

# Fine-Grained Visual Categorization via Multi-stage Metric Learning

Qi Qian<sup>1</sup> Rong Jin<sup>1</sup> Shenghuo Zhu<sup>2</sup> Yuanqing Lin<sup>3</sup>  
<sup>1</sup>Department of Computer Science and Engineering  
 Michigan State University, East Lansing, MI, 48824, USA  
<sup>2</sup>Alibaba Group, Seattle, WA, 98101, USA  
<sup>3</sup>NEC Laboratories America, Cupertino, CA, 95014, USA

{qianqi, rongjin}@cse.msu.edu, shenghuo.zhu@alibaba-inc.com, ylin@nec-labs.com

## Abstract

*Fine-grained visual categorization (FGVC) is to categorize objects into subordinate classes instead of basic classes. One major challenge in FGVC is the co-occurrence of two issues: 1) many subordinate classes are highly correlated and are difficult to distinguish, and 2) there exists the large intra-class variation (e.g., due to object pose). This paper proposes to explicitly address the above two issues via a distance metric learning (DML). DML addresses the first issue by learning an embedding so that data points from the same class will be pulled together while those from different classes should be pushed apart from each other; and it addresses the second issue by allowing the flexibility that only a portion of the neighbors (not all data points) from the same class need to be pulled together. However, feature representation of an image is often high dimensional, and DML is known to have difficulty in dealing with high dimensional feature vectors since it would require  $\mathcal{O}(d^2)$  for storage and  $\mathcal{O}(d^3)$  for optimization. To this end, we proposed a multi-stage metric learning framework that divides the large-scale high dimensional learning problem to a series of simple subproblems, achieving  $\mathcal{O}(d)$  computational complexity. The empirical study with FGVC benchmark datasets verifies that our method is both effective and efficient compared to the state-of-the-art FGVC approaches.*

## 1. Introduction

Fine-grained visual categorization (FGVC) aims to distinguish objects in subordinate classes. For example, dog images are classified into different breeds of dogs, such as “Chihuahua”, “Pug”, “Samoyed” and so on [17, 24]. One challenge of FGVC is that it has to handle the co-occurrence of two somewhat contradictory requirements: 1) it needs to distinguish many similar classes (e.g., the dog breeds that

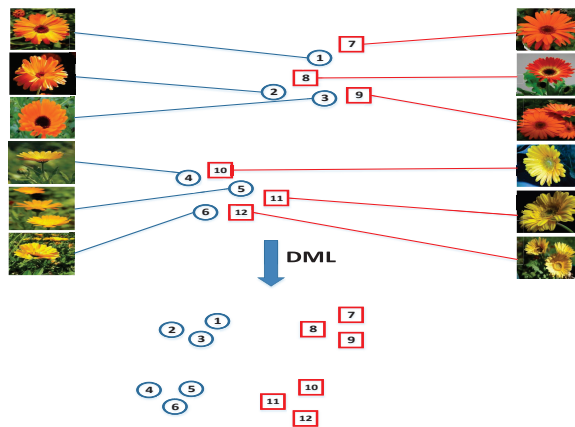


Figure 1. Illustration of how DML learns the embedding that pulls together the data points from the same class and pushes apart the data points from different classes. Blue points are from the class “English marigold” while red ones are “Barberton daisy”. An important note here is that our approach does not require to collapse all data points from each class to a single cluster and this allows the flexibility to model the intra-class variation.

only have subtle differences), and 2) it needs to deal with the large intra-class variation (e.g., caused by different poses, examples, etc.).

The popular pipeline for FGVC consists of two steps, feature extraction step and classification step. The feature extraction step, which sometimes combines with segmentation [1, 7, 24], part localization [2, 35] or both [6], is to extract image level representations, and popular choices include LLC features [1], Fisher vectors [14], etc. A recent development is to train the convolutional neural network (CNN) [18] on a large-scale image dataset (e.g., ImageNet [26]) and then use the trained model to extract features [12]. The so-called deep learning features have demonstrated the state-of-the-art performance on FGVC datasets [12]. Note that there has been some difficulties

in training CNN directly on FGVC datasets because the existing FGVC benchmarks are often too small [12] (only several tens of thousands of training images or less). In this paper, we simply take the state-of-the-art deep learning features without any other operators (e.g., segmentation) and focus on studying better classification approach to address the aforementioned two co-occurring requirements in FGVC.

For the classification step, many existing FGVC methods directly learn a single classifier for each fine-grained class using the one-vs-all strategy [1, 2, 7, 35]. Apparently, this strategy does not scale well to the number of fine-grained classes while the number of subordinate classes in FGVC could be very large (e.g., 200 classes in *birds11* dataset). Additionally, such one-vs-all scheme is only to address the first issue in the two issues, namely, it makes efforts to separate different classes without modeling intra-class variation. In this paper, we propose a distance metric learning (DML) approach, aiming to explicitly handle the two co-occurring requirements with a *single* metric. Fig. 1 illustrates how DML works for FGVC. It learns a distance metric that pulls neighboring data points of the same class close to each other and pushes data points from different classes far apart. By varying the neighborhood size when learning the metric, it is able to effectively handle the tradeoff between the inter-class and intra-class variation. With a learned metric, a  $k$ -nearest neighbor classifier will be applied to find the class assignment for a test image.

Although numerous algorithms have been developed for DML [8, 11, 32, 33], most of them are limited to low dimensional data (i.e. no more than a few hundred dimensions) while the dimensionality of image data representation is usually higher than 10,000 [1]. A straightforward approach toward high dimensional DML is to reduce the dimensionality of data by using the methods such as principle component analysis (PCA) [32] and random projection [29]. The main problem with most dimensionality reduction methods is that they are unable to take into account the supervised information, and as a result, the subspaces identified by the dimensionality reduction methods are usually suboptimal.

There are three challenges in learning a metric directly from the original high dimensional space:

- *Large number of constraints*: A large number of training constraints are usually required to avoid the overfitting of high dimensional DML. The total number of triplet constraints could be up to  $\mathcal{O}(n^3)$  where  $n$  is the number of examples.
- *Computational challenge*: DML has to learn a matrix of size  $d \times d$ , where  $d$  is the dimensionality of data and  $d = 134,016$  in our study. The  $\mathcal{O}(d^2)$  number of variables leads to two computational challenges in finding the optimal metric. First, it results in a slower convergence rate in solving the related optimization

problem [25]. Second, to ensure the learned metric to be positive semi-definitive (PSD), most DML algorithms require, at *every* iteration of optimization, projecting the intermediate solution onto a PSD cone, an expensive operation with complexity of  $\mathcal{O}(d^3)$  (at least  $\mathcal{O}(d^2)$ ).

- *Storage limitation*: It can be expensive to simply save  $\mathcal{O}(d^2)$  number of variables in memory. For example, in our study, it would take more than 130 GB to store the completed metric in memory, which adds more complexity to the already difficult optimization problem.

In this work, we propose a multi-stage metric learning framework for high dimensional DML that explicitly addresses these challenges. First, to deal with a large number of constraints used by high dimensional DML, we divide the original optimization problem into multiple stages. At each stage, only a small subset of constraints that are difficult to be classified by the currently learned metric will be adaptively sampled and used to improve the learned metric. By setting the regularizer appropriately, we can prove that the final solution is optimized over all appeared constraints. Second, to handle the computational challenge in each subproblem, we extend the theory of *dual random projection* [36], which was originally developed for linear classification problems, to DML. The proposed method enjoys the efficiency of random projection, and on the other hand learns a distance metric of size  $d \times d$ . This is in contrast to most dimensionality reduction methods that learn a metric in a *reduced* space. Finally, to handle the storage problem, we propose to maintain a low rank copy of the learned metric by a randomized algorithm for low rank matrix approximation. It not only accelerates the whole learning process but also regularizes the learned metric to avoid overfitting. Extensive comparisons on benchmark FGVC datasets verify the effectiveness and efficiency of the proposed method.

The rest of the paper is organized as follows: Section 2 summarizes related work for DML. Section 3 describes the details of the proposed method. Section 4 shows the results of the empirical study, and Section 5 concludes this work with future directions.

## 2. Related Work

Many algorithms have been developed for DML [11, 32, 33] and a detailed review can be found in two survey papers [19, 34]. Some of them are based on pairwise constraints [11, 33], while others focus on optimizing triplet constraints [8, 32]. In this paper, we adopt triplet constraints, which exactly serve our purpose for addressing the second issue of FGVC. Although numerous studies were devoted to DML, few examined the challenges of high dimensional DML. A common approach for high dimensional DML is to project data into a low dimensional space, and

learn a metric in the space of reduced dimension, which often leads to a suboptimal performance. An alternative approach is to assume  $M$  to be of low rank by writing  $M$  as  $M = LL^\top$  [10, 32], where  $L$  is a tall rectangle matrix and the rank of  $M$  is fixed in advance of applying DML methods. Instead of learning  $M$ , these approaches directly learn  $L$  from data. The main shortcoming of this approach is that it has to solve a non-convex optimization problem, making it computationally less attractive. Several recent studies [20, 25] address high dimensional DML by assuming  $M$  to be sparse. Although resolving the storage problem, they still suffer from high cost in optimizing  $\mathcal{O}(d^2)$  variables.

### 3. Multi-stage Metric Learning

The proposed DML algorithm focuses on triplet constraints so as to pull the small portion of nearest examples from the same class together [32]. Let  $\mathcal{D} = \{(\mathbf{x}_i, y_i), i = 1, \dots, n\}$  be a collection of  $n$  training images, where  $\mathbf{x}_i \in \mathbb{R}^d$  and  $y_i$  is the class assignment of  $\mathbf{x}_i$ . Given a distance metric  $M$ , the distance between two data points  $\mathbf{x}_i$  and  $\mathbf{x}_j$  is measured by

$$d_M(\mathbf{x}_i, \mathbf{x}_j) = (\mathbf{x}_i - \mathbf{x}_j)^\top M (\mathbf{x}_i - \mathbf{x}_j)$$

Let  $\{\mathbf{x}_i^t, \mathbf{x}_j^t, \mathbf{x}_k^t\} (t = 1, \dots, N)$  be a set of  $N$  triplet constraints derived from the training examples in  $\mathcal{D}$ . Since in each constraint  $(\mathbf{x}_i^t, \mathbf{x}_j^t, \mathbf{x}_k^t)$ ,  $\mathbf{x}_i^t$  and  $\mathbf{x}_j^t$  share the same class assignment which is different from that of  $\mathbf{x}_k^t$ , we expect  $d_M(\mathbf{x}_i^t, \mathbf{x}_j^t) < d_M(\mathbf{x}_i^t, \mathbf{x}_k^t)$ . As a result, the optimal distance metric  $M$  is learned by solving the following optimization problem

$$\min_{M \in S_d, M \succeq 0} \frac{\lambda}{2} \|M\|_F^2 + \sum_{t=1}^N \ell(d_M(\mathbf{x}_i^t, \mathbf{x}_k^t) - d_M(\mathbf{x}_i^t, \mathbf{x}_j^t)) \quad (1)$$

where  $S_d$  includes all  $d \times d$  real symmetric matrices and  $\ell(\cdot)$  is a loss function that penalizes the objective function when  $d_M(\mathbf{x}_i^t, \mathbf{x}_k^t)$  is not significantly larger than  $d_M(\mathbf{x}_i^t, \mathbf{x}_j^t)$ . In this study, we choose the smoothed hinge loss [28] that appears to be more effective for optimization than the hinge loss while keeping the benefit of large margin

$$\ell(x) = \begin{cases} 0 & : x > 1 \\ 1 - x - \gamma/2 & : x < 1 - \gamma \\ \frac{1}{2\gamma}(1 - x)^2 & : o.w. \end{cases}$$

One main computational challenge of DML comes from the PSD constraint  $M \succeq 0$  in (1). We address this challenge by following the one projection paradigm [8] that first learns a metric  $M$  without the PSD constraint and then projects  $M$  to the PSD cone at the very end of the learning process. Hence, in this study, we will focus on the following optimization problem for FGVC

$$\min_{M \in S_d} \frac{\lambda}{2} \|M\|_F^2 + \sum_{t=1}^N \ell(\langle A_t, M \rangle) \quad (2)$$

where  $A_t = (\mathbf{x}_i^t - \mathbf{x}_k^t)(\mathbf{x}_i^t - \mathbf{x}_k^t)^\top - (\mathbf{x}_i^t - \mathbf{x}_j^t)(\mathbf{x}_i^t - \mathbf{x}_j^t)^\top$  is introduced as a matrix representation for each triplet constraint, and  $\langle \cdot, \cdot \rangle$  represents the dot product between two matrices.

We will discuss the strategies to address the three challenges of high dimensional DML, and summarize the framework of high dimensional DML for FGVC at the end of this section.

### 3.1. Constraints Challenge: Multi-stage Division

In order to reliably determine the distance metric in a high dimensional space, a large number of training examples are needed to avoid the overfitting problem. Since the number of triplet constraints can be  $\mathcal{O}(n^3)$ , the number of summation terms in (2) can be extremely large, making it difficult to effectively solve the optimization problem in (2). Although learning with active set may help reduce the number of constraints [32], the number of active constraints can still be very large since many images in FGVC from different categories are visually similar, leading to many mistakes. To address this challenge, we divide the learning process into multiple stages. At the  $s$ -th stage, let  $M_{s-1}$  be the distance metric learned from the last stage. We sample a subset of active triplet constraints that are difficult to be classified by  $M_{s-1}$  (i.e., incur large hinge loss)<sup>1</sup>. Given  $M_{s-1}$  and the sampled triplet constraints  $\mathcal{N}_s$ , we update the distance metric by solving the following optimization problem

$$\min_{M \in S_d} \frac{\lambda}{2} \|M - M_{s-1}\|_F^2 + \sum_{t \in \mathcal{N}_s} \ell(\langle A_t, M \rangle) \quad (3)$$

Although only a small size of constraints is used to improve the metric at each stage, we have

**Theorem 1.** *The metric learned by solving the problem (3) also optimizes the following objective function*

$$\min_{M \in S_d} \frac{\lambda}{2} \|M\|_F^2 + \sum_{k=1}^s \sum_{t \in \mathcal{N}_k} \ell(\langle A_t, M \rangle)$$

*Proof.* Consider the objective function for the first  $s$  stages

$$\min_{M \in S_d} \frac{\lambda}{2} \|M\|_F^2 + \underbrace{\sum_{k=1}^{s-1} \sum_{t \in \mathcal{N}_k} \ell(\langle A_t, M \rangle)}_{:= \mathcal{L}_{s-1}(M)} + \sum_{t \in \mathcal{N}_s} \ell(\langle A_t, M \rangle) \quad (4)$$

It is obvious that  $\mathcal{L}_{s-1}$  is strongly convex, so we have (Chapter 9, [4])

$$\begin{aligned} \mathcal{L}_{s-1}(M) &= \mathcal{L}_{s-1}(M_{s-1}) + \langle \nabla \mathcal{L}_{s-1}(M_{s-1}), M - M_{s-1} \rangle \\ &\quad + \frac{1}{2} \langle (M - M_{s-1}) \nabla^2 \mathcal{L}_{s-1}(M'), M - M_{s-1} \rangle \end{aligned}$$

<sup>1</sup>The strategy of finding hard constraints at each stage is also applied by cutting plane methods [21] and active learning [27].

for some  $M'$  between  $M$  and  $M_{s-1}$ .

Since  $M_{s-1}$ , the solution obtained from the first  $s - 1$  stages, approximately optimizes  $\mathcal{L}_{s-1}(M)$  and  $\mathcal{L}_{s-1}$  is  $\lambda$ -strongly convex, then

$$\mathcal{L}_{s-1}(M) \approx \mathcal{L}_{s-1}(M_{s-1}) + \frac{\lambda}{2} \|M - M_{s-1}\|_F^2 \quad (5)$$

We finish the proof by replacing  $\mathcal{L}_{s-1}(M)$  in (4) with the approximation in (5).  $\square$

**Remark** This theorem demonstrates that the metric learned from the last stage is optimized over constraints from all stages. Therefore, the original problem could be divided into several subproblems and each of them has an affordable number of active constraints. Fig. 2 summaries the framework of the multi-stage learning procedure.




Figure 2. The framework of the proposed method.

### 3.2. Computational Challenge: Dual Random Projection

Now we try to solve the high dimensional subproblem by dual random projection technique. To simplify the analysis, we investigate the subproblem at the first stage and the following stages could be analyzed in the same way. By introducing the convex conjugate  $\ell_*$  for  $\ell$  in (3), the dual problem of DML is

$$\max_{\alpha \in \mathbb{R}^{|\mathcal{N}_1|}} - \sum_{t=1}^{|\mathcal{N}_1|} \ell_*(\alpha_t) - \frac{1}{2\lambda} \alpha^\top G \alpha \quad (6)$$

where  $\alpha_t$  is the dual variable for  $A_t$  and  $G$  is a matrix defined as  $G_{a,b} = \langle A_a, A_b \rangle$ .  $M_1 = -\frac{1}{\lambda} \sum_{t=1}^{|\mathcal{N}_1|} \alpha_t A_t$  by setting the gradient with respect to  $M_1$  to zero. Let  $R_1, R_2 \in \mathbb{R}^{d \times m}$  be two Gaussian random matrices, where  $m$  is the number of random projections ( $m \ll d$ ) and  $R_1^{i,j}, R_2^{i,j} \sim \mathcal{N}(0, 1/m)$ . For each triplet constraint, we project its representation  $A_t$  into the low dimensional space using the random matrices, i.e.  $\hat{A}_t = R_1^\top A_t R_2$ . By using double random projections, which is different from the single random projection in [36], we have

**Lemma 1.**  $\forall A_a, A_b$ , the double random projections preserve the pairwise similarity between them:  $\mathbb{E}[\langle \hat{A}_a, \hat{A}_b \rangle] = \langle A_a, A_b \rangle$

The proof is straightforward. According to the lemma, the dual variables in (6) can be estimated in the low dimensional space as

$$\max_{\hat{\alpha} \in \mathbb{R}^{|\mathcal{N}_1|}} - \sum_{t=1}^{|\mathcal{N}_1|} \ell_*(\hat{\alpha}_t) - \frac{1}{2\lambda} \hat{\alpha}^\top \hat{G} \hat{\alpha} \quad (7)$$

where  $\hat{G}(a, b) = \langle \hat{A}_a, \hat{A}_b \rangle$ . Then, by the definition of convex conjugate, each dual variable  $\hat{\alpha}_t$  in (7) can be further estimated by  $\ell'(\langle \hat{A}_t, \hat{M}_1 \rangle)$ , where  $\hat{M}_1 \in \mathbb{R}^{m \times m}$  is the metric learned in the reduced space. Generally,  $\hat{M}_s$  is learned by solving the following optimization problem

$$\min_{\hat{M}_s \in S_m} \frac{\lambda}{2} \|\hat{M}_s - \hat{M}_{s-1}\|_F^2 + \sum_{t=1}^{|\mathcal{N}_s|} \ell(\langle \hat{A}_t, \hat{M}_s \rangle) \quad (8)$$

Since the size of  $\hat{M}_s \in \mathbb{R}^{m \times m}$  is significantly smaller than that of  $M_s$ , (8) can be solved much more efficiently than (3). In our implementation, a simple stochastic gradient descent (SGD) method is developed to efficiently solve the optimization problem in (8). Given  $\hat{M}_1$ , the final distance metric  $M_1 \in \mathbb{R}^{d \times d}$  in the original space is estimated as

$$M_1 = -\frac{1}{\lambda} \sum_{t=1}^{|\mathcal{N}_1|} \hat{\alpha}_t A_t \quad (9)$$

### 3.3. Storage Challenge: Low Rank Approximation

Although (9) allows us to recover the distance metric  $M$  in original  $d$  dimensional space from the dual variables  $\{\alpha_t\}_{t=1}^{|\mathcal{N}_1|}$ , it is expensive, if not impossible, to save  $M$  in memory since  $d$  is very large in FGVC [1]. To address this challenge, instead of saving  $M$ , we propose to save the low rank approximation of  $M$ . More specifically, let  $\sigma_1, \dots, \sigma_r$  be the first  $r \ll d$  eigenvalues of  $M$ , and let  $\mathbf{u}_1, \dots, \mathbf{u}_r$  be the corresponding eigenvectors. We approximate  $M$  by a low rank matrix  $M' = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{u}_i^\top = LL^\top$ . Different from existing DML methods that directly optimize  $L$  [31], we obtain  $M$  first and then decompose it to avoid suboptimal solution. Unlike  $M$  that requires  $\mathcal{O}(d^2)$  storage space, it only takes  $\mathcal{O}(rd)$  space to save  $M'$  and  $r$  could be an arbitrary value. In addition, the low rank metric accelerates the sampling step by reducing the cost of computing distance from  $\mathcal{O}(d)$  to  $\mathcal{O}(r)$ . Low rank is also a popular regularizer to avoid overfitting when learning high dimensional metric [20]. However, the key issue is how to efficiently compute the eigenvectors and eigenvalues of  $M$  at each stage. This is particularly challenging in our case as  $M$  in (9) even can not be computed explicitly due to its large size.

To address this problem, first we investigate the structure

of the recovering step for the  $s$ -th stage as in (9)

$$\begin{aligned}
M_s &= M_{s-1} - \frac{1}{\lambda} \sum_{t=1}^{|\mathcal{N}_s|} \alpha_t^s A_t^s \\
&= M_{s-2} - \frac{1}{\lambda} \left( \sum_{t=1}^{|\mathcal{N}_s|} \alpha_t^s A_t^s + \sum_{t=1}^{|\mathcal{N}_{s-1}|} \alpha_t^{s-1} A_t^{s-1} \right) \\
&= -\frac{1}{\lambda} \sum_{k=1}^s \sum_{t=1}^{|\mathcal{N}_k|} \alpha_t^k A_t^k
\end{aligned}$$

Therefore, we can express the summation as matrix multiplications. In particular, for each triplet  $(\mathbf{x}_i^t, \mathbf{x}_j^t, \mathbf{x}_k^t)$ , we denote its dual variable by  $\alpha = \ell'(\langle \widehat{A}, \widehat{M} \rangle)$  and set the corresponding entries in a sparse matrix  $C$  as

$$\begin{aligned}
C(i, j) &= \frac{\alpha}{\lambda}, \quad C(j, i) = \frac{\alpha}{\lambda}, \quad C(j, j) = -\frac{\alpha}{\lambda} \\
C(i, k) &= -\frac{\alpha}{\lambda}, \quad C(k, i) = -\frac{\alpha}{\lambda}, \quad C(k, k) = \frac{\alpha}{\lambda}
\end{aligned} \quad (10)$$

It is easy to verify that  $M$  can be written as

$$M = X C X^\top \quad (11)$$

Second, we exploit the randomized theory [15] to efficiently compute the eigen-decomposition of  $M$ . More specifically, let  $R \in \mathbb{R}^{d \times q}$  ( $q \ll d$ ) be an Gaussian random matrix. According to [15], with an overwhelming probability, most of the top  $r$  eigenvectors of  $M$  lie in the subspace spanned by the column vectors in  $MR$  provided  $q \geq r + k$ , where  $k$  is a constant independent from  $d$ . The limitation of the method is that it requires the appearance of the matrix  $M$  for computing  $MR$  while keeping the whole matrix is unaffordable here. Fortunately, by replacing  $M$  with  $X C X^\top$  according to (11), we can approximate the top eigenvectors of  $M$  by those of  $X C X^\top R$  that is of size  $d \times q$  and can be computed efficiently since  $C$  is a sparse matrix. The overall computational cost of the proposed algorithm for low rank approximation is only  $\mathcal{O}(qnd)$ , which is linear in  $d$ . Note that the sparse matrix  $C$  is cumulated over all stages.

Alg. 1 summarizes the key steps of the proposed approach for low rank approximation, where  $qr$  and  $eig$  stand for QR and eigen decomposition of a matrix. Note that the distributed computing is particularly effective for the realization of the algorithm because the matrix multiplications  $X C X^\top R$  can be accomplished in parallel, which is helpful when  $n$  is also large.

Alg. 2 shows the whole picture of the proposed method.

## 4. Experiments

DeCAF features [12] are extracted as the image representations in the experiments. Although it is from the activation of a deep convolutional network, which is trained on

---

### Algorithm 1 An Efficient Algorithm for Recovering $M$ and Projecting It onto PSD Cone from $\widehat{M}$

---

- 1: **Input:** Dataset  $X \in \mathbb{R}^{d \times n}$ ,  $\widehat{M} \in \mathbb{R}^{m \times m}$ , the number of random combinations  $q$
  - 2: Compute a Gaussian random matrix  $R \in \mathbb{R}^{d \times q}$
  - 3: Compute the sparse matrix  $C$  using (10)
  - 4:  $Y = R \times X^\top$ ,  $Y = Y \times C$ ,  $Y = Y \times X$
  - 5:  $[Q, R] = qr(Y)$
  - 6:  $B = Q^\top \times X^\top$ ,  $B = B \times C$ ,  $B = B \times X$
  - 7:  $[U, \Sigma] = eig(B)$
  - 8:  $U = Q * U$
  - 9: **return**  $L = [\sqrt{\sigma_1} \mathbf{u}_1, \dots, \sqrt{\sigma_r} \mathbf{u}_r]$  and  $M = LL^\top$ , where  $\mathbf{u}_i$  is the  $i$ th column of  $U$  and  $\sigma_i$  is the  $i$ th positive diagonal element of  $\Sigma$
- 

---

### Algorithm 2 The Multi-stage Metric Learning Framework for High Dimensional DML (MsML)

---

- 1: **Input:** Dataset  $X \in \mathbb{R}^{d \times n}$ , the number of random projections  $m$ , the number of random combinations  $q$ , and the number of stages  $T$
  - 2: Compute two Gaussian random matrices  $R_1, R_2 \in \mathbb{R}^{d \times m}$
  - 3: Initialize  $\widehat{M}_0 = \mathbf{0} \in \mathbb{R}^{m \times m}$  and  $M_0 = \mathbf{0} \in \mathbb{R}^{d \times d}$
  - 4: **for**  $s = 1, \dots, T$  **do**
  - 5:   Sample one epoch active triplet constraints using  $M_{s-1}$
  - 6:   Estimate  $\widehat{M}_s$  by solving the optimization problem as in (8) with SGD
  - 7:   Recover the distance metric  $M_s$  in the  $d$  dimensional space using Alg. 1
  - 8: **end for**
  - 9: **return**  $M_T$
- 

ImageNet [18], it outperforms conventional visual features on many general tasks [12]. We concatenate features from the last three fully connected layers (i.e., DeCAF<sub>5+6+7</sub>) and the dimension of resulting features is 51, 456.

We apply the proposed algorithm to learn a distance metric and use the learned metric together with a smoothed  $k$ -nearest neighbor classifier, a variant of  $k$ -NN, to predict the class assignments for test examples. Different from conventional  $k$ -NN, it first obtains  $k$  reference centers for each class by clustering training images in each class into  $k$  clusters. Then, it computes the query's distance to each class as the soft min of the distances between the test image and corresponding reference centers, and assigns the test image to the class with the shortest distance. It is more efficient when predicting, especially for large-scale training set, and the performance is similar to that of conventional one. We refer to the classification approach based on the metric learned by the proposed algorithm and the smoothed  $k$ -NN as **MsML**,

and the smoothed  $k$ -NN with Euclidean distance in the original space as **Euclid**. Although the size of the covariance matrix is very large ( $51,456 \times 51,456$ ), its rank is low due to the small number of training examples, and thus PCA can be computed explicitly. The state-of-the-art DML algorithm, i.e. **LMNN** [32] with PCA as preprocess, is also included in comparison. The one-vs-all strategy, based on the implementation of **LIBLINEAR** [13], is used as a baseline for FGVC, with the regularization parameter varied in the range  $\{10^i\}(i = -2, \dots, 3)$ . We refer to it as **LSVM**. We also include the state-of-the-art results for FGVC in our evaluation. All the parameters used by MsML are set empirically, with the number of random projections  $m = 100$  and the number of random combinations  $q = 600$ . PCA is applied for LMNN to reduce the dimensionality to  $m$  before the metric is learned. LMNN is implemented by the code from the original authors and the recommended parameters are used<sup>2</sup>. To ensure that the baseline method fully exploits the training data, we set the maximum number of iterations for LMNN as  $10^4$ . These parameter values are used throughout all the experiments. All training/test splits are provided by datasets. **Mean accuracy**, a standard evaluation metric for FGVC, is used to evaluate the classification performance. All experiments are run on a single machine with 16 2.10GHz cores and 96GB memory.

#### 4.1. Oxford Cats&Dogs

*cats&dogs* contains 7,349 images from 37 cat and dog species [24]. There are about 100 images per class for training and the rest are for test. Table 1 summaries the results. First, we observe that MsML is more accurate than the baseline LSVM. This is not surprising because the distance metric is learned from the training examples of all class assignments. This is in contrast to the one-vs-all approach used in LSVM that the classification function for a class  $C$  is learned only by the examples with the class assignment of  $C$ . Second, our method performs significantly better than the baseline DML method, indicating that the unsupervised dimension reduction method PCA may result in suboptimal solutions for DML. Fig. 3 compares the images that are most similar to the query images using the metric learned by the proposed algorithm (Column 8-10) to those based on the metric learned by LMNN (Column 5-7) and Euclid (Column 2-4). We observe that more images from the same class as the query are found by the metric learned by MsML than LMNN. For example, MsML is able to capture the difference between two cat species (longhair v.s. shorthair) while LMNN returns the very similar images with wrong class assignments. Third, MsML has overwhelming performance compared to all state-of-the-art FGVC approaches. Although the method [24] using ground truth head bound-

<sup>2</sup>We did vary the parameter slightly from the recommended values and did not find any noticeable change in the classification accuracy.

ing box and segmentation achieves 59.21%, MsML is 20% better than it with only image information, which shows the advantage of the proposed method. Finally, it takes less than 0.2 second to extract DeCAF features per image based on a *CPU* implementation while a simple segmentation operator costs more than 2.5 seconds as reported in the study [1], making the proposed method for FGVC more appealing.

Table 1. Comparison of mean accuracy(%) on *cats&dogs* dataset. “#” means that more information (e.g., ground true segmentation) is used by the method.

Methods	Mean Accuracy (%)
Image only [24]	39.64
Det+Seg [1]	54.30
Image+Head+Body# [24]	59.21
Euclid	72.60
LSVM	77.63
LMNN	76.24
MsML	80.45
MsML+	81.18

To evaluate the performance of MsML for extremely high dimensional features, we concatenate conventional features by using the pipeline for visual feature extraction that is outlined in [1]. Specifically, we extract HOG [9] features at 4 different scales and encode them to  $8K$  dimensional feature dictionary by the LLC method [31]. A max pooling strategy is then used to aggregate local features into a single vector representation. Finally, 82,560 features are extracted from each image and the total dimension is up to 134,016. MsML with the combined features is denoted as **MsML+** and it further improves the performance by about 1% as in Table 1. Note that the time of extracting these high dimensional conventional features is only 0.5 second per image, which is still much cheaper than any segmentation or localization operator.

#### 4.2. Oxford 102 Flowers

*102flowers* is the Oxford flowers dataset for flower species [23], which consists of 8189 images from 102 classes. Each class has 20 images for training and rest for test. Table 2 shows the results from different methods. We have the similar conclusion for the baseline methods. That is, MsML outperforms LSVM and LMNN significantly. Although LSVM already performs very well, MsML further improves the accuracy. Additionally, it is observed that even the performances of state-of-the-art methods with segmentation operators are much worse than that of MsML. Note that GT [23] uses hand annotated segmentations followed by multiple kernel SVM, while MsML outperforms it about 3% without any supervised information, which confirms the effectiveness of the proposed method.

Fig. 4 illustrates the changing trend of test mean accu-

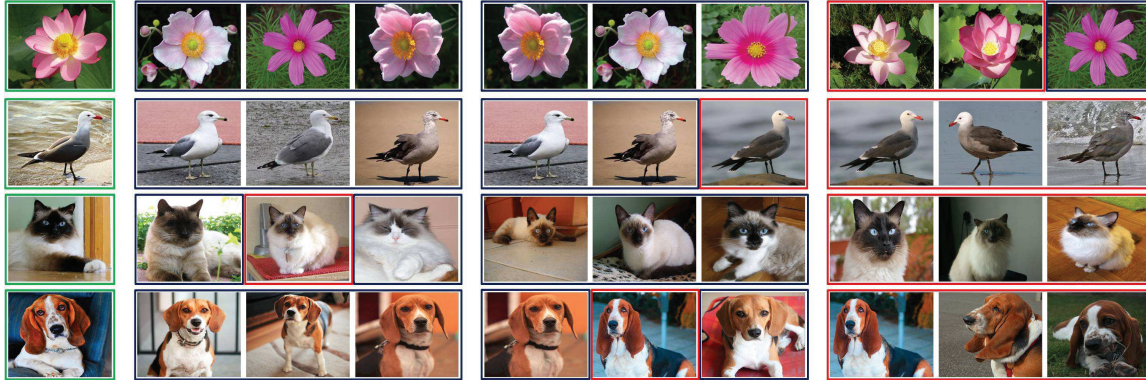


Figure 3. Examples of retrieved images. The first column indicates the query images highlighted by green bounding boxes. Columns 2-4 include the most similar images measured by Euclid. Columns 5-7 show those by the metric from LMNN. Columns 8-10 are from the metric of MsML. Images in columns 2-10 are highlighted by red bounding boxes when they share the same category as queries, and blue bounding boxes if they are not.

Table 2. Comparison of mean accuracy(%) on *102flowers* dataset. “#” means that more information (e.g., ground true segmentation) is used by the method.

Methods	Mean Accuracy (%)
Combined CoHoG [16]	74.80
Combined Features [22]	76.30
BiCoS-MT [5]	80.00
Det+Seg [1]	80.66
TriCoS [7]	85.20
GT# [23]	85.60
Euclid	76.21
LSVM	87.14
LMNN	81.93
MsML	88.39
MsML+	89.45

accuracy as the number of stages increases. We observe that MsML converges very fast, which verifies that multi-stage division is essential to the proposed framework.




Figure 4. Convergence curve Figure 5. Comparison with of the proposed method on different size of classes on *102flowers*. *birds11*.

### 4.3. Birds-2011

*birds11* is the Caltech-USCD-200-2011 birds dataset for bird species [30]. There are 200 classes with 11,788 images and each class has roughly 30 images for training. We use the version with ground truth bounding box. Table 3 compares the proposed method to the state-of-the-art baselines.

Table 3. Comparison of mean accuracy(%) on *birds11* dataset. “\*” denotes the method that mirrors training images.

Methods	Mean Accuracy (%)
Symb [6]	56.60
POOF [3]	56.78
Symb* [6]	59.40
Ali* [14]	62.70
DeCAF+DPD [12]	64.96
Euclid	46.85
LSVM	61.44
LMNN	51.04
MsML	65.84
MsML+	66.61
MsML+*	67.86

First, it is obvious that the performance of MsML is significantly better than all baseline methods as the observation above. Second, although Symb [6] combines segmentation and localization, MsML outperforms it by 9% without any time consuming operator. Third, Symb\* and Ali\* mirror the training images to improve their performances, while MsML is even better than them without this trick. Finally, MsML outperforms the method combining DeCAF features and DPD models [37], which is due to the fact that most of studies for FGVC ignore choosing the appropriate base classifier and simply adopt linear SVM with the one-vs-all strategy. For comparison, we also report the result mirroring training images which is denoted as **MsML+\***. It provides another 1% improvement over MsML+ as shown in Table 3.

To illustrate the capacity of MsML in exploring the correlation among classes, which makes it more effective than a simple one-vs-all classifier for FGVC, we conduct one additional experiment. We randomly select 50 classes from *birds11* as the *target classes* and use the test images from the target classes for evaluation. When learning the met-

ric, besides the training images from 50 target classes, we sample  $k$  classes from 150 unselected ones as the *auxiliary* classes, and use training images from the auxiliary classes as additional training examples for DML. Fig. 5 compares the performance of LSVM and MsML with the increasing number of auxiliary classes. It is not surprising to observe that the performance of LSVM decreases a little since it is unable to explore the supervision information in the auxiliary classes to improve the classification accuracy of target classes and more auxiliary classes just intensify the class imbalance problem. In contrast, the performance of MsML improves significantly with increasing auxiliary classes, indicating that MsML is capable of effectively exploring the training data from the auxiliary classes and therefore is particularly suitable for FGVC.

#### 4.4. Stanford Dogs

*S-dogs* is the Stanford dog species dataset [17]. It contains 120 classes and 20,580 images, where 100 images from each class is used for training. Since it is the subset of ImageNet [26], where DeCAF model is trained from, we just report the result in Table 4 as reference.

Table 4. Comparison of mean accuracy(%) on *S-dogs* dataset. “\*” denotes the method that mirrors training images.

Methods	Mean Accuracy (%)
SIFT [17]	22.00
Edge Templates [35]	38.00
Symb [6]	44.10
Symb* [6]	45.60
Ali* [14]	50.10
Euclid	59.22
LSVM	65.00
LMNN	62.17
MsML	69.07
MsML+	69.80
MsML+*	70.31

#### 4.5. Comparison of Efficiency

In this section, we compare the training time of the proposed algorithm for high dimensional DML to that of LSVM and LMNN. MsML is implemented by Julia, which is a little slower than C<sup>3</sup>, while LSVM uses the LIBLINEAR package, the state-of-the-art algorithm for solving linear SVM implemented mostly in C. The core part of LMNN is also implemented in C. The time for feature extraction is not included here because it is shared by all the methods in comparison. The running time for MsML includes all operational cost (i.e., the cost for sampling triplet constraints, computing random projections and low rank approximation).

<sup>3</sup>Detailed comparison can be found in <http://julialang.org>

Table 5. Comparison of running time (seconds).

Methods	<i>cats&amp;dogs</i>	<i>102flowers</i>	<i>birds11</i>	<i>S-dogs</i>
LSVM	196.2	309.8	1,417.0	1,724.8
LMNN	832.6	702.7	1,178.2	1,643.6
MsML	164.9	174.4	413.1	686.3
MsML+	337.2	383.7	791.3	1,229.7

Table 5 summarizes the results of the comparison. First, it takes MsML about 1/3 of the time to complete the computation compared to LMNN. This is because MsML employs a stochastic optimization method to find the optimal distance metric while LMNN is a batch learning method. Second, we observe that the proposed method is significantly more efficient than LSVM on most of datasets. The high computational cost of LSVM mostly comes from two aspects. First, LSVM has to train one classification model for each class, and becomes significantly slower when the number of classes is large. Second, the fact that images from different classes are visually similar makes it computationally difficult to find the optimal linear classifier that can separate images of one class from images from the other classes. In contrast, the training time of MsML is independently from the number of classes, making it more appropriate for FGVC. Finally, the running time of MsML+ with 134,016 features only doubles that of MsML, which verifies that the proposed method is linear in dimensionality ( $\mathcal{O}(d)$ ).

## 5. Conclusion

In this paper, we propose a multi-stage metric learning framework for high dimensional FGVC problem, which addresses the challenges arising from high dimensional DML. More specifically, it divides the original problem into multiple stages to handle the challenge arising from too many triplet constraints, extends the theory of dual random projection to address the computational challenge for high dimensional data, and develops a randomized low rank matrix approximation algorithm for the storage challenge. The empirical study shows that the proposed method with general purpose features yields the performance that is significantly better than the state-of-the-art approaches for FGVC. In the future, we plan to combine the proposed DML algorithm with segmentation and localization to further improve the performance of FGVC. Additionally, since the proposed method is a general DML approach, we will try to apply it for other applications with high dimensional features.

**Acknowledgments** Qi Qian and Rong Jin are supported in part by ARO (W911NF-11-1-0383), NSF (IIS-1251031) and ONR (N000141410631).



## References

- [1] A. Angelova and S. Zhu. Efficient object detection and segmentation for fine-grained recognition. In *CVPR*, 2013.
- [2] T. Berg and P. N. Belhumeur. Poof: Part-based one-vs-one features for fine-grained categorization, face verification, and attribute estimation. In *CVPR*, 2013.
- [3] T. Berg and P. N. Belhumeur. POOF: part-based one-vs.-one features for fine-grained categorization, face verification, and attribute estimation. In *CVPR*, pages 955–962, 2013.
- [4] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2009.
- [5] Y. Chai, V. S. Lempitsky, and A. Zisserman. Bicos: A bi-level co-segmentation method for image classification. In *ICCV*, pages 2579–2586, 2011.
- [6] Y. Chai, V. S. Lempitsky, and A. Zisserman. Symbiotic segmentation and part localization for fine-grained categorization. In *ICCV*, pages 321–328, 2013.
- [7] Y. Chai, E. Rahtu, V. S. Lempitsky, L. J. V. Gool, and A. Zisserman. Tricos: A tri-level class-discriminative co-segmentation method for image classification. In *ECCV*, pages 794–807, 2012.
- [8] G. Chechik, V. Sharma, U. Shalit, and S. Bengio. Large scale online learning of image similarity through ranking. *JMLR*, 11:1109–1135, 2010.
- [9] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, pages 886–893, 2005.
- [10] J. V. Davis and I. S. Dhillon. Structured metric learning for high dimensional problems. In *KDD*, pages 195–203, 2008.
- [11] J. V. Davis, B. Kulis, P. Jain, S. Sra, and I. S. Dhillon. Information-theoretic metric learning. In *ICML*, pages 209–216, 2007.
- [12] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *ICML*, pages 647–655, 2014.
- [13] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *JMLR*, 9:1871–1874, 2008.
- [14] E. Gavves, B. Fernando, C. G. M. Snoek, A. W. M. Smeulders, and T. Tuytelaars. Fine-grained categorization by alignments. In *ICCV*, pages 1713–1720, 2013.
- [15] N. Halko, P.-G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *ArXiv e-prints*, Sept. 2009.
- [16] S. Ito and S. Kubota. Object classification using heterogeneous co-occurrence features. In *ECCV*, pages 209–222, 2010.
- [17] A. Khosla, N. Jayadevaprakash, B. Yao, and F.-f. Li. Novel dataset for fine-grained image categorization. In *First Workshop on Fine-Grained Visual Categorization, CVPR*, 2011.
- [18] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1106–1114, 2012.
- [19] B. Kulis. Metric learning: A survey. *Foundations and Trends in Machine Learning*, 5(4):287–364, 2013.
- [20] D. K. H. Lim, B. McFee, and G. Lanckriet. Robust structural metric learning. In *ICML*, 2013.
- [21] Y. Nesterov. *Introductory lectures on convex optimization*, volume 87. Springer Science & Business Media, 2004.
- [22] M.-E. Nilsback. *An Automatic Visual Flora – Segmentation and Classification of Flowers Images*. PhD thesis, University of Oxford, 2009.
- [23] M.-E. Nilsback and A. Zisserman. Automated flower classification over a large number of classes. In *ICVGIP*, pages 722–729, 2008.
- [24] O. M. Parkhi, A. Vedaldi, A. Zisserman, and C. V. Jawahar. Cats and dogs. In *CVPR*, pages 3498–3505, 2012.
- [25] G.-J. Qi, J. Tang, Z.-J. Zha, T.-S. Chua, and H.-J. Zhang. An efficient sparse metric learning in high-dimensional space via  $l_1$ -penalized log-determinant regularization. In *ICML*, page 106, 2009.
- [26] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge, 2014.
- [27] B. Settles. Active learning literature survey. *University of Wisconsin, Madison*, 52(55-66):11, 2010.
- [28] S. Shalev-Shwartz and T. Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *CoRR*, abs/1209.1873, 2012.
- [29] G. Tsagkatakis and A. E. Savakis. Manifold modeling with learned distance in random projection space for face recognition. In *ICPR*, pages 653–656, 2010.
- [30] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical report, 2011.
- [31] J. Wang, J. Yang, K. Yu, F. Lv, T. S. Huang, and Y. Gong. Locality-constrained linear coding for image classification. In *CVPR*, pages 3360–3367, 2010.
- [32] K. Q. Weinberger and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. *JMLR*, 10:207–244, 2009.
- [33] E. P. Xing, A. Y. Ng, M. I. Jordan, and S. J. Russell. Distance metric learning with application to clustering with side-information. In *NIPS*, pages 505–512, 2002.
- [34] L. Yang and R. Jin. Distance metric learning: a comprehensive survey. 2006.
- [35] S. Yang, L. Bo, J. Wang, and L. G. Shapiro. Unsupervised template learning for fine-grained object recognition. In *NIPS*, pages 3131–3139, 2012.
- [36] L. Zhang, M. Mahdavi, R. Jin, T.-B. Yang, and S. Zhu. Recovering optimal solution by dual random projection. In *arXiv:1211.3046*, 2013.
- [37] N. Zhang, R. Farrell, F. N. Iandola, and T. Darrell. Deformable part descriptors for fine-grained recognition and attribute prediction. In *ICCV*, pages 729–736, 2013.