,j} = \|\mathbf{x}_i - \mathbf{x}_j\|_2$ is the Euclidean distance between embedding \mathbf{x}_i and \mathbf{x}_j . Triplet loss is usually used with applying L_2 -normalization to the embedding feature [22].

Lifted structured loss [21] is proposed to take full advantage of the training batches in the neural network training. Given a training batch, it aims to pull one positive point as close as possible and pushes all negative points corresponding to the positive points farther than a fixed margin of m :

$$L_{lifted} = \frac{1}{2|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left[\log \left\{ \sum_{k: \mathbf{y}[k] \neq \mathbf{y}[i]} \exp(m - d_{i,k}) \right\} + \sum_{k: \mathbf{y}[k] \neq \mathbf{y}[j]} \exp(m - d_{j,k}) \right] + d_{i,j} \quad (2)$$

Similar to triplet loss, lifted structured loss also uses L_2 -normalization to the embedding feature [21].

N-pair loss [25] allows joint comparison among more than one negative points to generalize triplet loss. More specifically, it aims to pull one positive pair and push away $N - 1$ negative points from $N - 1$ negative classes:

$$L_{npair} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left\{ \log \left[1 + \sum_{k: \mathbf{y}[k] \neq \mathbf{y}[i]} \exp(s_{i,k} - s_{i,j}) \right] \right\}, \quad (3)$$

where $s_{i,j} = \mathbf{x}_i^T \mathbf{x}_j$ is the similarity of embedding points \mathbf{x}_i and \mathbf{x}_j . N-pair loss does not apply L_2 -normalization to the embedding features because it leads to optimization difficulty for the loss [25]. Instead, it regularizes the L_2 norm of the embedding features to be small.

Multi-Similarity loss [35] (MS loss) is one of the latest works for metric learning loss. It is proposed to jointly measure both self-similarity and relative similarities of a pair, which enables the model to collect and weight informative pairs. MS loss performs pair mining for both positive and negative pairs. A negative pair of $\{\mathbf{x}_i, \mathbf{x}_j\}$ is selected with the condition of

$$s_{i,j}^- > \min_{\mathbf{y}[k] = \mathbf{y}[i]} s_{i,k} - \epsilon, \quad (4)$$

and a positive pair of $\{\mathbf{x}_i, \mathbf{x}_j\}$ is selected with the condition of

$$s_{i,j}^+ < \max_{\mathbf{y}[k] \neq \mathbf{y}[i]} s_{i,k} + \epsilon, \quad (5)$$

where the ϵ is a given margin. For an anchor x_i , we denote the index set of its selected positive and negative pairs as $\bar{\mathcal{P}}_i$ and $\bar{\mathcal{N}}_i$, respectively. Then, MS loss can be formulated as:

$$L_{ms} = \frac{1}{N} \sum_{i=1}^N \left[\frac{1}{\alpha} \log \left\{ 1 + \sum_{k \in \bar{\mathcal{P}}_i} e^{-\alpha(s_{i,k} - \lambda)} \right\} + \frac{1}{\beta} \log \left\{ \sum_{k \in \bar{\mathcal{N}}_i} e^{\beta(s_{i,k} - \lambda)} \right\} \right], \quad (6)$$

where α , β , and λ are hyper-parameters, and N denotes the number of training samples. MS loss uses L_2 -normalization on the embedding features.

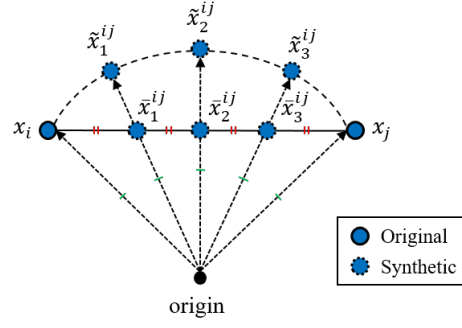


Figure 2. Illustration of generating synthetic points. Given two feature points $\{x_i, x_j\}$ from the same class, embedding expansion generates synthetic points which are internally dividing points into $n + 1$ equal parts $\{\bar{x}_1^{ij}, \bar{x}_2^{ij}, \bar{x}_3^{ij}\}$, where $n = 3$ in the figure. For the metric learning losses that use L_2 -normalization, the synthetic points are applied with L_2 -normalization and generates $\{\tilde{x}_1^{ij}, \tilde{x}_2^{ij}, \tilde{x}_3^{ij}\}$. Circles with plane line are original points and circles with dotted line are synthetic points.

4. Embedding Expansion

This section introduces the proposed embedding expansion consisting of two steps: synthetic points generation and hard negative pair mining.

4.1. Synthetic Points Generation

QE and DBA techniques in image retrieval generate synthetic points by combining feature points in an embedding space in order to exploit additional relevant information [7, 6, 29, 3]. Inspired by these techniques, the proposed embedding expansion generates multiple synthetic points by combining feature points from the same class in an embedding space to augment information for the metric learning losses. To be specific, embedding expansion performs linear interpolation in a linear interpolant between two feature points and generates synthetic points that are internally dividing points into $n + 1$ equal parts, as illustrated in Figure 2.

Given two feature points $\{\mathbf{x}_i, \mathbf{x}_j\}$ from the same class in an embedding space, the proposed method generates internally dividing points $\bar{\mathbf{x}}_k^{ij}$ into $n + 1$ equal parts and obtains a set of the synthetic points $\bar{\mathcal{S}}^{ij}$ as:

$$\bar{\mathbf{x}}_k^{ij} = \frac{k\mathbf{x}_i + (n - k)\mathbf{x}_j}{n}, \quad (7)$$

$$\bar{\mathcal{S}}^{ij} = \{\bar{\mathbf{x}}_1^{ij}, \bar{\mathbf{x}}_2^{ij}, \dots, \bar{\mathbf{x}}_n^{ij}\}, \quad (8)$$

where n is the number of points to generate. For the metric learning losses that use L_2 -normalization, such as triplet loss, lifted structured loss, and MS loss, L_2 -normalization

has to be applied to the synthetic points:

$$\tilde{\mathbf{x}}_k^{ij} = \frac{\tilde{\mathbf{x}}_k^{ij}}{\|\tilde{\mathbf{x}}_k^{ij}\|_2}, \quad (9)$$

$$\tilde{\mathcal{S}}^{ij} = \{\tilde{\mathbf{x}}_1^{ij}, \tilde{\mathbf{x}}_2^{ij}, \dots, \tilde{\mathbf{x}}_n^{ij}\}, \quad (10)$$

where $\tilde{\mathbf{x}}_k^{ij}$ is a L_2 -normalized synthetic point, and $\tilde{\mathcal{S}}^{ij}$ is a set of the L_2 -normalized synthetic points. These L_2 -normalized synthetic points will be located on the hypersphere space with the same norm. The way of generating synthetic points shares a similar spirit with mixup augmentation methods [41, 28, 31], and the comparison is given in supplementary material.

There are three advantages of generating points that are internally dividing points into $n + 1$ equal parts in an embedding space. (i) Given a pair of feature points from each class in well-clustered embedding space, the similarity of the hardest negative pair will be the shortest distance between line segments of each pair from each class (i.e., $\overline{\mathbf{x}_i \mathbf{x}_j} \leftrightarrow \overline{\mathbf{x}_k \mathbf{x}_l}$ in Figure 1). However, it is computationally expensive to compute the shortest distance between segments of finite length in a high-dimensional space [18, 24]. Instead, by computing distances between internally dividing points of each class, we can approximate the problem with less computation. (ii) The labels of synthetic points have a high degree of certainty because they are included inside the class cluster. Previous work [44] of sample generation method exploited a fully connected layer and softmax loss to control the labels of synthetic points, while the proposed method makes it certain by considering geometrical relations. We further investigate the certainty of labels of synthetic points with an experiment in Section 5.2.1. (iii) The proposed method of generating synthetic points requires a trivial amount of training speed and memory because we perform a simple linear interpolation in an embedding space. We further discuss the training speed and memory in Section 5.4.2.

4.2. Hard Negative Pair Mining

The second step of the proposed method is to perform hard negative pair mining among the synthetic and original points to ignore trivial pairs and train with informative pairs, as illustrated in Figure 1. The hard pair mining is only performed on negative pairs, and original points are used for positive pairs. The reason is that hard positive pair mining among original and synthetic points will always be a pair of original points because the synthetic points are internally dividing points of the pair. We formulate the combination of representative metric learning losses with the proposed embedding expansion.

EE + Triplet loss [36, 22] can be formulated by adding min-pooling on the negative pairs because the hardest pair

for triplet loss is a pair with the smallest Euclidean distance:

$$L_{triplet}^{EE} = \frac{1}{|\mathcal{P}|} \sum_{\substack{(i,j) \in \mathcal{P} \\ k: \mathbf{y}[k] \neq \mathbf{y}[i]}} \left[d_{i,j}^2 - \min_{(p,n) \in \tilde{\mathcal{N}}_{\mathbf{y}[i], \mathbf{y}[k]}} d_{p,n}^2 + m \right]_+, \quad (11)$$

where $\tilde{\mathcal{N}}_{\mathbf{y}[i], \mathbf{y}[k]}$ is a set of negative pairs with a positive point from the class $\mathbf{y}[i]$ and a negative point from the class $\mathbf{y}[k]$ including synthetic points.

EE + Lifted structured loss [21] also has to use min-pooling of Euclidean distance of negative pairs to add embedding expansion. The combined loss consists of minimizing the following hinge loss,

$$L_{lifted}^{EE} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left[\log \left\{ \sum_{k: \mathbf{y}[k] \neq \mathbf{y}[i]} \exp(m - \min_{(p,n) \in \tilde{\mathcal{N}}_{\mathbf{y}[i], \mathbf{y}[k]}} d_{p,n}) \right\} + d_{i,j} \right]_+^2. \quad (12)$$

EE + N-pair loss [25] can be formulated by using max-pooling on the negative pairs because the hardest pair for n-pair loss is a pair with the largest similarity, unlike triplet and lifted structured loss:

$$L_{npair}^{EE} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left\{ \log \left[1 + \sum_{k: \mathbf{y}[k] \neq \mathbf{y}[i]} \exp\left(\max_{(p,n) \in \tilde{\mathcal{N}}_{\mathbf{y}[i], \mathbf{y}[k]}} s_{p,n} - s_{i,j}\right) \right] \right\}. \quad (13)$$

EE + Multi-Similarity loss [35] contains two kinds of hard negative pair mining: one from the embedding expansion, and the other one from the MS loss. We integrate both hard negative pair mining by modifying the condition of Equation 4. A negative pair of $\{\mathbf{x}_i, \mathbf{x}_j\}$ is selected with the condition of

$$\max_{(p,n) \in \tilde{\mathcal{N}}_{\mathbf{y}[i], \mathbf{y}[j]}} s_{p,n} > \min_{\mathbf{y}[k] = \mathbf{y}[i]} s_{i,k} - \epsilon, \quad (14)$$

and we define the index set of selected negative pairs of an anchor x_i as $\tilde{\mathcal{N}}'_i$. Then, the combination of embedding expansion and MS loss can be formulated as:

$$L_{ms}^{EE} = \frac{1}{N} \sum_{i=1}^N \left[\frac{1}{\alpha} \log \left\{ 1 + \sum_{k \in \tilde{\mathcal{P}}_i} e^{-\alpha(s_{i,k} - \lambda)} \right\} + \frac{1}{\beta} \log \left\{ \sum_{k \in \tilde{\mathcal{N}}'_i} e^{\beta(s_{i,k} - \lambda)} \right\} \right]. \quad (15)$$

5. Experiments

5.1. Datasets and Settings

Datasets We evaluate the proposed method with two small benchmark datasets (CUB200-2011 [32],

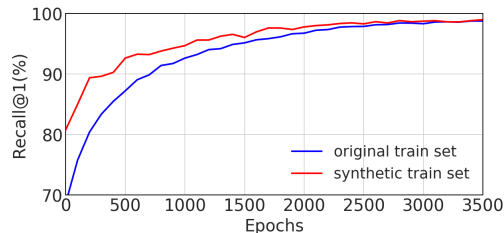
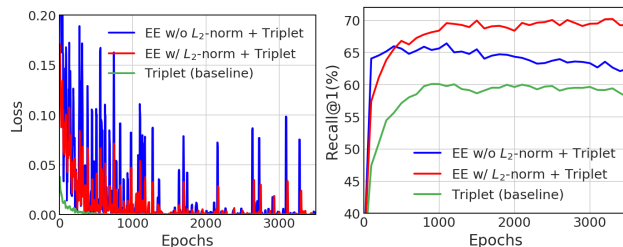


Figure 3. Recall@1(%) curve evaluated with original and synthetic points from the train set, trained with EE + triplet loss on CARS196.

CARS196 [16]), and one large benchmark dataset (Stanford Online Products [21]). We follow the conventional way of train and test splits used by [21, 44]. (i) CUB200-2011 [32] (CUB200) contains 200 different bird species with 11,788 images in total. The first 100 classes with 5,864 images are used for training, and the other 100 classes with 5,924 images are used for testing. (ii) CARS196 [16] contains 196 different types of cars with 16,185 images. The first 98 classes with 8,054 images are used for training, and the other 98 classes with 8,131 images are used for testing. (iii) Stanford Online Products [21] (SOP) is one of the largest benchmarks for the metric learning task. It consists of 22,634 classes of online products with 120,053 images, where 11,318 classes with 59,551 images are used for training, and the other 11,316 classes with 60,052 images are used for testing. For CUB200 and CARS196, we evaluate the proposed method without bounding box information.

Metrics Following the standard metrics in image retrieval and clustering [21, 34], we report the image clustering performance with F_1 and normalized mutual information (NMI) metrics [23] and image retrieval performance with Recall@K score.

Experimental Settings We implement our proposed method with the TensorFlow [1] framework on a Tesla P40 GPU with 24GB memory. Input images are resized to 256×256 , horizontally flipped, and randomly cropped to 227×227 . We use a 512-dimensional embedding size for all feature vectors. All models are trained with an ImageNet [8] pre-trained GoogLeNet [26] and a randomly initialized fully connected layer using the Xavier method [10]. We use the learning rate of 10^{-4} with the Adam optimizer [15] and set a batch size of 128 for every dataset. For the baseline metric learning loss, we use triplet loss with hard positive and hard negative mining (HPHN) [13, 38] and its combination of EE with the number of synthetic points $n = 2$ across all experiments, unless otherwise noted in the experiment.



(a) Loss value (b) Recall@1(%) performance

Figure 4. Loss value and recall@1(%) performance of training and test set from CARS196. It compares three models: triplet loss as baseline, EE without L_2 -normalization + triplet loss, and EE with L_2 -normalization + triplet loss.

5.2. Analysis of Synthetic Points

5.2.1 Labels of Synthetic Points

The main advantage of exploiting the internally dividing point is that the labels of synthetic points are expected to have a high degree of certainty because they are placed inside the class cluster. Thus, they can contribute to training a network as synthetic points with augmented information other than outliers. To investigate the certainty of the synthetic points during the training phase, we conduct an experiment that the synthetic and original points from the train set are evaluated at each epoch. For the evaluation of the synthetic points, we used the synthetic points as the query side and the original points as the database side. The score of synthetic points at the beginning is above 80%, which is enough for training, and it starts increasing by the training epoch. Overall, recall@1 of synthetic points from train sets are always higher than those of original points, and they maintain a high degree of certainty to be used as augmented feature points.

5.2.2 Impact of L_2 -normalization

Metric learning losses, such as triplet, lifted structured, and MS loss, apply L_2 -normalization to the last feature embeddings so that every feature embedding will be projected onto the hyper-sphere space with the same norm. Generating internally dividing points between L_2 -normalized feature embeddings will not be on the hyper-sphere space and will have a different norm. Thus, we proposed applying L_2 -normalization to the synthetic points to keep the continuity of the norm for these kinds of metric learning losses. To investigate the impact of L_2 -normalization, we conduct an experiment of EE with and without L_2 -normalization, including a baseline of triplet loss, as illustrated in Figure 4. Interestingly, EE without L_2 -normalization achieved better performance than the baseline. However, the model’s loss

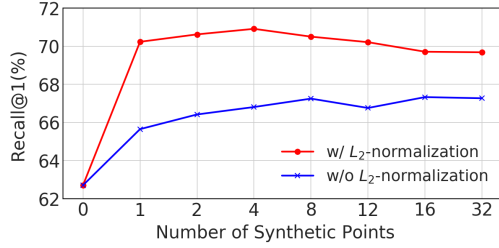


Figure 5. Recall@1(%) performance by the number of synthetic points with and without L_2 -normalization. Each model is trained with EE + triplet loss on the CARS196.

value fluctuates greatly, which can be caused by the different norms between original and synthetic points. The baseline and EE without L_2 -normalization start decreasing after peak points, which indicates the models are overfitting on the training set. On the other hand, the performance graph of EE with L_2 -normalization keeps increasing because of training with augmented information, which enables one to obtain a more generalized model.

5.2.3 Impact of Number of Synthetic Points

The number of synthetic points to generate is the sole hyper-parameter of our proposed method. As illustrated in Figure 5, we conduct an experiment by differentiating the number of synthetic points on EE with and without L_2 -normalization to see its impact. For EE without L_2 -normalization, the performance keeps increasing until about 8 synthetic points and maintains performance. In the case of the EE with L_2 -normalization, the peak of the performance is between 2 and 8 synthetic points, after which it starts decreasing. We speculate that it is because generating too many synthetic points can cause the model to be distracted by the synthetic points.

5.3. Analysis of Hard Negative Pair Mining

5.3.1 Selection Ratio of Synthetic Points

The proposed method performs hard negative pair mining among synthetic and original points in order to learn the metric learning loss with the most informative feature representations. To see the impact of synthetic points, we compute the ratio of synthetic and original points selected in the hard negative pair mining, as illustrated in Figure 6. At the beginning of training, more than 20% of synthetic points are selected for the hard negative pair. The ratio of synthetic points decreases as the clustering ability increases because many synthetic points are generated inside of the cluster. By increasing the number of n , the ratio of synthetic points increases. Throughout the training, a greater number of original points are selected than synthetic points.

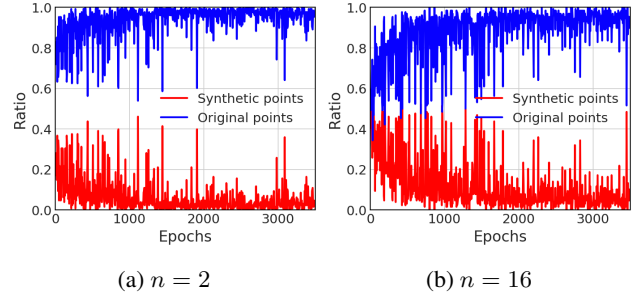


Figure 6. Ratio of synthetic and original points which are selected during hard negative pair mining of EE + triplet loss. We generate n synthetic points for EE and train the model with CARS196. The ratio of synthetic point is calculated as $ratio^{(syn)} = \frac{\# \text{ of synthetic}}{\# \text{ of synthetic} + \# \text{ of original}}$, while the ratio of original point is calculated as $ratio^{(ori)} = 1 - ratio^{(syn)}$.

This way, the synthetic points work as assistive augmented information instead of distracting the model training.

5.3.2 Effect of Hard Negative Pair Mining

We visualized distance matrices of triplet loss as a baseline, and EE + triplet loss to see the effect of hard negative pair mining as illustrated in Figure 7. By increasing the training epoch, the main diagonal of the heatmaps get redder, and the entries outside the diagonal get bluer in both triplet and EE + triplet loss. This indicates that the distances of positive pairs get smaller with smaller intra-class variation, and the distances of negative pairs get larger with a larger inter-class variation. In a comparison of triplet and EE + triplet loss on the same epoch, heatmaps of EE + triplet loss are filled with more yellow and red colors than the baseline of triplet, especially at the beginning of the training as shown in Figure 7e and Figure 7f. Even at the end of the training, as Figure 7h, the heatmap of EE + triplet loss still contains a greater number of hard negative pairs than triplet loss does. It shows that combining the proposed embedding expansion with the metric learning loss allows training with harder negative pairs with augmented information. A more detailed analysis of the training process is presented in supplementary material and video² with a t-SNE visualization [19] of embedding space at certain epoch.

5.4. Analysis of Model

5.4.1 Robustness

To see the improvement of the model robustness, we evaluate performance by putting occlusion on the input images in two ways. Center occlusion fills zeros in a center hole, and

²<https://youtu.be/5msMSXyQZ5U>

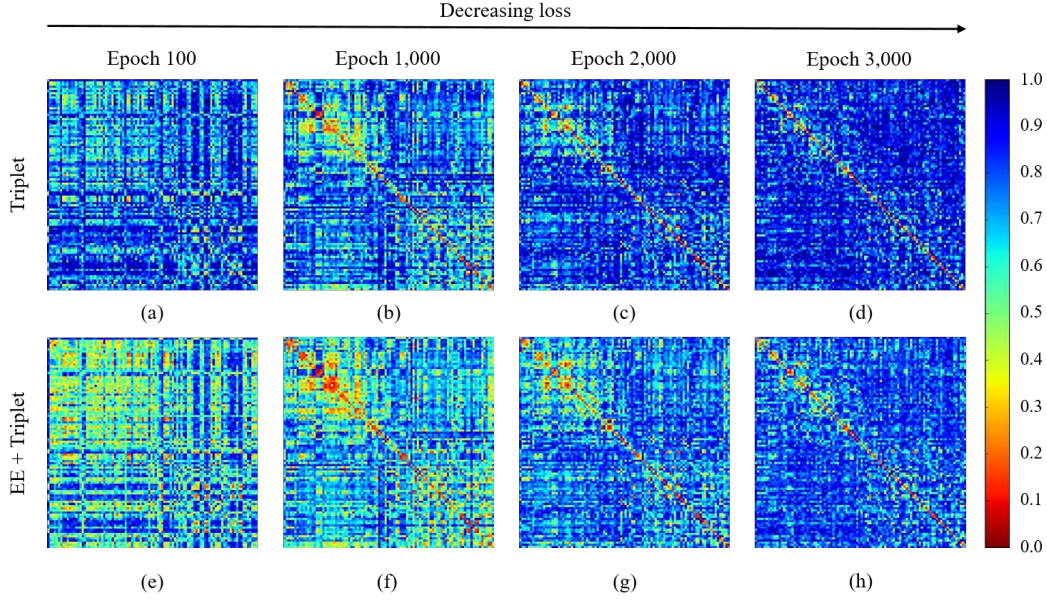
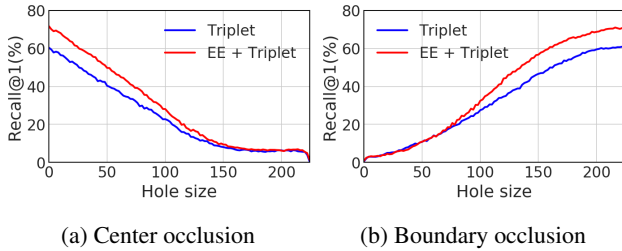


Figure 7. Comparison of Euclidean distance heatmaps between triplet and EE + triplet loss during training CARS196 dataset. In each heatmap, given two samples from each class, all the rows and columns are the first and the second samples from each class, respectively. The main diagonal is the distance of positive pairs, where the entries outside the diagonal are the distance of negative pairs. The smaller distance of negative pair (yellow and red) indicates the harder negative pairs, where all distance is normalized between 0 and 1.

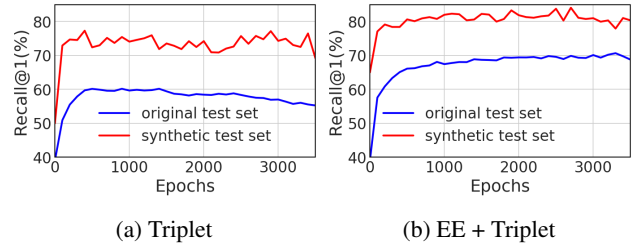


(a) Center occlusion (b) Boundary occlusion

Figure 8. Evaluation of occluded images on CARS196 test set.

boundary occlusion fills zeros outside of the hole. As shown in Figure 8, we measure Recall@1 performance by increasing the size of the hole from 0 to 227. The result shows that EE + triplet loss achieves significant improvements in robustness for both occlusion cases.

If embedding features are well clustered and contain the key representations of the class, the combination of embedding features from the same class would be included in the same class cluster with the same effect of QE [7, 6] and DBA [29, 3]. To see the robustness of the embedding features, we generate synthetic points by combining randomly selected test points of the same class as a query side and evaluate Recall@1 performance with original test points as a database side. As shown in Figure 9, EE + triplet loss



(a) Triplet (b) EE + Triplet

Figure 9. Evaluation of synthetic features from CARS196 test set.

improves the performance of evaluation with synthetic test sets compared to triplet loss, which shows that the feature robustness is improved by forming well-structured clusters.

5.4.2 Training Speed and Memory

Additional training time and memory of the proposed method are negligible. To see the additional training time of the proposed method, we compared the training time of baseline triplet and EE + triplet loss. As shown in Table 1, generating synthetic points takes from 0.0002 ms to 0.0023 ms longer. The total computing time of EE + triplet loss takes just 0.0038 ms to 0.0159 ms longer than the baseline ($n = 0$). Even though the number of points to generate in-

n	0	2	4	8	16	32
Gen	0	0.0002	0.0004	0.0008	0.0013	0.0023
Total	0.2734	0.2772	0.2775	0.2822	0.2885	0.2893

Table 1. Computation time (ms) of EE by the number of synthetic points n . Gen is time for generating synthetic points, and total is time for computing triplet loss, including generation and hard negative pair mining.

creases, the additional time for computing is negligible because it can be done by simple linear algebra. For memory requirements, if N embedding points with a $N \times N$ similarity matrix are necessary for computing the triplet loss, EE + triplet loss requires $(n + 1)N$ embedding points with a $(n + 1)N \times (n + 1)N$ similarity matrix, which are trivial.

5.5. Comparison with State-of-the-Art

We compare the performance of our proposed method with the state-of-the-art metric learning losses and previous sample generation methods on clustering and image retrieval tasks. As shown in Table 2, all combinations of EE with metric learning losses achieve significant performance boost on both tasks. The maximum improvements of Recall@1 and NMI are 12.1% and 7.4% from triplet loss in the CARS196 dataset, respectively. The minimum improvements of Recall@1 and NMI are 1.1% from MS loss in the CARS196 dataset and 0.1% from HPHN triplet loss in the SOP dataset, respectively. By comparison with previous sample generation methods, the proposed method outperforms every dataset and loss, except for the N-pair loss on the CARS196 dataset. In the large-scale datasets with enormous categories like SOP, the performance improvement of the proposed method is more competitive, even when taking those methods with generative networks into consideration. Performance with different capacity of networks and comparison with HTG are presented in the supplementary material.

6. Conclusion

In this paper, we proposed embedding expansion for augmentation in the embedding space that can be used with existing pair-based metric learning losses. We do so by generating synthetic points within positive pairs, and performing hard negative pair mining among synthetic and original points. Embedding expansion is simple, easy to implement, no computational overheads, but adjustable on various pair-based metric learning losses. We demonstrated that the proposed method significantly improves the performance of existing metric learning losses, and also outperforms the previous sample generation methods in both image retrieval and clustering tasks.

Method	Clustering		Retrieval			
	NMI	F1	R@1	R@2	R@4	R@8
Triplet	49.8	15.0	35.9	47.7	59.1	70.0
Triplet [†]	58.1	24.2	48.3	61.9	73.0	82.3
DAML (Triplet)	51.3	17.6	37.6	49.3	61.3	74.4
HDML (Triplet)	55.1	21.9	43.6	55.8	67.7	78.3
EE + Triplet	55.7	22.4	44.3	57.0	68.1	78.9
EE + Triplet [†]	60.5	27.0	51.7	63.5	74.5	82.5
N-pair	60.2	28.2	51.9	64.3	74.9	83.2
DAML (N-pair)	61.3	29.5	52.7	65.4	75.5	84.3
HDML (N-pair)	62.6	31.6	53.7	65.7	76.7	85.7
EE + N-pair	62.7	32.4	55.2	67.4	77.7	86.4
Lifted-Struct	56.4	22.6	46.9	59.8	71.2	81.5
EE + Lifted	61.2	28.2	54.2	66.6	76.7	85.2
MS	62.8	31.2	56.2	68.3	79.1	86.5
EE + MS	63.3	32.5	57.4	68.7	79.5	86.9

(a) CUB200-2011 dataset

Method	Clustering		Retrieval			
	NMI	F1	R@1	R@2	R@4	R@8
Triplet	52.9	17.9	45.1	57.4	69.7	79.2
Triplet [†]	57.4	22.6	60.3	73.4	83.5	90.5
DAML (Triplet)	56.5	22.9	60.6	72.5	82.5	89.9
HDML (Triplet)	59.4	27.2	61.0	72.6	80.7	88.5
EE + Triplet	60.3	25.1	57.2	70.5	81.3	88.2
EE + Triplet [†]	63.1	32.0	71.6	80.7	87.5	92.2
N-pair	62.7	31.8	68.9	78.9	85.8	90.9
DAML (N-pair)	66.0	36.4	75.1	83.8	89.7	93.5
HDML (N-pair)	69.7	41.6	79.1	87.1	92.1	95.5
EE + N-pair	63.4	32.6	72.5	81.1	87.6	92.5
Lifted-Struct	57.8	25.1	59.9	70.4	79.6	87.0
EE + Lifted	59.1	27.2	65.2	76.4	85.6	89.5
MS	62.4	30.2	75.0	83.1	89.5	93.6
EE + MS	63.5	33.5	76.1	84.2	89.8	93.8

(b) CARS196 dataset

Method	Clustering		Retrieval		
	NMI	F1	R@1	R@10	R@100
Triplet	86.3	20.2	53.9	72.1	85.7
Triplet [†]	91.4	43.3	75.5	88.8	95.4
DAML (Triplet)	87.1	22.3	58.1	75.0	88.0
HDML (Triplet)	87.2	22.5	58.5	75.5	88.3
EE + Triplet	87.4	24.8	62.4	79.0	91.0
EE + Triplet [†]	91.5	43.6	77.2	89.6	95.5
N-pair	87.9	27.1	66.4	82.9	92.1
DAML (N-pair)	89.4	32.4	68.4	83.5	92.3
HDML (N-pair)	89.3	32.2	68.7	83.2	92.4
EE + N-pair	90.6	38.8	73.5	87.5	94.4
Lifted-Struct	87.2	25.3	62.6	80.9	91.2
EE + Lifted	89.6	35.3	70.6	85.5	93.6
MS	91.3	42.7	75.9	89.3	95.6
EE + MS	91.9	46.1	78.1	90.3	95.8

(c) Stanford Online Products dataset

Table 2. Clustering and retrieval performance (%) on three benchmarks in comparison with other methods. [†] denotes the HPHN triplet loss, and bold numbers indicate the best score within the same loss.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016. [5](#)
- [2] Ejaz Ahmed, Michael Jones, and Tim K Marks. An improved deep learning architecture for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3908–3916, 2015. [2](#)
- [3] Relja Arandjelović and Andrew Zisserman. Three things everyone should know to improve object retrieval. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2911–2918. IEEE, 2012. [1](#), [2](#), [3](#), [7](#)
- [4] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017. [2](#)
- [5] Sumit Chopra, Raia Hadsell, Yann LeCun, et al. Learning a similarity metric discriminatively, with application to face verification. In *CVPR (1)*, pages 539–546, 2005. [1](#)
- [6] Ondřej Chum, Andrej Mikulík, Michal Perdoch, and Jiří Matas. Total recall ii: Query expansion revisited. In *CVPR 2011*, pages 889–896. IEEE, 2011. [1](#), [2](#), [3](#), [7](#)
- [7] Ondrej Chum, James Philbin, Josef Sivic, Michael Isard, and Andrew Zisserman. Total recall: Automatic query expansion with a generative feature model for object retrieval. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007. [1](#), [2](#), [3](#), [7](#)
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [5](#)
- [9] Yueqi Duan, Wenzhao Zheng, Xudong Lin, Jiwen Lu, and Jie Zhou. Deep adversarial metric learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2780–2789, 2018. [1](#), [2](#)
- [10] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010. [5](#)
- [11] Albert Gordo, Jon Almazán, Jerome Revaud, and Diane Larlus. Deep image retrieval: Learning global representations for image search. In *European conference on computer vision*, pages 241–257. Springer, 2016. [1](#)
- [12] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742. IEEE, 2006. [1](#)
- [13] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017. [2](#), [5](#)
- [14] John R Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe. Deep clustering: Discriminative embeddings for segmentation and separation. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 31–35. IEEE, 2016. [1](#)
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [16] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 554–561, 2013. [2](#), [5](#)
- [17] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. Deep-reid: Deep filter pairing neural network for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 152–159, 2014. [1](#)
- [18] Vladimir J Lumelsky. On fast computation of distance between line segments. *Information Processing Letters*, 21(2):55–61, 1985. [4](#)
- [19] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008. [6](#)
- [20] Hyun Oh Song, Stefanie Jegelka, Vivek Rathod, and Kevin Murphy. Deep metric learning via facility location. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5382–5390, 2017. [1](#)
- [21] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4004–4012, 2016. [1](#), [2](#), [3](#), [4](#), [5](#)
- [22] Florian Schroff, Dmitry Kalenichenko, and James 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. [1](#), [2](#), [4](#)
- [23] Hinrich Schütze, Christopher D Manning, and Prabhakar Raghavan. Introduction to information retrieval. In *Proceedings of the international communication of association for computing machinery conference*, page 260, 2008. [5](#)
- [24] William Griswold Smith. *Practical descriptive geometry*. McGraw-Hill, 1916. [4](#)
- [25] Kihyuk Sohn. Distance metric learning with n-pair loss, Aug. 10 2017. US Patent App. 15/385,283. [1](#), [3](#), [4](#)
- [26] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. [5](#)
- [27] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1701–1708, 2014. [1](#)
- [28] Yuji Tokozume, Yoshitaka Ushiku, and Tatsuya Harada. Between-class learning for image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5486–5494, 2018. [4](#)
- [29] Panu Turcot and David G Lowe. Better matching with fewer features: The selection of useful features in large database recognition problems. In *ICCV Workshops*, volume 2, 2009. [1](#), [2](#), [3](#), [7](#)

- [30] Evgeniya Ustinova and Victor Lempitsky. Learning deep embeddings with histogram loss. In *Advances in Neural Information Processing Systems*, pages 4170–4178, 2016. 1
- [31] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Aaron Courville, Ioannis Mitliagkas, and Yoshua Bengio. Manifold mixup: Learning better representations by interpolating hidden states. 2018. 4
- [32] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 2, 4, 5
- [33] Jiang Wang, Yang Song, Thomas Leung, Chuck Rosenberg, Jingbin Wang, James Philbin, Bo Chen, and Ying 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. 1
- [34] Jian Wang, Feng Zhou, Shilei Wen, Xiao Liu, and Yuanqing Lin. Deep metric learning with angular loss. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2593–2601, 2017. 1, 5
- [35] Xun Wang, Xintong Han, Weilin Huang, Dengke Dong, and Matthew R Scott. Multi-similarity loss with general pair weighting for deep metric learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5022–5030, 2019. 3, 4
- [36] Kilian Q Weinberger and Lawrence K Saul. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research*, 10(Feb):207–244, 2009. 1, 2, 4
- [37] Chao-Yuan Wu, R Manmatha, Alexander J Smola, and Philipp Krahenbuhl. Sampling matters in deep embedding learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2840–2848, 2017. 1, 2
- [38] Hong Xuan, Abby Stylianou, and Robert Pless. Improved embeddings with easy positive triplet mining. *arXiv preprint arXiv:1904.04370*, 2019. 5
- [39] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Deep metric learning for person re-identification. In *2014 22nd International Conference on Pattern Recognition*, pages 34–39. IEEE, 2014. 1
- [40] Rui Yu, Zhiyong Dou, Song Bai, Zhaoxiang Zhang, Yongchao Xu, and Xiang Bai. Hard-aware point-to-set deep metric for person re-identification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 188–204, 2018. 2
- [41] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 4
- [42] Yiru Zhao, Zhongming Jin, Guo-jun Qi, Hongtao Lu, and Xian-sheng Hua. An adversarial approach to hard triplet generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 501–517, 2018. 1, 2
- [43] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE international conference on computer vision*, pages 1116–1124, 2015. 1
- [44] Wenzhao Zheng, Zhaodong Chen, Jiwen Lu, and Jie Zhou. Hardness-aware deep metric learning. In *Proceedings of the*

IEEE Conference on Computer Vision and Pattern Recognition, pages 72–81, 2019. 1, 2, 4, 5