
Geometric Mean Metric Learning

Pourya Habib Zadeh

Reshad Hosseini

School of ECE, College of Engineering, University of Tehran, Tehran, Iran

Suvrit Sra

Massachusetts Institute of Technology, Cambridge, MA, USA

P.HABIBZADEH@UT.AC.IR

RESHAD.HOSSEINI@UT.AC.IR

SUVRIT@MIT.EDU

Abstract

We revisit the task of learning a Euclidean metric from data. We approach this problem from first principles and formulate it as a surprisingly simple optimization problem. Indeed, our formulation even admits a closed form solution. This solution possesses several very attractive properties: (i) an innate geometric appeal through the Riemannian geometry of positive definite matrices; (ii) ease of interpretability; and (iii) computational speed several orders of magnitude faster than the widely used LMNN and ITML methods. Furthermore, on standard benchmark datasets, our closed-form solution consistently attains higher classification accuracy.

1. Introduction

Many machine learning algorithms require computing distances between input data points, be it for clustering, classification, or search. Selecting the distance measure is, therefore, an important concern; though the answer is task specific. When supervised or weakly supervised information is available, selection of the distance function can itself be cast as a learning problem called “metric learning” (Kulis, 2012; Weinberger & Saul, 2009).

In its most common form, metric learning seeks to learn a Euclidean metric. An abstract approach is to take input data in \mathbb{R}^n and learn a linear map $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^m$, so that the Euclidean distance $\|\Phi(\mathbf{x}) - \Phi(\mathbf{y})\|$ can be used to measure the distance between points $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. More generally, the map Φ can also be nonlinear.

The problem of learning linear maps was introduced in (Xing et al., 2002) as “Mahalanobis metric learning.”

Since then metric learning has witnessed a sequence of improvements both in modeling and algorithms (see related work). More broadly, the idea of linearly transforming input features is a bigger theme across machine learning and statistics; encompassing whitening transforms, linear dimensionality reduction, Euclidean metric learning, and more (Kulis, 2012; Cunningham & Ghahramani, 2015).

We revisit the task of learning a Euclidean metric. Like most Euclidean metric learning methods, we also seek to learn a *Mahalanobis distance*¹

$$d_{\mathbf{A}}(\mathbf{x}, \mathbf{x}') = (\mathbf{x} - \mathbf{x}')^T \mathbf{A}(\mathbf{x} - \mathbf{x}'), \quad (1)$$

where $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$ are input vectors, and \mathbf{A} is a $d \times d$ real, symmetric positive definite (SPD) matrix². Like other metric learning approaches we also assume weak-supervision, which is provided through the sets of pairs

$$\begin{aligned} \mathcal{S} &:= \{(\mathbf{x}_i, \mathbf{x}_j) \mid \mathbf{x}_i \text{ and } \mathbf{x}_j \text{ are in the same class}\} \\ \mathcal{D} &:= \{(\mathbf{x}_i, \mathbf{x}_j) \mid \mathbf{x}_i \text{ and } \mathbf{x}_j \text{ are in different classes}\}. \end{aligned}$$

Unlike other Euclidean metric learning methods, however, we follow a much simpler yet fresh new approach.

Specifically, we make the following main contributions:

- **Formulation.** We formulate Euclidean metric learning from first principles following intuitive geometric reasoning; we name our setup “Geometric Mean Metric Learning” (GMML) and cast it as an *unconstrained* smooth, strictly convex optimization problem.
- **Solution & insights.** We show that our formulation admits a closed form solution, which not only also enjoys connections to the Riemannian geometry of SPD matrices (and thus explains the name GMML) but also has important empirical consequences.

Proceedings of the 33rd International Conference on Machine Learning, New York, NY, USA, 2016. JMLR: W&CP volume 48. Copyright 2016 by the author(s).

¹This is actually a squared distance. The true metric is $\sqrt{d_{\mathbf{A}}}$; but in accord with metric learning literature we call (1) a distance.

²Do not confuse SPD with positive semi-definite matrices.

– **Validation.** We consider multi-class classification using the learned metrics, and validate GMML by comparing it against widely used metric learning methods. GMML runs up to three orders of magnitude faster while consistently delivering equal or higher classification accuracy.

1.1. Related work

We recall below some related work to help place GMML in perspective. We omit a discussion of nonlinear methods, and other variations of the basic Euclidean task outlined above; for these, we refer the reader to both kernelized metric learning (Davis et al., 2007) and other techniques as summarized in the recent surveys of Kulis (2012) and Bellet et al. (2013).

Probably the earliest work to formulate metric learning is (Xing et al., 2002), sometimes referred to as MMC. This method minimizes the sum of distances over similar points while trying to ensure that dissimilar points are far away from each other. Using the sets \mathcal{S} and \mathcal{D} , MMC solves the optimization problem

$$\begin{aligned} \min_{\mathbf{A} \succeq 0} \quad & \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{S}} d_{\mathbf{A}}(\mathbf{x}_i, \mathbf{x}_j) \\ \text{such that} \quad & \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{D}} \sqrt{d_{\mathbf{A}}(\mathbf{x}_i, \mathbf{x}_j)} \geq 1. \end{aligned} \quad (2)$$

Xing et al. (2002) use $\sqrt{d_{\mathbf{A}}}$ instead of the distance $d_{\mathbf{A}}$ because under $d_{\mathbf{A}}$, problem (2) has a trivial rank-one solution. To optimize (2) Xing et al. (2002) use a gradient-descent algorithm combined with a projection onto the set of positive semi-definite matrices. The term $\sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{S}} d_{\mathbf{A}}(\mathbf{x}_i, \mathbf{x}_j)$ is also used in the other metric learning methods like LMNN (Weinberger & Saul, 2009) and MCML (Globerson & Roweis, 2005) as a part of their cost functions.

Information-Theoretic Metric Learning (ITML) (Davis et al., 2007), aims to satisfy the similarity and dissimilarity constraints while staying as “close” as possible to a predefined matrix. This closeness is measured using the *LogDet divergence* $D_{\text{ld}}(\mathbf{A}, \mathbf{A}_0) := \text{tr}(\mathbf{A}\mathbf{A}_0^{-1}) - \log \det(\mathbf{A}\mathbf{A}_0^{-1}) - d$; and ITML is formulated as follows:

$$\begin{aligned} \min_{\mathbf{A} \succeq 0} \quad & D_{\text{ld}}(\mathbf{A}, \mathbf{A}_0) \\ \text{such that} \quad & d_{\mathbf{A}}(\mathbf{x}, \mathbf{y}) \leq u, \quad (\mathbf{x}, \mathbf{y}) \in \mathcal{S}, \\ & d_{\mathbf{A}}(\mathbf{x}, \mathbf{y}) \geq l, \quad (\mathbf{x}, \mathbf{y}) \in \mathcal{D}, \end{aligned} \quad (3)$$

where $u, v \in \mathbb{R}$ are threshold parameters, chosen to encourage distance between similar points to be small and between dissimilar points be large. Similar to ITML, Meyer

et al. (2011) propose the formulation

$$\begin{aligned} \min_{\mathbf{A} \succeq 0} \quad & \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{S}} \max(0, l - d_{\mathbf{A}}(\mathbf{x}_i, \mathbf{x}_j))^2 \\ & + \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{D}} \max(0, d_{\mathbf{A}}(\mathbf{x}_i, \mathbf{x}_j) - u)^2, \end{aligned} \quad (4)$$

for which they use Riemannian techniques to minimize the cost function. Although (4) does not use any regularizer, the authors observed good classification performance.

There exist several attempts for achieving high scalability with both the dimensionality and the number of constraints in the metric learning methods; some examples include (Shalev-Shwartz et al., 2004; Jain et al., 2009; Weinberger & Saul, 2008; Shalit et al., 2012).

However, the focus of our paper is different: we are concerned with the formulation of Euclidean metric learning. Remarkably, our new formulation admits a closed form solution, which turns out to be 3 orders of magnitude faster than established competing methods.

2. GMML: formulation and solution

As discussed above, the guiding idea behind Euclidean metric learning is to ultimately obtain a metric that yields “small” distances for similar points and “big” ones for dissimilar ones. Different metric learning methods try to fulfill this guideline either implicitly or explicitly.

The main idea that we introduce below is in how we choose to include the impact of the dissimilar points. Like one of earliest metric learning methods MMC, we propose to find a matrix \mathbf{A} that decreases the sum of distances over all the similar points, but unlike all previous methods, instead of treating dissimilar points asymmetrically, we propose to measure their interpoint distances using \mathbf{A}^{-1} , and to add their contribution to the overall objective. More precisely, we propose the following novel objective function:

$$\sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{S}} d_{\mathbf{A}}(\mathbf{x}_i, \mathbf{x}_j) + \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{D}} d_{\mathbf{A}^{-1}}(\mathbf{x}_i, \mathbf{x}_j). \quad (5)$$

In the sequel, we write $\hat{d}_{\mathbf{A}} \equiv d_{\mathbf{A}^{-1}}$ for brevity.

2.1. Insights

Let us provide some intuition behind our proposed objective (5). These insights are motivated by the idea that we may increase the Mahalanobis distance between dissimilar points $d_{\mathbf{A}}(\mathbf{x}, \mathbf{y})$ by decreasing $\hat{d}_{\mathbf{A}}(\mathbf{x}, \mathbf{y})$. The first idea is the simple observation that the distance $d_{\mathbf{A}}(\mathbf{x}, \mathbf{y})$ increases monotonically in \mathbf{A} , whereas the distance $\hat{d}_{\mathbf{A}}(\mathbf{x}, \mathbf{y})$ decreases monotonically in \mathbf{A} . This observation follows from the following well-known result:

Proposition 1. Let \mathbf{A}, \mathbf{B} be (strictly) positive definite matrices such that $\mathbf{A} \succ \mathbf{B}$. Then, $\mathbf{A}^{-1} \prec \mathbf{B}^{-1}$.

The second idea (which essentially reaffirms the first) is that the gradients of $d_{\mathbf{A}}$ and $\hat{d}_{\mathbf{A}}$ point in nearly opposite directions. Therefore, infinitesimally decreasing $d_{\mathbf{A}}$ leads to an increase in $\hat{d}_{\mathbf{A}}$. Indeed, the (Euclidean) gradient of $d_{\mathbf{A}}(\mathbf{x}, \mathbf{y})$ is

$$\frac{\partial d_{\mathbf{A}}}{\partial \mathbf{A}} = \mathbf{u}\mathbf{u}^T,$$

where $\mathbf{u} = \mathbf{x} - \mathbf{y}$; this is a rank-one positive semi-definite matrix. The gradient of $\hat{d}_{\mathbf{A}}(\mathbf{x}, \mathbf{y})$ is

$$\frac{\partial \hat{d}_{\mathbf{A}}}{\partial \mathbf{A}} = -\mathbf{A}^{-1}\mathbf{u}\mathbf{u}^T\mathbf{A}^{-1},$$

which is a rank-one matrix with a negative eigenvalue. It is easy to see that the inner product of these two gradients is negative, as desired.

2.2. Optimization problem and its solution

In the following, we further simplify the objective in (5). Rewriting the Mahalanobis distance using traces, we turn (5) into the optimization problem

$$\begin{aligned} \min_{\mathbf{A} \succ 0} \quad & \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{S}} \text{tr}(\mathbf{A}(\mathbf{x}_i - \mathbf{x}_j)(\mathbf{x}_i - \mathbf{x}_j)^T) \\ & + \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{D}} \text{tr}(\mathbf{A}^{-1}(\mathbf{x}_i - \mathbf{x}_j)(\mathbf{x}_i - \mathbf{x}_j)^T). \end{aligned} \quad (6)$$

We define now the following two important matrices:

$$\begin{aligned} \mathbf{S} &:= \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{S}} (\mathbf{x}_i - \mathbf{x}_j)(\mathbf{x}_i - \mathbf{x}_j)^T, \\ \mathbf{D} &:= \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{D}} (\mathbf{x}_i - \mathbf{x}_j)(\mathbf{x}_i - \mathbf{x}_j)^T, \end{aligned} \quad (7)$$

which denote the similarity and dissimilarity matrices, respectively. The matrices \mathbf{S} and \mathbf{D} are scaled second sample moments of the differences between similar points and the differences between dissimilar points. In the rest of this subsection, we assume that \mathbf{S} is a SPD matrix, which is a realistic assumption in many situations. For the cases where \mathbf{S} is just a positive semi-definite matrix, the regularized version can be used; we treat this case in Section 2.3.

Using (7), the minimization problem (6) yields the basic optimization formulation of GMML, namely

$$\min_{\mathbf{A} \succ 0} h(\mathbf{A}) := \text{tr}(\mathbf{A}\mathbf{S}) + \text{tr}(\mathbf{A}^{-1}\mathbf{D}). \quad (8)$$

The GMML cost function (8) has several remarkable properties, which may not be apparent at first sight. Below we highlight some of these to help build greater intuition, as well as to help us minimize it.

The first key property of $h(\mathbf{A})$ is that it is both strictly convex and strictly geodesically convex. Therefore, if $\nabla h(\mathbf{A}) = 0$ has a solution, that solution will be the global minimizer. Before proving this key property of h , let us recall some material that is also helpful for the remainder of the section.

Geodesic convexity is the generalization of ordinary (linear) convexity to (nonlinear) manifolds and metric spaces (Papadopoulos, 2005; Rapcsák, 1991). On Riemannian manifolds, geodesics are curves with zero acceleration that at the same time locally minimize the Riemannian distance between two points. The set of SPD matrices forms a Riemannian manifold of nonpositive curvature (Bhatia, 2007, Ch. 6). We denote this manifold by \mathbb{S}_+ . The geodesic curve joining \mathbf{A} to \mathbf{B} on the SPD manifold is denoted by

$$\mathbf{A}\sharp_t\mathbf{B} = \mathbf{A}^{1/2}(\mathbf{A}^{-1/2}\mathbf{B}\mathbf{A}^{-1/2})^t\mathbf{A}^{1/2}, \quad t \in [0, 1].$$

This notation for geodesic is customary, and in the literature, $\gamma(t)$ is also used. Moreover, the entire set of SPD matrices is geodesically convex, as there is a geodesic between every two points in the set. On this set, one defines geodesically convex functions as follows.

Definition 2. A function f on a geodesically convex subset of a Riemannian manifold is *geodesically convex*, if for all points \mathbf{A} and \mathbf{B} in this set, it satisfies

$$f(\mathbf{A}\sharp_t\mathbf{B}) \leq tf(\mathbf{A}) + (1-t)f(\mathbf{B}), \quad t \in [0, 1].$$

If for $t \in (0, 1)$ the above inequality is strict, the function is called strictly geodesically convex.

We refer the reader to (Sra & Hosseini, 2015) for more on geodesic convexity for SPD matrices. We are ready to state a simple but key convexity result.

Theorem 3. The cost function h in (8) is both strictly convex and strictly geodesically convex on the SPD manifold.

Proof. The first term in (8) is linear, hence convex, while the second term is strictly convex (Boyd & Vandenberghe, 2004, Ch. 3), viewing SPD matrices as a convex cone (see Rockafellar, 1970, Thm. 2.6). Thus, strict convexity of $h(\mathbf{A})$ is obvious. Therefore, we concentrate on proving its strict geodesic convexity. Using continuity, it suffices to show *midpoint strict convexity*, namely

$$h(\mathbf{A}\sharp_{1/2}\mathbf{B}) < \frac{1}{2}h(\mathbf{A}) + \frac{1}{2}h(\mathbf{B}).$$

It is well-known (Bhatia, 2007, Ch. 4) that for two distinct SPD matrices, we have the operator inequality

$$\mathbf{A}\sharp_{1/2}\mathbf{B} \prec \frac{1}{2}\mathbf{A} + \frac{1}{2}\mathbf{B}. \quad (9)$$

Since \mathbf{S} is SPD, it immediately follows that

$$\text{tr}((\mathbf{A}\sharp_{1/2}\mathbf{B})\mathbf{S}) < \frac{1}{2}\text{tr}(\mathbf{A}\mathbf{S}) + \frac{1}{2}\text{tr}(\mathbf{B}\mathbf{S}). \quad (10)$$

From the definition of \sharp_t , a brief manipulation shows that

$$(\mathbf{A}\sharp_t\mathbf{B})^{-1} = \mathbf{A}^{-1}\sharp_t\mathbf{B}^{-1}.$$

Thus, in particular for the midpoint (with $t = 1/2$) we have

$$\text{tr}((\mathbf{A}\sharp_{1/2}\mathbf{B})^{-1}\mathbf{D}) < \frac{1}{2}\text{tr}(\mathbf{A}^{-1}\mathbf{D}) + \frac{1}{2}\text{tr}(\mathbf{B}^{-1}\mathbf{D}). \quad (11)$$

Adding (10) and (11), we obtained the desired result. \square

Solution via geometric mean. The optimal solution to (8) will reveal one more reason why we invoke geodesic convexity. Since the constraint set of (8) is open and the objective is strictly convex, to find its global minimum, it is enough to find a point where the gradient ∇h vanishes. Differentiating with respect to \mathbf{A} , this yields

$$\nabla h(\mathbf{A}) = \mathbf{S} - \mathbf{A}^{-1}\mathbf{D}\mathbf{A}^{-1}.$$

Setting this gradient to zero results in the equation

$$\nabla h(\mathbf{A}) = 0 \implies \mathbf{A}\mathbf{S}\mathbf{A} = \mathbf{D}. \quad (12)$$

Equation (12) is a Riccati equation whose unique solution is nothing but the midpoint of the geodesic joining \mathbf{S}^{-1} to \mathbf{D} (see e.g., Bhatia (2007, 1.2.13)). Indeed,

$$\mathbf{A} = \mathbf{S}^{-1}\sharp_{1/2}\mathbf{D} = \mathbf{S}^{-1/2}(\mathbf{S}^{1/2}\mathbf{D}\mathbf{S}^{1/2})^{1/2}\mathbf{S}^{-1/2}.$$

Observe by construction this solution is SPD, therefore, the constraint of optimization is satisfied.

It is this fact that the solution to GMML is given by the midpoint of the geodesic joining the inverse of the second moment matrix of similar points to the second moment matrix of dissimilar points, which gives GMML its name: the midpoint of this geodesic is known as the *matrix geometric mean* and is a very important object in the study of SPD matrices (Bhatia, 2007, Ch. 6).

2.3. Regularized version

We have seen that the solution of our method is the geometric mean between \mathbf{S}^{-1} and \mathbf{D} . However, in practice the matrix \mathbf{S} might sometimes be non-invertible or near-singular. To address this concern, we propose to add a regularizing term to the objective function. This regularizer term can also be used to incorporate prior knowledge about the distance function. In particular, we propose to use

$$\min_{\mathbf{A} \succ 0} \lambda D_{\text{slid}}(\mathbf{A}, \mathbf{A}_0) + \text{tr}(\mathbf{A}\mathbf{S}) + \text{tr}(\mathbf{A}^{-1}\mathbf{D}), \quad (13)$$

where \mathbf{A}_0 is the ‘‘prior’’ (SPD matrix) and $D_{\text{slid}}(\mathbf{A}, \mathbf{A}_0)$ is the symmetrized LogDet divergence: $D_{\text{ld}}(\mathbf{A}, \mathbf{A}_0) + D_{\text{ld}}(\mathbf{A}_0, \mathbf{A})$, which is equal to

$$D_{\text{slid}}(\mathbf{A}, \mathbf{A}_0) := \text{tr}(\mathbf{A}\mathbf{A}_0^{-1}) + \text{tr}(\mathbf{A}^{-1}\mathbf{A}_0) - 2d, \quad (14)$$

where d is the dimensionality of the data. Interestingly, using (14) and following the argument as above, we see that the minimization problem in (13) with this regularizer also has a closed form solution. After straightforward computations, we obtain the following solution

$$\mathbf{A}_{\text{reg}} = (\mathbf{S} + \lambda\mathbf{A}_0^{-1})^{-1}\sharp_{1/2}(\mathbf{D} + \lambda\mathbf{A}_0), \quad (15)$$

the regularized geometric mean of suitably modified \mathbf{S} and \mathbf{D} matrices. Observe that as the regularization parameter $\lambda \geq 0$ increases, \mathbf{A}_{reg} becomes more similar to \mathbf{A}_0 .

2.4. Extension to weighted geometric mean

The geodesic viewpoint is also key to deciding how one may assign different ‘‘weights’’ to the matrices \mathbf{S} and \mathbf{D} when computing the GMML solution. This viewpoint is important because merely scaling the cost in (8) to change the balance between \mathbf{S} and \mathbf{D} is not meaningful as it only scales the resulting solution \mathbf{A} by a constant.

Given the geometric nature of the GMML’s solution, we replace the linear cost in (8) by a nonlinear one guided by Riemannian geometry of the SPD manifold. The key insight into obtaining a weighted version of GMML comes from a crucial geometric observation. *The minimum of (8) is also the minimum to the following optimization problem:*

$$\min_{\mathbf{A} \succ 0} \delta_R^2(\mathbf{A}, \mathbf{S}^{-1}) + \delta_R^2(\mathbf{A}, \mathbf{D}), \quad (16)$$

where δ_R denotes the Riemannian distance

$$\delta_R(\mathbf{X}, \mathbf{Y}) := \|\log(\mathbf{Y}^{-1/2}\mathbf{X}\mathbf{Y}^{-1/2})\|_{\text{F}} \quad \text{for } \mathbf{X}, \mathbf{Y} \succ 0,$$

on SPD matrices and $\|\cdot\|_{\text{F}}$ denotes the Frobenius norm.

Once we identify the solution of (8) with that of (16), the generalization to the weighted case becomes transparent. We introduce a parameter that characterizes the degree of balance between the cost terms of similarity and dissimilarity data. The weighted GMML formulation is then

$$\min_{\mathbf{A} \succ 0} h_t(\mathbf{A}) := (1-t)\delta_R^2(\mathbf{A}, \mathbf{S}^{-1}) + t\delta_R^2(\mathbf{A}, \mathbf{D}), \quad (17)$$

where t is a parameter that determines the balance. Unlike (8), which we observed to be strictly convex as well as strictly geodesically convex, problem (17) is *not* (Euclidean) convex. Fortunately, it is still geodesically convex, because δ_R itself is geodesically convex. The proof of the geodesic convexity of δ_R is more involved than that of Theorem 3, and we refer the reader to (Bhatia, 2007, Ch. 6) for complete details.

It can be shown, see e.g., (Bhatia, 2007, Ch. 6) that the unique solution to (17) is the weighted geometric mean

$$\mathbf{A} = \mathbf{S}^{-1}\sharp_t\mathbf{D}, \quad (18)$$




Figure 1. The solution of GMML is located in the geodesic between matrices S^{-1} and D on the manifold of SPD matrices.

Algorithm 1 Geometric Mean Metric Learning

Input: S : set of similar pairs, D : set of dissimilar pairs, t : step length of geodesic, λ : regularization parameter, A_0 : prior knowledge

Compute the similarity and dissimilarity matrices:

$$S = \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{S}} (\mathbf{x}_i - \mathbf{x}_j)(\mathbf{x}_i - \mathbf{x}_j)^T$$

$$D = \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{D}} (\mathbf{x}_i - \mathbf{x}_j)(\mathbf{x}_i - \mathbf{x}_j)^T$$

Return the distance matrix:

$$A = (S + \lambda A_0^{-1})^{-1} \#_t (D + \lambda A_0)$$

that is, a point on the geodesic from S^{-1} and D . Figure 1 illustrates this fact about the solution of GMML.

The regularized form of the previous solution is given by

$$A_{\text{reg}} = (S + \lambda A_0^{-1})^{-1} \#_t (D + \lambda A_0),$$

for $t \in [0, 1]$. In the cases where $t = 1/2$, it is equal to (15). This solution is our final and complete proposed solution to the linear metric learning problem. The summary of our GMML algorithm for metric learning is presented in Algorithm 1. Empirically, we have observed that the generalized solution (with free t) can significantly outperform the ordinary solution.

There are several approaches for fast computation of Riemannian geodesics for SPD matrices, for instance, Cholesky-Schur and scaled Newton methods (Iannazzo, 2011). We use Cholesky-Schur method in our paper to expedite the computation of Riemannian geodesics.

3. Results

In this section, we compare the performance of the proposed method GMML (Algorithm 1) to some well-known metric learning algorithms:

- ITML (Davis et al., 2007);
- LMNN (Weinberger & Saul, 2009); and
- FlatGeo with batch flat geometry (Meyer et al., 2011).

We exploit the commonly used criterion for comparing the performance of different methods, that is, the rate of the classification error for a k -NN classifier on different datasets. We choose $k = 5$, and estimate a full-rank matrix A in all methods.

3.1. Experiment 1

Assume c to be the number of classes, it is common in practice to generate $40c(c - 1)$ number of constraints by randomly choosing $40c(c - 1)$ pairs of points in a dataset. In our first experiment, shown in Figure 2, we use this number of constraints in our method in addition to ITML and FlatGeo methods. The LMNN method does not have this number of constraints parameter and we used a new version of its toolbox that uses Bayesian optimization for optimizing the model hyper-parameters. We use the default parameters used in ITML and FlatGeo, except we also use a minimum iterations of 10^4 for the FlatGeo method, because we observed that sometimes FlatGeo stops prematurely leading to a very poor performance. ITML has a regularization parameter that is set by using cross-validation.

Figure 2 reports the results for the smaller datasets. The datasets are obtained from the well-known UCI repository (Asuncion & Newman, 2007). In the plot, the baseline of using Euclidean distance for classification is shown in yellow. It can be seen that GMML outperforms the other three metric learning methods.

The figure reports 40 runs of a two-fold splitting of the data. In each run, the data is randomly divided into two equal sets. The regularization parameter λ is set to zero for most of the datasets. We only add a small value of λ when the similarity matrix S becomes singular. For example, since the similarity matrix of the Segment data is near singular, we use the regularized version of our method with $\lambda = 0.1$ and A_0 equals to the identity matrix.

We use five-fold cross-validation for choosing the best parameter t . We tested 18 different values for t in a two-step method. In the first step the best t is chosen among the values $\{0.1, 0.3, 0.5, 0.7, 0.9\}$. Then in the second step, 12 values of t are tested within an interval of length 0.02 in the window around the previously selected point.




Figure 2. Classification error rates of k -nearest neighbor classifier via different learned metrics for different small datasets. Numbers below each correspond to the dimensionality of feature space in the data (d), number of classes (c) and number of total data (n).

Figure 3 shows the effect of the parameter t on the average accuracy of k -NN classifier for five datasets. These datasets are also appeared in Figure 2. It is obvious that in some datasets, going from the ordinary version to the extended version can make the GMML’s performance substantially better. Observe that each curve has a convex-like shape with some wiggling. That is why we choose the above approach for finding the best t , and we can verify its precision by comparing Figures 2 and 3.




Figure 3. Classification error rates of k -nearest neighbor classifier along with GMML for different values of the parameter t . We analyze five datasets here, which is also appeared in Figure 2.

3.2. Experiment 2

To evaluate the performance of our method on larger datasets, we conduct a second set of experiments. The results can be summarized in Figure 4. The datasets in this experiment are Isolet, Letters (Asuncion & Newman, 2007), MNIST³ (LeCun et al., 1998) and USPS (Le Cun et al., 1990).

Figure 4 reports the average classification error over 5 runs of random splitting of the data. We use three-fold cross-validation for adjusting the parameter t . Since the similarity matrices of the MNIST data were not invertible, we use the regularized version of our method with regularization parameter $\lambda = 0.1$. The prior matrix \mathbf{A}_0 is set to the identity matrix.

On two of the large datasets, Letters and USPS, our method achieves the same performance as the best competing method that is LMNN. For one of the datasets our method significantly outperforms LMNN, and in one dataset it is significantly outdone by LMNN. We also observed that by using more data pairs for generating the similarity and dissimilarity matrices, the performance of our method on Isolet and MNIST datasets improves. We tested $1000c(c-1)$ for these two datasets, with which we achieve about 1 percent better accuracy for Isolet leading to slightly better performance than FlatGeo approach. For MNIST data, we achieved about 0.5 percent better accuracy.

³We used a smaller version of the MNIST dataset available in www.cad.zju.edu.cn/home/dengcai/Data/MLData.html




Figure 4. Classification error rates of k -nearest neighbor classifier via different learned metrics for large datasets.

Table 1. Running time (in seconds) of metric learning methods

DATA SET	GMML	LMNN	ITML	FLATGEO
SEGMENT	0.0054	77.595	0.511	63.074
LETTERS	0.0137	401.90	7.053	13543
USPS	0.1166	811.2	16.393	17424
ISOLET	1.4021	3331.9	1667.5	24855
MNIST	1.6795	1396.4	1739.4	26640

The average running times of the methods on all large data sets and one small dataset are shown in Table 1. The running time of different methods is reported for only one run of each algorithm for fixed values of hyper-parameters; that means, the reported run times *do not* include the time required to select the hyper-parameters. All methods were implemented on MATLAB R2014a (64-bit), and the simulations were run on a personal laptop with an Intel Core i5 (2.5Ghz) processor under the OS X Yosemite operating system.

It can be seen that our method is several order of magnitudes faster than other methods. In addition to obtaining good classification accuracy using the proposed method, the computational complexity of our method is another nice property making it an interesting candidate for large-scale metric learning.

4. Conclusion and future work

We revisited the task of learning a Euclidean metric from weakly supervised data given as pairs of similar and dissimilar points. Building on geometric intuition, we approached the task of learning a symmetric positive definite matrix by formulating it as a smooth, strictly convex optimization problem (thus, ensuring a unique solution). Remarkably, our formulation was shown to have a closed form solution. We also viewed our formulation as an optimization problem on the Riemannian manifold of SPD matri-

ces, a viewpoint that proved crucial to obtaining a weighted generalization of the basic formulation. We also presented a regularized version of our problem. In all cases, the solution could be obtained as a closed form “matrix geometric mean”, thus explaining our choice of nomenclature.

We experimented with several datasets, both large and small, in which we compared the classification accuracy of a k -NN classifier using metric learned via various competing methods. In addition to good classification accuracy and global optimality, our proposed method for solving the metric learning problem has other nice properties like being fast and being scalable with regard to both the dimensionality d and the number of training samples n .

Given the importance of metric learning to a vast number of applications, we believe that the new understanding offered by our formulation, its great simplicity, and its tremendous speedup over widely used methods make it attractive.

4.1. Future work

Several avenues of future work are worth pursuing. We list some most promising directions below:

- To view our metric learning methods as a dimensionality reduction method; here the connections in (Cunningham & Ghahramani, 2015) may be helpful.
- Extensions of our simple geometric framework to learn nonlinear and local metrics.
- Applying the idea of using concurrently the Mahalanobis distance d_A with its counterpart \hat{d}_A on the other machine learning problems.

References

- Asuncion, Arthur and Newman, David. UCI machine learning repository, 2007.
- Bellet, Aurélien, Habrard, Amaury, and Sebban, Marc. A survey on metric learning for feature vectors and structured data. *arXiv preprint arXiv:1306.6709*, 2013.
- Bhatia, R. *Positive definite matrices*. Princeton University Press, 2007.
- Boyd, Stephen and Vandenberghe, Lieven. *Convex optimization*. Cambridge University Press, 2004.
- Cunningham, John P and Ghahramani, Zoubin. Linear dimensionality reduction: Survey, insights, and generalizations. *Journal of Machine Learning Research*, 2015.
- Davis, Jason V, Kulis, Brian, Jain, Prateek, Sra, Suvit, and Dhillon, Inderjit S. Information-theoretic metric learning. In *Proceedings of the 24th international conference on Machine learning*, pp. 209–216. ACM, 2007.

- Globerson, Amir and Roweis, Sam T. Metric learning by collapsing classes. In *Advances in neural information processing systems*, pp. 451–458, 2005.
- Iannazzo, Bruno. The geometric mean of two matrices from a computational viewpoint. *arXiv preprint arXiv:1201.0101*, 2011.
- Jain, Prateek, Kulis, Brian, Dhillon, Inderjit S, and Grauman, Kristen. Online metric learning and fast similarity search. In *Advances in neural information processing systems*, pp. 761–768, 2009.
- Kulis, Brian. Metric learning: A survey. *Foundations and Trends in Machine Learning*, 5(4):287–364, 2012.
- Le Cun, B Boser, Denker, John S, Henderson, D, Howard, Richard E, Hubbard, W, and Jackel, Lawrence D. Handwritten digit recognition with a back-propagation network. In *Advances in neural information processing systems*, 1990.
- LeCun, Yann, Bottou, Léon, Bengio, Yoshua, and Haffner, Patrick. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Meyer, Gilles, Bonnabel, Silvère, and Sepulchre, Rodolphe. Regression on fixed-rank positive semidefinite matrices: A riemannian approach. *The Journal of Machine Learning Research*, 12:593–625, 2011.
- Papadopoulos, A. *Metric spaces, convexity and nonpositive curvature*. European Mathematical Society, 2005.
- Rapcsák, T. Geodesic convexity in nonlinear optimization. *Journal of Optimization Theory and Applications*, 69(1): 169–183, 1991.
- Rockafellar, Ralph Tyrell. *Convex analysis*. Princeton University Press, 1970.
- Shalev-Shwartz, Shai, Singer, Yoram, and Ng, Andrew Y. Online and batch learning of pseudo-metrics. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 94, 2004.
- Shalit, Uri, Weinshall, Daphna, and Chechik, Gal. Online learning in the embedded manifold of low-rank matrices. *The Journal of Machine Learning Research*, 13(1):429–458, 2012.
- Sra, Suvrit and Hosseini, Reshad. Conic geometric optimization on the manifold of positive definite matrices. *SIAM Journal on Optimization*, 25(1):713–739, 2015.
- Weinberger, Kilian Q and Saul, Lawrence K. Fast solvers and efficient implementations for distance metric learning. In *Proceedings of the 25th international conference on Machine learning*, pp. 1160–1167, 2008.
- Weinberger, Kilian Q and Saul, Lawrence K. Distance metric learning for large margin nearest neighbor classification. *The Journal of Machine Learning Research*, 10: 207–244, 2009.
- Xing, Eric P, Jordan, Michael I, Russell, Stuart, and Ng, Andrew Y. Distance metric learning with application to clustering with side-information. In *Advances in neural information processing systems*, pp. 505–512, 2002.