

Symmetrical Synthesis for Deep Metric Learning

Geonmo Gu,* Byungsoo Ko*

Clova Vision, NAVER Corp.
{korgm403, kobiso62}@gmail.com

Abstract

Deep metric learning aims to learn embeddings that contain semantic similarity information among data points. To learn better embeddings, methods to generate synthetic hard samples have been proposed. Existing methods of synthetic hard sample generation are adopting autoencoders or generative adversarial networks, but this leads to more hyper-parameters, harder optimization, and slower training speed. In this paper, we address these problems by proposing a novel method of synthetic hard sample generation called symmetrical synthesis. Given two original feature points from the same class, the proposed method firstly generates synthetic points with each other as an axis of symmetry. Secondly, it performs hard negative pair mining within the original and synthetic points to select a more informative negative pair for computing the metric learning loss. Our proposed method is hyper-parameter free and plug-and-play for existing metric learning losses without network modification. We demonstrate the superiority of our proposed method over existing methods for a variety of loss functions on clustering and image retrieval tasks.

1 Introduction

The objective of deep metric learning is to learn an embedding space where semantically similar images are embedded close together, and semantically dissimilar images are embedded far apart. Many recent deep metric learning approaches are built on similarity or distance between pairs of samples. Contrastive loss (Chopra et al. 2005) and triplet loss (Weinberger and Saul 2009) are both conventional losses that consider pair and triplet feature points, respectively. Recent works (Sohn 2016; Oh Song et al. 2016; Wang et al. 2017) have modified the structures of loss functions to contain richer information by considering multiple feature points, and they have achieved competitive performances. Along with the importance of the loss function, sampling strategy is also known to be essential for effective training. Different sampling strategies can lead to drastically different performances for the same loss function.

* Authors contributed equally.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

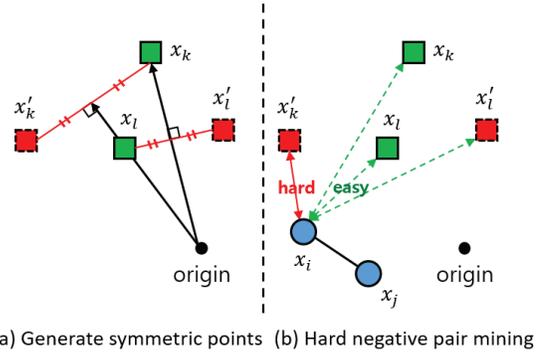


Figure 1: Illustration of our proposed symmetrical synthesis with two steps. First, given positive points (x_i, x_j) and negative points (x_k, x_l) in an embedding space, the negative points generate their synthetic points (x'_k, x'_l) with each other as an axis of symmetry. Secondly, it selects the hardest negative pair within the four feature points: two original points and two synthetic points. In the figure, x'_k will be selected. Rectangles and circles represent two different classes. Green and blue points are original features while red points with dotted boundary are synthetic features.

This has motivated recent works to focus on sampling strategies, such as hard negative pair mining (Hermans, Beyer, and Leibe 2017), semi-hard negative pair mining (Schroff, Kalenichenko, and Philbin 2015), and soft-hard mining (Yu et al. 2018). However, mining strategies could lead to a biased model because they usually account for a small selected minority and a large non-selected majority (Wu et al. 2017; Schroff, Kalenichenko, and Philbin 2015; Zheng et al. 2019).

To address this problem, recent works (Duan et al. 2018; Zhao et al. 2018; Zheng et al. 2019) have proposed using generative adversarial networks and autoencoders to generate synthetic hard samples. These methods enable non-selected majorities to be exploited by synthesizing them into hard samples and training a model with augmented information. Despite a performance boost of deep metric learning, it also suffers from several limitations. First, along with a

model for metric learning, an additional sub-network is required to generate synthetic hard samples, which leads to an increase in model size, hyper-parameters, and training time. Moreover, deploying a generative model, such as a generative adversarial network, can result in optimization difficulty (Arjovsky, Chintala, and Bottou 2017).

In this paper, we propose a simple yet powerful method for synthetic hard sample generation called symmetrical synthesis to address aforementioned limitations. As illustrated in Figure 1, given two feature points within the same class, our proposed method generates symmetrical synthetic points with each other as an axis of symmetry. Then, it selects the hardest negative pair within the original and synthetic points. This allows to train a model to push away samples of different classes with a stronger power. In contrast to previous methods, as our method only requires simple algebraic computation to generate synthetic points, it is hyper-parameter free and can be applied to existing metric learning losses in a plug-and-play manner, without any modification of network architecture. Furthermore, deploying our proposed method does not influence the training speed and optimization difficulty. We demonstrate that deploying our proposed method gives a significant improvement in image clustering and retrieval tasks on CUB-200-2011, CARS196, and Stanford Online Products. The proposed method outperforms previous methods by wide margins.

2 Related Work

Our work is related to three lines of active research: (1) metric learning, (2) hard negative pair mining, and (3) hard sample generation.

Metric Learning Metric learning losses have been proposed based on similarity and distance using the feature representation. One of the simplest losses is the triplet loss (Weinberger and Saul 2009), which takes triplets of samples to separate the negative pair more than the positive pair with a fixed relative margin. Despite its success, it has been reported to require expensive sampling methods to provide non-trivial samples for an efficient training (Chechik et al. 2010; Cui et al. 2016). To address this problem, N-pair loss (Sohn 2016) is proposed to expand the idea of triplet loss by considering $N - 1$ negative samples of different classes. Similar to N-pair loss, lifted structure loss (Oh Song et al. 2016) is proposed to train the embedding function by incorporating all negative samples within a batch. Angular loss (Wang et al. 2017) considers that the distance metric is sensitive to scale and only considers second-order information between samples. To circumvent these problems, angular loss constraints the angle at the negative point of triplet triangles.

Hard Negative Pair Mining Hard negative pair mining has played an essential role in the performance of deep metric learning. The purpose of this strategy is to progressively select false positive samples, which can give more information in the training process. For example, offline hard negative pair mining (Ahmed, Jones, and Marks 2015) is

proposed to iteratively fine-tune a model with hard negative samples selected by a previously trained model. Online hard negative pair mining (Hermans, Beyer, and Leibe 2017) proposes the selection of the hardest positive and negative within a batch to compute the triplet loss. Semi-hard negative pair mining (Schroff, Kalenichenko, and Philbin 2015) is proposed to avoid too confusing samples, such as the hardest positives and negatives, which may often be noise in data. One of the limitations is that mining strategies usually focus on the selected minority and overlook the non-selected majority, which can lead to a biased model (Wu et al. 2017; Schroff, Kalenichenko, and Philbin 2015; Zheng et al. 2019).

Hard Sample Generation Recently, there have been attempts to generate synthetic hard samples for exploiting a large number of easy negatives and training a model with extra semantic information. For example, the deep adversarial metric learning (DAML) framework (Duan et al. 2018) is proposed to generate synthetic hard samples from the easy negative samples in an adversarial manner. Similarly, an adversarial network for hard triplet generation (Zhao et al. 2018) is proposed to train a model with synthetic hard samples. The hardness-aware deep metric learning (HDML) framework (Zheng et al. 2019) exploits an autoencoder architecture to generate label-preserving synthetics in the embedding space and manipulate their hard levels. Nevertheless, all above-mentioned methods require additional generative networks, which result in a bigger model, slower training speed, and more hyper-parameters. Our work re-defines the core component of generation by geometrical approach with simple algebraic computation in the embedding space instead of using the generative networks. We show it can be easily used to existing metric learning losses without additional hyper-parameter, training speed decrease, and network modification.

3 Proposed Method

In this section, we present a novel method of synthetic hard sample generation called symmetrical synthesis (Symm). As illustrated in Figure 1, the proposed method follows two steps: (1) symmetrical synthetic generation and, (2) hard negative pair mining.

3.1 Symmetrical Synthesis

The first step for the proposed method is to generate symmetrical synthetic points in the embedding space. Let \mathcal{I} be the data space and \mathcal{X} be the d -dimensional embedding space. We define $f : \mathcal{I} \xrightarrow{f} \mathcal{X}$ be the mapping from the data space to the embedding space parameterized by a deep neural network. We sample a set of feature points $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]$ where each point \mathbf{x}_i has label $l_i \in \{1, \dots, C\}$. As illustrated in Figure 2, given two feature points $(\mathbf{x}_k, \mathbf{x}_l)$ from the same class, synthetic points $(\mathbf{x}'_k, \mathbf{x}'_l)$ can be generated with each other as an axis of symmetry. In order to get \mathbf{x}'_k , we define \mathbf{r}^l_k , which is a projection of \mathbf{x}_k onto \mathbf{x}_l , i.e.,

$$\mathbf{r}^l_k = (\mathbf{x}_k \cdot \mathbf{u}_{\mathbf{x}_l}) \mathbf{u}_{\mathbf{x}_l}, \quad (1)$$

Similarly to the triplet loss, the combination of symmetrical synthesis and lifted structure loss can be formulated by using min-pooling as follows:

$$\mathcal{L}_{lifted}^{Symm} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left[\log \left\{ \sum_{k:l_i \neq l_k} \exp \left(m - \min_{(p,n) \in \tilde{\mathcal{N}}_{l_i, l_k}} D_{p,n} \right) + D_{i,j} \right\} \right]^2, \quad (6)$$

For N-pair loss, additional negative samples are considered into triplets, and the triplet is turned into an N-tuplet. The loss is defined as:

$$\mathcal{L}_{npair} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left\{ \log \left[1 + \sum_{k:l_i \neq l_k} \exp(S_{i,k} - S_{i,j}) \right] \right\}, \quad (7)$$

where $S_{i,j} = \mathbf{x}_i^T \mathbf{x}_j$ is the similarity between embedding \mathbf{x}_i and \mathbf{x}_j . We formulate N-pair loss with symmetrical synthesis by adding max-pooling because of cosine similarity, and perform hard negative pair mining on every negative class in a mini-batch:

$$\mathcal{L}_{npair}^{Symm} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left\{ \log \left[1 + \sum_{k:l_i \neq l_k} \exp \left(\max_{(p,n) \in \tilde{\mathcal{N}}_{l_i, l_k}} S_{p,n} - S_{i,j} \right) \right] \right\}, \quad (8)$$

Angular loss is proposed to encode the third-order relation to triplet in terms of the angle at the negative point:

$$\mathcal{L}_{ang} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left\{ \log \left[1 + \sum_{k:l_i \neq l_k} \exp(f_{i,j,k}^n - f_{i,j}^p) \right] \right\}, \quad (9)$$

where $f_{i,j}^p = 2(1 + \tan^2 \alpha) \mathbf{x}_i^T \mathbf{x}_j$ and $f_{i,j,k}^n = 4 \tan^2 \alpha (\mathbf{x}_i + \mathbf{x}_j)^T \mathbf{x}_k$. Similarly to the N-pair loss, we can combine the symmetrical synthesis with angular loss by adding max-pooling for hard negative pair mining on every negative class as follows:

$$\mathcal{L}_{ang}^{Symm} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \left\{ \log \left[1 + \sum_{k:l_i \neq l_k} \exp \left(\max_{(p,q,r) \in \tilde{\mathcal{N}}_{l_i, l_k}} f_{p,q,r}^n - f_{i,j}^p \right) \right] \right\}, \quad (10)$$

where $\tilde{\mathcal{N}}_{l_i, l_k}$ is the set of triplets with two positive points from the class l_i and one negative point from the class l_k which is utilized in $f_{i,j,k}^n$.

Metric learning with the proposed symmetrical synthesis has two effects. First, using synthetic feature points leads to a more generalized model, because trivial samples, which could have been ignored by mining strategies, can be exploited by generating synthetic points and training the model with augmented information. Secondly, hard negative pair mining within the original and synthetic points allows metric learning losses to push away between different classes with greater force. This leads to higher inter-class variation with better clustering in the embedding space.

4 Experiments

In this section, we report experimental results from the proposed symmetrical synthesis on both image clustering and retrieval tasks. To evaluate quantitative performance, we use the standard F_1 and NMI metrics (Manning, Raghavan, and Schütze 2010) for the image clustering task, and Recall@K score for the image retrieval task.

4.1 Datasets

We evaluate our proposed method on the widely used three benchmarks by following the conventional protocol of train and test splits used by (Zheng et al. 2019; Oh Song et al. 2016). (1) CUB-200-2011 (CUB200) (Wah et al. 2011) has 11,788 images of 200 bird species, where the first 5,864 images of 100 species are used for training and the remaining 5,924 images of 100 species are used for testing. (2) CARS196 (Krause et al. 2013) has 16,185 car images of 196 classes. We use the first 8,054 images of 98 classes for training and the remaining 8,131 images of 98 classes for testing. (3) Standard Online Products (SOP) (Oh Song et al. 2016) datasets contains 120,053 product images of 22,634 classes, where the first 59,551 images of 11,318 classes are used for training and the remaining 60,502 images of 11,316 classes are used for testing. For CUB200 and CARS196, our method is evaluated without the bounding box information.

4.2 Experimental Setting

Throughout the experiments, TensorFlow (Abadi et al. 2016) framework is used on a Tesla P40 GPU with 24GB memory. All images are normalized to 256×256 , horizontal flipped and randomly cropped to 227×227 . The embedding size is set to 512-dimensional for all feature vectors. Triplet and lifted structure loss use l_2 -normalized features with Euclidean distance, and N-pair and angular loss use non- l_2 -normalized features with cosine similarity. We use ImageNet (Deng et al. 2009) pre-trained GoogLeNet (Szegedy et al. 2015) and the Xavier method (Glorot and Bengio 2010) to random initialize a fully connected layer. We set the learning rate to 10^{-4} with the Adam optimizer (Kingma and Ba 2014). The batch size of 128 is used for every dataset.

4.3 Experimental Results

We perform experiments to analyze the effect of our proposed method. The following experiments are conducted on the CARS196 dataset with N-pair loss in the image clustering and retrieval task.

Impact of Similarity and Norm As mentioned in Section 3.1, we generate synthetics with symmetric property so that the similarity and norm can be maintained. To see the impact of similarity and norm, we conduct experiments by differentiating α and β in Eq. 2. As illustrated in Figure 2, differentiating α gives a point with different cosine similarity and norm, but we force the norm to be same with the original point by multiplying $\|x_k\|/\|x'_k\|$ for the experiments. With the same norm, larger cosine similarity ($\alpha = 1.5$, $\beta = 1.0$) is not trainable and smaller cosine similarity ($\alpha = 2.5$, $\beta = 1.0$) results in dramatic performance reduction. Differentiating β gives a point with the same cosine

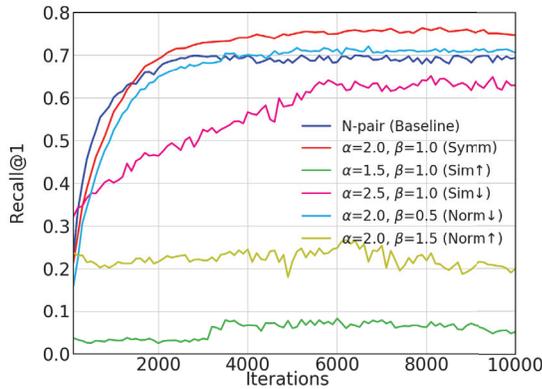


Figure 4: Recall@1 curve for comparison of different similarity and norm. The baseline is N-pair loss, while the rest is Symm + N-pair loss with different α and β on CARS196 dataset.

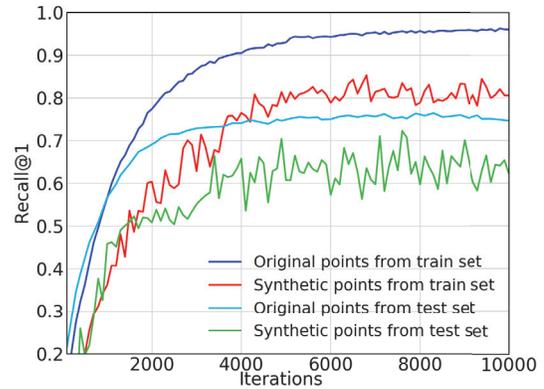


Figure 6: Recall@1 curve for comparison of original points and synthetic points from train and test set. It is trained and evaluated with Symm + N-pair loss on CARS196 dataset.

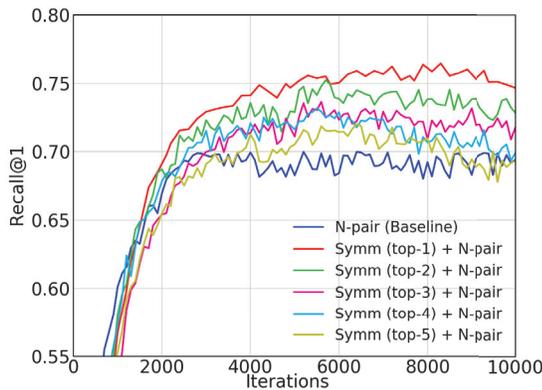


Figure 5: Recall@1 curve for comparison of top- k hard negative pair mining. It is trained and evaluated with N-pair loss for baseline and Symm + N-pair loss on CARS196 dataset.

similarity and different norm. With the same cosine similarity, larger norm ($\alpha = 2.0, \beta = 1.5$) is not trainable and smaller norm ($\alpha = 2.0, \beta = 0.5$) shows similar performance with N-pair, but lower than the proposed symmetrical synthetics ($\alpha = 2.0, \beta = 1.0$). This demonstrates that maintaining similarity and norm by generating symmetrical synthetics is essential for network optimization and converged performance.

Level of Hardness In the proposed symmetrical synthesis, the similarity of every possible negative pair is computed and the hardest negative pair is selected. This hard negative pair mining strategy is designed to use the most informative pairs for training. To analyze the effect of the hard negative pair mining, we conduct experiments to compare among the proposed methods with different top- k hardest negative pair mining with N-pair loss and the baseline

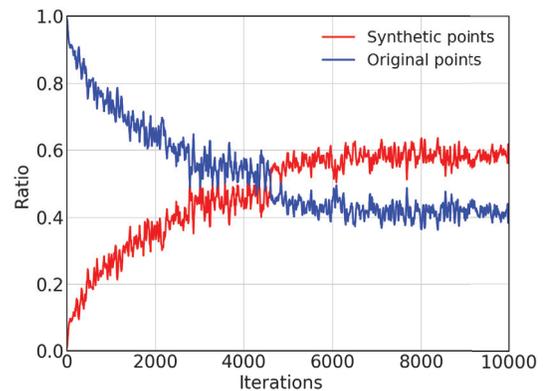


Figure 7: Ratio of selected feature points during hard negative pair mining between original and synthetic points. It is trained with Symm + N-pair loss on CARS196 dataset.

model with only N-pair loss. Figure 5 shows the learning curves of each model setting in the retrieval task. We observe that the harder the pair selected, the higher the performance, where every model with the proposed method outperforms the baseline model in both tasks. This is because harder pairs are more informative; thus, using the symmetrical synthesis with harder pairs pushes different classes away from each other with stronger power.

Label of Synthetics As shown in Figure 6, we conduct experiments to estimate where the synthetic points are generated. We evaluate the recall performance of the original points and synthetic points from the training and test sets. Synthetic points are generated from the original points with randomly selected points in the same class as the axis of symmetry. We speculate that synthetic points do not have to be inside the same cluster because they will only be used as hard samples to push the other class with stronger power.

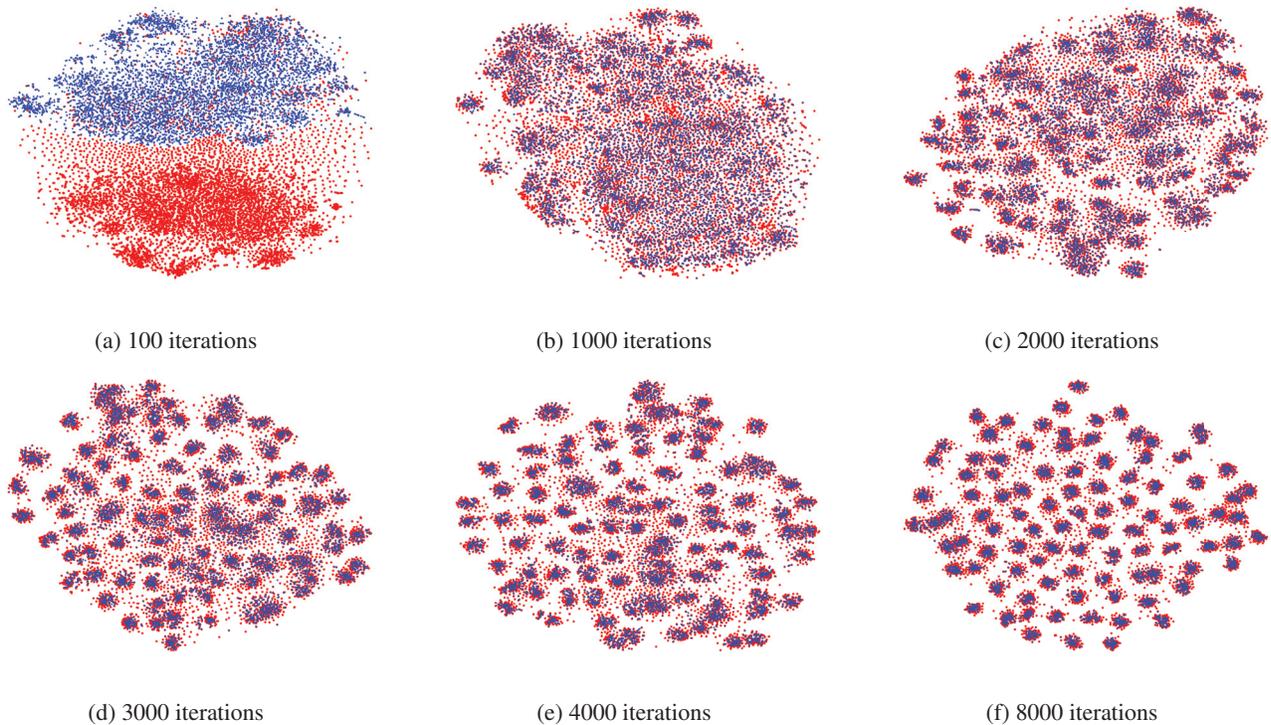


Figure 8: A t-SNE visualization of Symm + N-pair loss with the original (blue) and symmetrical synthetic (red) feature points from the training set of CARS196 dataset.

However, we expect the synthetic points would lie around the boundary of the same cluster to work as hard samples. We observe that the curve of synthetic points has more fluctuation than that of the original points in both the training and test sets. Besides, the performance of the original points is always higher than that of the synthetic points. We believe this is because a high portion of synthetic points is generated around the boundary of the cluster working as hard samples. Also, the performance increase of synthetic points indicates that synthetic points lie in meaningful spots due to the increased clustering ability of the model.

Visualization and Ratio of Feature Points The ideal place of generated symmetrical synthetic points is around the boundary of the class cluster so that the synthetic points can work as hard feature points during hard negative pair mining. To see the geographical location of the symmetrical synthetic points compared to the original points, we visualize the embedding space of each training step with the Barnes-Hut t-SNE (Van Der Maaten 2014), as shown in Figure 8. Moreover, we conduct an experiment to see how many original and synthetic feature points are selected to be the hardest negative pair in each training step, as illustrated in Figure 7. The ratio of synthetic points is calculated as $ratio^{(syn)} = \frac{\# \text{ of synthetic}}{\# \text{ of original} + \# \text{ of synthetic}}$, while the ratio of original points is calculated as $ratio^{(ori)} = 1 - ratio^{(syn)}$.

At the beginning of the training, original feature points are scattered without forming clusters, and the similarities

of positive pairs will be relatively small, as shown in Figure 8a. This causes the symmetrical synthetic points are generated on the meaningless place far from the positive pairs, that will be hardly selected during hard negative pair mining. Hence, the original points are selected mostly over the synthetic points at first, as illustrated in Figure 7. The better clustering ability the model has, the higher the chance that the synthetic points will be generated around the boundary of the cluster to become the hard feature points. Generated synthetic points start lying around the boundary of the class cluster from the 3000 steps, as shown in Figure 8d. After the 4000 steps, more than half of synthetic points are selected over the original points during hard negative pair mining, as illustrated in Figure 7. These selected synthetic points will work as hard negatives to train the model with richer information. Finally, we obtain clean and well-clustered embeddings of the original feature points, as shown in Figure 8f. More details of visualization and ratio of feature points are given in the supplementary video.

Training Speed and Memory The computational cost and memory consumption of symmetrical synthesis are negligible. On a Tesla P40 GPU with a batch size of 128, forward and backward pass of training the baseline N-pair loss takes 8.852×10^{-1} seconds, while Symm + N-pair takes 8.866×10^{-1} seconds per batch. In detail, computing the baseline N-pair loss takes 0.2454 ms, while generating symmetrical synthetic points, hard negative pair mining, and computing N-pair loss takes only 0.2497 ms. For mem-

| 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 [†] | 53.4 | 17.9 | 40.6 | 52.3 | 64.2 | 75.0 |
| 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 |
| Symm+Triplet | 59.6 | 26.2 | 51.4 | 63.0 | 74.4 | 84.1 |
| Symm+Triplet [†] | 63.3 | 32.1 | 55.0 | 67.3 | 77.5 | 86.0 |
| 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 |
| Symm+N-pair | 63.6 | 32.5 | 55.9 | 67.6 | 78.3 | 86.2 |
| Angular | 61.0 | 30.2 | 53.6 | 65.0 | 75.3 | 83.7 |
| Symm+Angular | 62.3 | 30.5 | 54.9 | 66.9 | 77.3 | 86.0 |
| Lifted-Struct | 56.4 | 22.6 | 46.9 | 59.8 | 71.2 | 81.5 |
| Symm+Lifted | 62.1 | 28.7 | 54.9 | 66.4 | 76.4 | 85.3 |

Table 1: Experimental results (%) of clustering and retrieval performance on CUB200-2011 dataset in comparison with other methods. [†] denotes the semi-hard triplet.

| 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 [†] | 55.7 | 22.4 | 53.2 | 65.4 | 74.3 | 83.6 |
| 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 |
| Symm+Triplet | 62.4 | 31.8 | 69.7 | 78.7 | 86.1 | 91.4 |
| Symm+Triplet [†] | 61.7 | 31.1 | 68.5 | 78.5 | 85.8 | 90.9 |
| 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 |
| Symm+N-pair | 66.3 | 36.6 | 76.5 | 84.3 | 90.4 | 94.1 |
| Angular | 62.4 | 31.8 | 71.3 | 80.7 | 87.0 | 91.8 |
| Symm+Angular | 66.1 | 35.9 | 75.5 | 84.0 | 90.0 | 94.0 |
| Lifted-Struct | 57.8 | 25.1 | 59.9 | 70.4 | 79.6 | 87.0 |
| Symm+Lifted | 59.9 | 28.5 | 66.6 | 77.2 | 84.7 | 89.9 |

Table 2: Experimental results (%) of clustering and retrieval performance on CARS196 dataset in comparison with other methods. [†] denotes the semi-hard triplet.

| Method | Clustering | | Retrieval | | |
|---------------------------|-------------|-------------|-------------|-------------|-------------|
| | NMI | F1 | R@1 | R@10 | R@100 |
| Triplet | 86.3 | 20.2 | 53.9 | 72.1 | 85.7 |
| Triplet [†] | 86.7 | 22.1 | 57.8 | 75.3 | 88.1 |
| DAML (Triplet) | 87.1 | 22.3 | 58.1 | 75.0 | 88.0 |
| HDML (Triplet) | 87.2 | 22.5 | 58.5 | 75.5 | 88.3 |
| Symm+Triplet | 88.9 | 30.6 | 65.7 | 81.4 | 91.7 |
| Symm+Triplet [†] | 89.5 | 33.9 | 68.5 | 82.4 | 91.3 |
| 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 |
| Symm+N-pair | 90.7 | 38.7 | 73.2 | 86.7 | 94.8 |
| Angular | 87.8 | 26.5 | 67.9 | 83.2 | 92.2 |
| Symm+Angular | 90.5 | 38.4 | 73.1 | 86.6 | 94.0 |
| Lifted-Struct | 87.2 | 25.3 | 62.6 | 80.9 | 91.2 |
| Symm+Lifted | 90.4 | 38.2 | 72.3 | 86.6 | 94.2 |

Table 3: Experimental results (%) of clustering and retrieval performance on SOP dataset in comparison with other methods. [†] denotes the semi-hard triplet.

ory consumption of symmetrical synthesis, our proposed

method requires the same size of additional matrix as the original feature points to save synthetic feature points and 16 times a similarity matrix for saving 16 possible positive and negative pairs, which are trivial.

Comparison with State-of-the-Art We compare our proposed method with famous metric learning losses including triplet loss, triplet loss with semi-hard negative pair mining, N-pair loss, angular loss and lifted structure loss, as well as hard sample generation methods, including DAML and HDML. We deploy our proposed method with triplet loss, triplet loss with semi-hard negative pair mining, N-pair loss, angular loss, and lifted structure loss. For fair comparison, we use the same pre-trained CNN model and hyper-parameters as DAML and HDML.

The experimental results on the CUB200, CARS196, and SOP datasets are listed in Tables 1, 2, and 3, respectively. Bold numbers indicate the best score within the same type of loss, and numbers highlighted in gray indicate the best score within the dataset. In comparison with the metric learning losses, combining our proposed method leads to a performance boost with high margins among all baseline losses and datasets in both clustering and retrieval tasks. When triplet loss and lifted structure loss use Euclidean distance with l_2 -normalized features, and N-pair loss and angular loss use cosine similarity with non- l_2 -normalized features during training, the experimental results show that our proposed method is applicable to both cases. Our proposed method outperforms all hard sample generation methods for every loss and dataset except one. In the CARS196 with N-pair loss, HDML (N-pair) shows better performance than Symm + N-pair. We speculate that this is because we use the same hyper-parameters with HDML for the fair comparison without hyper-parameter tuning. On the other hand, the performance improvements of the existing hard sample generation methods with a large training set (i.e., SOP) are relatively smaller than with small training sets (i.e., CUB200 and CARS196), which can be critical for practical usage. While they achieve a 2.0 to 4.6% performance gain in Recall@1 score on the SOP dataset, our proposed method achieve a 5.2 to 11.8% performance boost. This demonstrates that our proposed method gives a competitive performance boost with any training set size.

5 Conclusion

We propose a novel method for generating synthetic hard samples, symmetrical synthesis, for deep metric learning. Applying our method on existing metric learning losses has significantly improved performance by exploiting trivial samples with augmented information and pushing different classes away with stronger power. We demonstrate the effectiveness of symmetrical synthesis with extensive experiments on the three famous benchmarks for image clustering and retrieval tasks.

References

Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G. S.; Davis, A.; Dean, J.; Devin,

- M.; et al. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.
- Ahmed, E.; Jones, M.; and Marks, T. K. 2015. An improved deep learning architecture for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3908–3916.
- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein gan. *arXiv preprint arXiv:1701.07875*.
- Chechik, G.; Sharma, V.; Shalit, U.; and Bengio, S. 2010. Large scale online learning of image similarity through ranking. *Journal of Machine Learning Research* 11(Mar):1109–1135.
- Chopra, S.; Hadsell, R.; LeCun, Y.; et al. 2005. Learning a similarity metric discriminatively, with application to face verification. In *CVPR (1)*, 539–546.
- Cui, Y.; Zhou, F.; Lin, Y.; and Belongie, S. 2016. Fine-grained categorization and dataset bootstrapping using deep metric learning with humans in the loop. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1153–1162.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Duan, Y.; Zheng, W.; Lin, X.; Lu, J.; and Zhou, J. 2018. Deep adversarial metric learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2780–2789.
- Glorot, X., and Bengio, Y. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 249–256.
- Hermans, A.; Beyer, L.; and Leibe, B. 2017. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 554–561.
- Manning, C.; Raghavan, P.; and Schütze, H. 2010. Introduction to information retrieval. *Natural Language Engineering* 16(1):100–103.
- Oh Song, H.; Xiang, Y.; Jegelka, S.; and Savarese, S. 2016. Deep metric learning via lifted structured feature embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4004–4012.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 815–823.
- Sohn, K. 2016. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in Neural Information Processing Systems*, 1857–1865.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1–9.
- Van Der Maaten, L. 2014. Accelerating t-sne using tree-based algorithms. *The Journal of Machine Learning Research* 15(1):3221–3245.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The caltech-ucsd birds-200-2011 dataset.
- Wang, J.; Zhou, F.; Wen, S.; Liu, X.; and Lin, Y. 2017. Deep metric learning with angular loss. In *Proceedings of the IEEE International Conference on Computer Vision*, 2593–2601.
- Weinberger, K. Q., and Saul, L. K. 2009. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research* 10(Feb):207–244.
- Wu, C.-Y.; Manmatha, R.; Smola, A. J.; and Krahenbuhl, P. 2017. Sampling matters in deep embedding learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 2840–2848.
- Yu, R.; Dou, Z.; Bai, S.; Zhang, Z.; Xu, Y.; and Bai, X. 2018. Hard-aware point-to-set deep metric for person re-identification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 188–204.
- Zhao, Y.; Jin, Z.; Qi, G.-j.; Lu, H.; and Hua, X.-s. 2018. An adversarial approach to hard triplet generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 501–517.
- Zheng, W.; Chen, Z.; Lu, J.; and Zhou, J. 2019. Hardness-aware deep metric learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 72–81.