Random Projections on Manifolds of Symmetric Positive Definite Matrices for Image Classification

Azadeh Alavi, Arnold Wiliem, Kun Zhao, Brian C. Lovell, Conrad Sanderson

NICTA, GPO Box 2434, Brisbane, QLD 4001, Australia
University of Queensland, School of ITEE, QLD 4072, Australia
Queensland University of Technology, Brisbane, QLD 4000, Australia

Abstract

Recent advances suggest that encoding images through Symmetric Positive Definite (SPD) matrices and then interpreting such matrices as points on Riemannian manifolds can lead to increased classification performance. Taking into account manifold geometry is typically done via (1) embedding the manifolds in tangent spaces, or (2) embedding into Reproducing Kernel Hilbert Spaces (RKHS). While embedding into tangent spaces allows the use of existing Euclidean-based learning algorithms, manifold shape is only approximated which can cause loss of discriminatory information. The RKHS approach retains more of the manifold structure, but may require non-trivial effort to kernelise Euclidean-based learning algorithms. In contrast to the above approaches, in this paper we offer a novel solution that allows SPD matrices to be used with unmodified Euclidean-based learning algorithms, with the true manifold shape well-preserved. Specifically, we propose to project SPD matrices using a set of random projection hyperplanes over RKHS into a random projection space, which leads to representing each matrix as a vector of projection coefficients. Experiments on face recognition, person re-identification and texture classification show that the proposed approach outperforms several recent methods, such as Tensor Sparse Coding, Histogram Plus Epitome, Riemannian Locality Preserving Projection and Relational Divergence Classification.

1. Introduction

Covariance matrices have recently been employed to describe images and videos [22, 13, 35], as they are known to provide compact and informative feature description [7, 3]. Non-singular covariance matrices are naturally symmetric positive definite matrices (SPD) which form connected Riemannian manifolds when endowed with a Riemannian metric over tangent spaces [20]. As such, the Riemannian geometry needs to be considered for solving learning tasks [35].

One of the most widely used metrics for SPD matrices is the Affine Invariant Riemannian Metric (AIRM) [22].

The AIRM induces Riemannian structure which is invariant to inversion and similarity transforms. Despite its properties, learning methods using this approach have to deal with computational challenges, such as employing computationally expensive non-linear operators.

To address the above issue, two lines of research have been proposed: (1) embedding manifolds into tangent spaces [21, 24, 27, 35, 36]; (2) embedding into Reproducing Kernel Hilbert Spaces (RKHS), induced by kernel functions [2, 14, 16, 17, 28, 30, 34]. The former approaches in effect map manifold points to Euclidean spaces, thereby enabling the use of existing Euclidean-based learning algorithms. This comes at the cost of disregarding some of the manifold structure. The latter approach addresses this by implicitly mapping points on the manifold into RKHS, which can be considered as a high dimensional Euclidean space. Training data can be used to define a space that preserves manifold geometry [17]. The downside is that existing Euclidean-based learning algorithms need to be kernelised, which may not be trivial. Furthermore, the resulting methods can still have high computational load, making them impractical to use in more complex scenarios.

Contributions. In this paper we offer a novel approach for analysing SPD matrices which combines the main advantage of tangent space approaches with the discriminatory power provided by kernel space methods. We adapt a recent idea from techniques specifically designed for learning tasks in very large image datasets [12, 19]. In this domain, image representations are mapped into a reduced space wherein the similarities are still well-preserved [19]. In our proposed approach, we employ such a mapping technique to create a space which preserves the manifold geometry while can be considered as Euclidean.

Specifically, we first embed SPD manifold points into RKHS via the Stein Divergence Kernel [33]. We then generate random projection hyperplanes in RKHS and project the embedded points via the method proposed in [19]. Finally, as the underlying space can be thought as Euclidean, any appropriate Euclidean-based learning machinery can be applied. In this paper, we study the efficacy of this embedding method for classification tasks. We show that the

space is only as effective as the completeness of the training data generating the random projection hyperplanes, and address this through the use of synthetic data to augment training data. Experiments on several vision tasks (person re-identification, face recognition and texture recognition), show that the proposed approach outperforms several state-of-the-art methods.

We continue the paper as follows. Section 2 provides a brief overview of the manifold structure and its associated kernel function. We then detail the proposed approach in Section 3. Section 4 presents results on the study of the random projection space discriminability as well as comparisons with the state-of-the-art results in various visual classification tasks. The main findings and possible future directions are summarised in Section 5.

2. Manifold Structure and Stein Divergence

Consider $\{X_1 \dots X_n\} \in Sym_+^d$ to be a set of non-singular $d \times d$ -sized covariance matrices, which are symmetric positive definite (SPD) matrices. These matrices belong to a smooth differentiable topological space, known as SPD manifolds. In this work, we endow the SPD manifold with the AIRM to induce the Riemannian structure [22]. As such, a point on manifold $\mathcal{M}$ can be mapped to a tangent space using:

$$\log_{X_i} X_j = X_i^{\frac{1}{2}} \log(X_i^{-\frac{1}{2}} X_j X_i^{-\frac{1}{2}}) X_i^{\frac{1}{2}} \quad (1)$$

where $X_i, X_j \in Sym_+^d$, X_i is the point where the tangent space is located (*i.e.* tangent pole) and X_j is the point that we would like to map into the tangent space $\mathcal{T}_{X_i}\mathcal{M}$; $\log(\cdot)$ is the matrix logarithm. The inverse function of this maps points on a particular tangent space into the manifold is:

$$\exp_{X_i} y = X_i^{\frac{1}{2}} \exp(X_i^{-\frac{1}{2}} y X_i^{-\frac{1}{2}}) X_i^{\frac{1}{2}} \quad (2)$$

where $X_i \in Sym_+^d$ is again the tangent pole; $y \in \mathcal{T}_{X_i}\mathcal{M}$ is a point in the tangent space $\mathcal{T}_{X_i}\mathcal{M}$; $\exp(\cdot)$ is the matrix exponential.

From the above functions, we now define the shortest distance between two points on the manifold. The distance, here called geodesic distance, is represented as the minimum length of the curvature path that connects two points [22]:

$$d_g^2(X_i, X_j) = \text{trace} \left\{ \log^2(X_i^{-\frac{1}{2}} X_j X_i^{-\frac{1}{2}}) \right\} \quad (3)$$

The above mapping functions can be computationally expensive. We can also use the recently introduced Stein divergence [33] to determine similarities between points on the SPD manifold. Its symmetrised form is:

$$J_\phi(X, Y) \triangleq \log \left(\det \left(\frac{X + Y}{2} \right) \right) - \frac{1}{2} \log(\det(XY)) \quad (4)$$

The Stein divergence kernel can then be defined as:

$$K(X, Y) = \exp\{-\sigma J_\phi(X, Y)\} \quad (5)$$

under the condition of $\sigma \in \{\frac{1}{2}, \frac{2}{2}, \dots, \frac{d-1}{2}\}$ to ensure that the kernel matrix formed by Eqn. (5) is positive definite [15].

3. Random Projection on RKHS

We aim to address classification tasks, originally formulated on the manifold, by embedding them into a random projection space, which can be considered as Euclidean, while still honouring the manifold geometry structure. To this end, we propose to use random projection on RKHS with the aid of the Stein divergence kernel.

Random projection is an approximation approach for estimating distances between pairs of points in a high-dimensional space [1]. In essence, the projection of a point $u \in \mathbb{R}^d$ can be done via a set of randomly generated hyperplanes $\{r_1 \dots r_k\} \in \mathbb{R}^d$:

$$f(u) = u^\top R \quad (6)$$

where $R \in \mathbb{R}^{d \times k}$ is the matrix wherein each column contains a single hyperplane r_i ; $f(\cdot)$ is the mapping function which maps any point in $\mathbb{R}^d$ into a random projection space $\mathbb{R}^k$. According to the Johnson-Lindenstrauss lemma [1], it is possible to map a set of high-dimensional points into much lower dimensional space wherein the pairwise distance between two points are well-preserved:

Lemma 3.1. Johnson-Lindenstrauss Lemma. *For any ϵ such that $\frac{1}{2} > \epsilon > 0$, and any set of points $S \in \mathbb{R}^d$ with $|S| = n$ upon projection to a uniform random k -dimension subspace where $k = O(\log n)$, the following property holds with probability at least $\frac{1}{2}$ for every pair $u, v \in S$, $(1 - \epsilon)||u - v||^2 \leq ||f(u) - f(v)||^2 \leq (1 + \epsilon)||u - v||^2$, where $f(u), f(v)$ are projection of u, v .*

Despite the success of numerous approaches using this lemma to accomplish various computer vision tasks, most of them restrict the distance function to the ℓ_p norm, Mahalanobis metric or inner product [5, 8, 18], which makes them incompatible for non-Euclidean spaces. Recently, Kulis and Grauman [19] proposed a method that allows the distance function to be evaluated over RKHS. Thus, it is possible to apply the lemma for any arbitrary kernel $\mathbb{K}(i, j) = K(X_i, X_j) = \phi(X_i)^\top \phi(X_j)$ for an unknown embedding $\phi(\cdot)$ which maps the points to a Hilbert space $\mathcal{H}$ [19]. This approach makes it possible for one to construct a random projection space on an SPD manifold, where the manifold structure is well-preserved.

The main idea of our proposed approach, denoted as **Random Projection On SPD manifold for Image Classification (ROSE)**, is to first map all points on the manifold into RKHS, with implicit mapping function $\phi(\cdot)$, via the Stein divergence kernel. This is followed by mapping all the points in the RKHS $\phi(X_i) \in \mathcal{H}$ into a random projection space $\mathbb{R}^k$. To achieve this we follow the Kulis-Grauman approach [19] by randomly generating a set of hyperplanes over the RKHS $\{r_1 \dots r_k\} \in \mathcal{H}$ which is restricted to be approximately Gaussian. As the embedding function $\phi(\cdot)$

is unknown, then the generation process is done indirectly via a weighted sum of the subset of the given training sets.

To this end, consider each data point $\phi(\mathbf{X}_i)$ from the training set as a vector from some underlying distribution D with unknown mean μ and unknown covariance Σ . Let S be a set of t training exemplars chosen i.i.d. from D , then $\mathbf{z}_t = \frac{1}{t} \sum_{i \in S} \phi(\mathbf{X}_i)$ is defined over S . According to the central limit theorem for sufficiently large t , the random vector $\tilde{\mathbf{z}}_t = \sqrt{t}(\mathbf{z}_t - \mu)$ is distributed according to the multi-variate Gaussian $\mathcal{N}(\mu, \Sigma)$ [26]. Then if a whitening transform is applied, it results in $\mathbf{r}_i = \Sigma^{-\frac{1}{2}} \tilde{\mathbf{z}}_t$ which follows $\mathcal{N}(0, \mathbf{I})$ distribution in Hilbert space $\mathcal{H}$. Therefore the i -th coefficient of each vector in the random projection space is defined as:

$$\phi(\mathbf{X}_i)^T \Sigma^{-\frac{1}{2}} \tilde{\mathbf{z}}_t \quad (7)$$

The mean μ and covariance Σ need to be approximated from training data. A set of p objects is chosen to form the first p items of a reference set: $\phi(\mathbf{X}_1), \dots, \phi(\mathbf{X}_p)$. Then the mean is implicitly estimated as $\mu = \frac{1}{p} \sum_{i=1}^p \phi(\mathbf{X}_i)$, and the covariance matrix Σ is also formed over the p samples. Eqn. (7) can be solved using a similar approach as for Kernel PCA, which requires projecting onto the eigenvectors of the covariance matrix [19]. Let the eigendecomposition of Σ be UVU^T , then $\Sigma^{-\frac{1}{2}} = UV^{\frac{1}{2}}U^T$, and therefore Eqn. (7) can be rewritten as [19]:

$$\phi(\mathbf{X}_i)^T UV^{\frac{1}{2}}U^T \tilde{\mathbf{z}}_t \quad (8)$$

Let then define $\mathbf{K}$ as a kernel matrix over the p randomly selected training points, where $\mathbf{K} = \Lambda \Theta \Lambda^T$. Based on the fact that the non zero eigenvalues of V are equal to the non-zero eigenvalues of Θ , Kulis-Grauman [19] showed that Eqn. (8) is equivalent to:

$$\sum_{i=1}^p \omega(i) (\phi(\mathbf{X}_i)^T \phi(\mathbf{X})) \quad (9)$$

where

$$\omega(i) = \frac{1}{t} \sum_{j=1}^p \sum_{l \in S} K_{ij}^{-\frac{1}{2}} K_{jl} - \frac{1}{p} \sum_{j=1}^p \sum_{k=1}^p K_{ij}^{-\frac{1}{2}} K_{jk} \quad (10)$$

where, for S , a set of t points are randomly selected from the p sampled points. The expression $w(i)$ in Eqn. (10) can be further simplified by defining $\mathbf{e}$ as a vector of all ones, and $\mathbf{e}_S$ as a zero vector with ones in the entries corresponding to the indices of S [19]:

$$\mathbf{w} = \mathbf{K}^{\frac{1}{2}} \left(\frac{1}{t} \mathbf{e}_S - \frac{1}{p} \mathbf{e} \right) \quad (11)$$

In terms of calculating the computational complexity of the training algorithm, according to Eqns. (9) and (11), the most expensive step is in the single offline computation of $\mathbf{K}^{\frac{1}{2}}$, which takes $O(p^3)$. The computational complexity of classifying a query point depends then on three factors: computing the kernel vector which requires $O(pd^3)$

operations, projecting the resulting kernel vector into random hyperplane which demands $O(pt)$ operations (where $t < p$), and finally applying a classifier in the projection space which can be done with one versus all support vector machine $O(nb)$ operations, where n is the number of classes and b is the number of hyperplanes used in defining the hyperplane [19]. Hence the complexity of classification for a single query data is equal to $O(p^3 + pt + nb)$ which is more efficient than Relational Divergence based Classification (RDC) [2], which is later shown to be the second best approach in the experiment part. The RDC represents Riemannian points as similarity vectors to a set of training points. As similarity vectors are in Euclidean space, RDC then employs Linear Discriminant Analysis as a classifier.

3.1. Synthetic Data

As later shown in the experiment section (for instance the result shown in Fig. 2), the discriminative power of the random projection space depends heavily on the training set which generates the random projection hyperplanes. To overcome this limitation, we propose to use generated synthetic SPD matrices $\mathbf{X}_1, \dots, \mathbf{X}_n \in \text{Sym}_+^d$ centred around the mean of the data (denoted by $\mathbf{X}_\mu$), where the mean of the training set can be determined intrinsically via the Karcher mean algorithm [22].

We relate the synthetic data to the training set, by enforcing the following condition on the synthetic SPD matrices:

$$\forall \mathbf{X}_j \in S \text{ and } \forall \mathbf{X}_i \in G : \quad (12)$$

$$d_g(\mathbf{X}_\mu, \mathbf{X}_j) \leq \max(d_g(\mathbf{X}_\mu, \mathbf{X}_i))$$

where G is the training set, S is a set of t training exemplars chosen i.i.d. from some underlying distribution D , $\mathbf{X}_\mu$ is the mean of the training set and $\mathbf{X}_j$ is a generated synthetic point.

The constraint in Eqn. (12) considers a ball around the mean of the training data, with the radius equal to the longest calculated distance between mean and the given training points:

$$r = \max(d_g(\mathbf{X}_\mu, \mathbf{X}_i)) \quad (13)$$

Then we need to generate SPD matrices which are located within r radius from the mean (Eqn. 13). It is not trivial to generate SPD matrices which follow Eqn. (12), as it establishes a relation between generate SPD matrices and the original training points. To address this, we apply the relationship between the geodesic distance and the given Riemannian metric in a tangent space. Let $\mathbf{X}_i, \mathbf{X}_j \in \text{Sym}_+^d$ be two points on the manifold and $\mathbf{x}_i, \mathbf{x}_j \in \mathcal{T}_{\mathbf{X}_i} \mathcal{M}$ be the corresponding points on the tangent space $\mathcal{T}_{\mathbf{X}_i} \mathcal{M}$. The norm of vector $\bar{\mathbf{x}}_i \mathbf{x}_j$ is equivalent to $d_g(\mathbf{X}_i, \mathbf{X}_j)$ [22]. Therefore, it is possible to find a point $\mathbf{Y}_i$ along the geodesic $\mathbf{X}_i$ and $\mathbf{X}_j$ whose geodesic distance to $\mathbf{X}_i$ satisfies (12).

Along with the above definitions, we introduce the following definition and proposition:

Definition 3.2. Any point on an SPD manifold $\mathbf{X}_i \in \text{Sym}_+^d$ is said to have normalised geodesic distance with respect to $\mathbf{X}_j \in \text{Sym}_+^d$ if and only if $d_g(\mathbf{X}_i, \mathbf{X}_j) = 1$.

Proposition 3.3. For any two SPD matrices $\mathbf{X}, \mathbf{X}_\mu \in \text{Sym}_+^d$, there exists $\mathbf{X}_g$ on the geodesic curve defined on $\mathbf{X}$ and $\mathbf{X}_\mu$, which has normalised geodesic distance with respect to $\mathbf{X}_\mu$. The point $\mathbf{X}_g$ can be determined via: $\mathbf{X}_\mu^{\frac{1}{2}} \left(\mathbf{X}_\mu^{-\frac{1}{2}} \mathbf{X} \mathbf{X}_\mu^{-\frac{1}{2}} \right)^c \mathbf{X}_\mu^{\frac{1}{2}}$, where $c = \frac{\zeta}{d_g(\mathbf{X}, \mathbf{X}_\mu)}$, for $\zeta = 1$.

To prove the above proposition, we let $\mathbf{X}, \mathbf{X}_\mu \in \text{Sym}_+^d$ to be two given points on an SPD manifold. In order to normalise the geodesic distance of $\mathbf{X}$ with respect to $\mathbf{X}_\mu$, we map point $\mathbf{X}$ into tangent space $\mathcal{T}_{\mathbf{X}_\mu} \mathcal{M}$. As a tangent space is considered as Euclidean space where the distance between $\mathbf{X}$ and tangent pole $\mathbf{X}_\mu$ is preserved, Euclidean vector normalisation can be applied. Finally the normalised point is mapped back to the manifold. These steps can be presented as:

$$\mathbf{X}_g = \exp_{\mathbf{X}_\mu} \left(\frac{\zeta}{d_g(\mathbf{X}_\mu, \mathbf{X})} \log_{\mathbf{X}_\mu}(\mathbf{X}) \right) \quad (14)$$

By plugging in (1) and (2) we obtain:

$$\mathbf{X}_g = \mathbf{X}_\mu^{\frac{1}{2}} \exp \left(\frac{\zeta}{d_g(\mathbf{X}_\mu, \mathbf{X})} \log(\mathbf{X}_\mu^{-\frac{1}{2}} \mathbf{X} \mathbf{X}_\mu^{-\frac{1}{2}}) \right) \mathbf{X}_\mu^{\frac{1}{2}}$$

If we let $c = \frac{\zeta}{d_g(\mathbf{X}_\mu, \mathbf{X})}$, based on the fact that $\mathbf{X}$ and $\mathbf{X}_\mu$ are SPD matrices, we arrive at:¹

$$\mathbf{X}_g = \mathbf{X}_\mu^{\frac{1}{2}} \exp \left(\log \left((\mathbf{X}_\mu^{-\frac{1}{2}} \mathbf{X} \mathbf{X}_\mu^{-\frac{1}{2}})^c \right) \right) \mathbf{X}_\mu^{\frac{1}{2}}$$

which proves that:

$$\mathbf{X}_g = \mathbf{X}_\mu^{\frac{1}{2}} \left(\mathbf{X}_\mu^{-\frac{1}{2}} \mathbf{X} \mathbf{X}_\mu^{-\frac{1}{2}} \right)^c \mathbf{X}_\mu^{\frac{1}{2}} \quad (15)$$

Having ζ is equal to 1 results a normalised geodesic distance with respect to $\mathbf{X}_\mu$. However in our case to satisfy Eqn. (12), we use $\zeta = \delta \times \max(d_g(\mathbf{X}_\mu, \mathbf{X}_j))$, where $\delta \in [0, 1]$ is randomly generated number according to uniform distribution.

4. Experiments and Discussion

We consider three computer vision classification tasks: (1) texture classification [25]; (2) face recognition [23] and (3) person re-identification [9]. We first detail the experiment set up for each application and discuss our results for the comprehensive study of the random projection space discriminability on the tasks. To this end, we first embed the SPD matrices into RKHS via the Stein divergence kernel, followed by mapping the embedded data points into a random projection space. The resulting vectors are then

¹ See the appendix for proof of $\log \mathbf{X}^c = c \log \mathbf{X}$, where $\mathbf{X} \in \text{Sym}_+^d$.

Figure 1. Top row: example of pedestrians in the ETHZ dataset [29]. Middle row: example images from the Brodatz texture dataset [25]. Bottom row: examples of closely-cropped faces from the FERET ‘b’ subset [23].

fed to a linear Support Vector Machine classifier, which uses a one-versus-all configuration for multi-class classification [10, 30]

The parameter settings are as follows. As suggested in [19], we have used $t = \min(30, \frac{1}{4}n)$, where n is the number of samples chosen to create each hyperplane. For the number of the random hyperplanes we have used validation data to choose one of n , $2n$ or $3n$. Based on empirical observations on validation sets, the number of synthetic samples was chosen as either n or m , where m be the number of samples per class. In a similar manner, the value of σ in Eqn. (5) was chosen from $\{1, 2, \dots, 20\}$.

We compare our proposed method, here denoted as **Random Projection On SPD manifold for Image Classification (ROSE)**, with several other embedding approaches (tangent spaces, RKHS and hashing) as well as several state-of-the-art methods. We also evaluate the effect of augmenting the training data with synthetic data points, and refer to this approach as **ROSE with Synthetic data (ROSES)**.

For the person re-identification task we used the modified version [29] of the ETHZ dataset [9]. The dataset was captured from a moving camera, with the images of pedestrians containing occlusions and wide variations in appearance. Sequence 1 contains 83 pedestrians, and Sequence 2 contains 35 (Fig. 1). Following [2], we first downsampled all the images and then created the training set using 10 randomly selected images, while the rest were used to shape the test set. The random selection of the training and testing data was repeated 20 times. Each image was represented as a covariance matrix of feature vectors obtained at each pixel location:

$$F_{x,y} = [x, y, R_{x,y}, G_{x,y}, B_{x,y}, R'_{x,y}, G'_{x,y}, B'_{x,y}, R''_{x,y}, G''_{x,y}, B''_{x,y}]$$

where x and y represent the position of a pixel, while $R_{x,y}$, $G_{x,y}$ and $B_{x,y}$ represent the corresponding colour information. $C'_{x,y}$, $C''_{x,y}$ represent the gradient and Laplacian for colour C , respectively.

For the task of texture classification, we use the Brodatz dataset [25]. See Fig. 1 for examples. We follow the test protocol presented in [32]. Accordingly, nine test scenarios with various number of classes were generated. To create SPD matrices, we follow [2] by down-sampling each image and then splitting it into 64 regions. A feature vector for each pixel $I(x, y)$ is calculated as $F(x, y) = [I(x, y), |\frac{\partial I}{\partial x}|, |\frac{\partial I}{\partial y}|, |\frac{\partial^2 I}{\partial x^2}|, |\frac{\partial^2 I}{\partial y^2}|]$. Each region is described by a covariance matrix formed from these vectors. For each test scenario, we randomly select 25 covariance matrices per class to construct training set and the rest is used to create the testing set. The random selection was repeated 10 times and the mean results are reported.

For face recognition task, the ‘b’ subset of the FERET dataset [23] is used. Each image is first closely cropped to include only the face and then downsampled (Fig. 1). The tests with various pose angles were created to evaluate the performance of the method. The training set consists of frontal images with illumination, expression and small pose variations. Non-frontal images are used to create the test sets. Each face image is represented by a covariance matrix, where for every pixel $I(x, y)$ the following feature vector is computed:

$$F_{x,y} = [I(x, y), x, y, |G_{0,0}(x, y)|, \dots, |G_{4,7}(x, y)|] \quad (16)$$

where $G_{u,v}(x, y)$ is the response of a 2D Gabor wavelet centred at x, y .

4.1. Random Projection Space Discriminability

We first compare the performance of the proposed ROSE method with several other embedding methods: (1) Kernel SVM (KSVM) using the Stein divergence kernel, (2) Kernelised Locality-Sensitive Hashing (KLSH) [19], and (3) Riemannian Spectral Hashing (RSH), a hashing method specifically designed for smooth manifolds [6].

Tables 1, 2 and 3 report the results for each dataset. ROSE considerably outperforms the other embedding methods on the texture and person re-identification applications, while being on par with KLSH on the face recognition task. This suggests that the random projection space constructed

Table 1. Recognition accuracy (in %) for the person re-identification task on Seq. 1 and Seq. 2 of the ETHZ dataset; KSVM: Kernel SVM; KLSH: Kernelised Locality-Sensitive Hashing; RSH: Riemannian Spectral Hashing. ROSE is the proposed method, and ROSES is ROSE augmented with synthetic data.

	KSVM	KLSH	RSH	ROSE	ROSES
Seq.1	72.0	81.0	58.5	90.7	92.5
Seq.2	79.0	84.0	62.7	91.5	94.0
Average	75.5	82.5	60.6	91.2	93.2

by the random hyperplanes over RKHS offers sufficient discrimination for the classification tasks. In fact, as we use linear SVM for the classifier, the results presented here follow the theoretical results from [31] which suggest that the margin for the SVM classifier is still well-preserved after the random projection.

We apply the ROSES method (ROSE augmented with synthetic data) on the three tasks in order to take a closer look at the contribution of the training data generating the random projection hyperplanes for space discriminability. As shown in the results, there is notable improvement over ROSE in the ETHZ person re-identification as well as Brodatz texture classification datasets. However, using synthetic points gives adverse effect on the FERET face recognition dataset.

The results suggest that the training data contributes to space discriminability. This is probably due to the fact that each random projection hyperplane is represented as a linear combination of randomly selected training points. As such, variations and completeness of the training data may have significant contributions to the resulting space. The performance loss suffered on the FERET face recognition dataset is probably caused by the skewed data distribution of this particular dataset. Hence adding synthetic points would significantly alter the original data distribution which in turn affects space discriminability. From our empirical observation (while working with RSH), we found that all data points are grouped together when an intrinsic clustering method was applied to the dataset. The very poor performance of RSH on this dataset supports our view.

Table 2. Recognition accuracy (in %) for the texture recognition task on BRODATZ dataset.

	KSVM	KLSH	RSH	ROSE	ROSES
5c	99.3	88.7	96.6	99.3	99.8
5m	85.8	43.6	81.9	90.1	88.4
5v	86.2	82.6	76.9	91.6	88.6
5v2	89.4	52.0	80.9	90.5	92.7
5v3	87.4	73.0	79.1	88.6	91.3
10	81.3	47.0	72.5	86.7	87.0
10v	81.5	48.0	69.3	88.1	88.5
16c	79.6	33.7	65.7	84.1	85.7
16v	73.4	35.5	59.0	77.1	79.8
Average	84.88	56.0	75.8	88.5	89.1

Table 3. Recognition accuracy (in %) for the face recognition task on the ‘b’ subset of the FERET dataset.

	KSVM	KLSH	RSH	ROSE	ROSES
bd	39.0	70.0	13.5	70.5	52.0
bg	58.5	80.5	31.5	80.5	61.5
Average	48.8	75.2	22.5	75.5	56.8

Figure 2. Sensitivity of random projection space discriminability to the number of selected data points for generating the random hyperplanes, as well as the effect of adding synthetic data points for improving space discriminability. The graphs compare the performance of ROSE and ROSES on 5c (top) and 5m (bottom) sets of the BRODATZ texture recognition dataset respectively.

To further highlight the proposed ROSES method we set an experiment on two sequences ‘5c’ and ‘5m’ of the BRODATZ dataset. In this experiment we reduce the number of data required for creating the mapping function step by step. First step we use all the provided training data to construct the random projection space. Then, we progressively discard training data points from a particular class to construct the space. We repeat this process until there is only one class left. Both ‘5c’ and ‘5m’ have a total of 5 classes where each class has 5 samples for training. We ran the experiment on every single combination for each case

(e.g. when two classes are excluded, there are 10 combinations) and present the average accuracy.

As shown in Fig. 2, there is a significant performance difference between the ROSE and ROSES methods, highlighting the importance of the training data generating the random projection hyperplanes. This performance difference is more pronounced when more classes are excluded from the training data. We note that this training set is different from the training set to train the classifier. Although we exclude some classes in the training set for constructing the random projection space, we still use all the provided training data to train the classifier.

4.2. Comparison with Recent Methods

Table 4 shows that on the FERET face recognition dataset the proposed ROSE method obtains considerably better results than several recent methods: log-Euclidean sparse representation (logE-SR) [13, 37], Tensor Sparse Coding (TSC) [32], Locality Preserving Projection (RLPP) [17], and Relational Divergence Classification (RDC) [2].

Table 5 contrasts the performance of the ROSES method (ROSE augmented with synthetic data) on the BRODATZ texture recognition task against the above methods. We note that in this case the use of synthetic data is necessary in order to achieve improved performance. On average, ROSES achieves higher performance than the other methods, with top performance obtained in 7 out of 9 tests.

Finally, we compared the ROSES method with several state-of-the-art algorithms for person re-identification on the ETHZ dataset: Histogram Plus Epitome (HPE) [4], Symmetry-Driven Accumulation of Local Features (SDALF) [11], RLPP [17] and RDC [2]. The performance of TSC [32] was not evaluated due to the method’s high computational demands: it would take approximately 200 hours to process the ETHZ dataset. We do not report the results for LogE-SR due to its low performance on the other two datasets. The results shown in Table 6 indicate that the proposed ROSES method obtains better performance. As in the previous experiment, the use of synthetic data is necessary to obtain improved performance.

Table 4. Recognition accuracy (in %) for the face recognition task using log-Euclidean sparse representation (logE-SR) [13, 37], Tensor Sparse Coding (TSC) [32], Riemannian Locality Preserving Projection (RLPP) [17], Relational Divergence Classification (RDC) [2], and the proposed ROSE method.

	LogE-SR	TSC	RLPP	RDC	ROSE
bd	35.0	36.0	47.0	59.0	70.0
bg	47.0	45.0	58.0	71.0	80.5
Average	41.0	40.5	52.5	65.0	75.2

Table 5. Performance on the Brodatz texture dataset [25] for LogE-SR [13, 37], Tensor Sparse Coding (TSC) [32], Riemannian Locality Preserving Projection (RLPP) [17], Relational Divergence Classification (RDC) [2], and the proposed ROSES method.

	LogE-SR	TSC	RLP	RDC	ROSES
5c	89.0	99.7	99.2	98.2	99.8
5m	53.5	72.5	86.2	88.0	88.4
5v	73.5	86.3	86.4	87.0	88.6
5v2	70.8	86.1	90.0	89.0	92.7
5v3	63.6	83.1	89.7	87.0	91.3
10	60.6	81.3	84.7	84.0	87.0
10v	63.4	67.9	83.0	86.0	88.5
16c	67.1	75.1	82.0	88.0	85.7
16v	55.4	66.6	74.0	81.0	79.8
Average	66.3	79.8	86.1	87.6	89.1

Table 6. Recognition accuracy (in %) for the person re-identification task on Seq.1 and Seq.2 of the ETHZ dataset. HPE: Histogram Plus Epitome [4]; SDALF: Symmetry-Driven Accumulation of Local Features [11]; RLPP: Riemannian Locality Preserving Projection [17]; RDC: Relational Divergence Classification [2].

	HPE	SDALF	RLPP	RDC	ROSES
Seq.1	79.5	84.1	88.2	88.7	92.5
Seq.2	85.0	84.0	89.8	89.8	94.0
Average	82.2	84.0	89.0	89.2	93.2

5. Main Findings and Future Directions

The key advantage of representing images in forms of non-singular covariance matrix groups is that superior performance can be achieved when the underlying structure of the group is considered. It has been shown that when endowed with the Affine Invariant Riemannian Metric (AIRM), the matrices form a connected, smooth and differentiable Riemannian manifold. Working directly on the manifold space via AIRM poses many computational challenges. Typical ways of addressing this issue include embedding the manifolds to tangent spaces, and embedding into Reproducing Kernel Hilbert Spaces (RKHS). Embedding the manifolds to tangent spaces considerably simplifies further analysis, at the cost of disregarding some of the manifold structure. Embedding via RKHS can better preserve the manifold structure, but adds the burden of extending existing Euclidean-based learning algorithms into RKHS.

In this work, we have presented a novel solution which embeds the data points into a random projection space by first generating random hyperplanes in RKHS and then projecting the data in RKHS into the random projection space. We presented a study of space discriminability for various computer vision classification tasks and found that the space has superior discriminative power to the typical approaches

outlined above. In addition, we found that the space discriminative power depends on the completeness of training data generating the random hyperplanes. To address this issue, we proposed to augment training data with synthetic data.

Experiments on face recognition, person re-identification and texture classification show that the proposed method (combined with a linear SVM) outperforms state-of-the-art approaches such as Tensor Sparse Coding, Histogram Plus Epitome, Riemannian Locality Preserving Projection and Relational Divergence Classification. To our knowledge this is the first time random projection space has been applied to solve classification tasks in manifold space. We envision that the proposed method can be used to bring superior discriminative power of manifold spaces to more general vision tasks, such as object tracking.

Appendix

Here we provide more details to support Proposition 3.3: we show that for SPD matrices, $\log(\mathbf{X}^c) = c \times \log(\mathbf{X})$. For this proof we let c be a discrete number, however we note that it can be extended to continuous c .

$$\log(\mathbf{X}^c) = \log(\underbrace{\mathbf{X} \times \mathbf{X} \times \dots \times \mathbf{X}}_{c \text{ instances of } \mathbf{X}}) \quad (17)$$

Replacing $\mathbf{X} \in \text{Sym}_+^d$ with its singular value decomposition (SVD) as $\mathbf{X} = \mathbf{U}\mathbf{V}\mathbf{U}^\top$, the above equation becomes:

$$\log(\mathbf{X}^c) = \log(\mathbf{U}\mathbf{V}\mathbf{U}^\top \times \dots \times \mathbf{U}\mathbf{V}\mathbf{U}^\top) \quad (18)$$

As $\mathbf{X} \in \text{Sym}_+^d$, the eigenvalue matrix $\mathbf{U}$ is orthonormal, and hence $\mathbf{U}^\top \mathbf{U} = \mathbf{I}$. As such, the following equation is valid:

$$\log(\mathbf{X}^c) = \log(\mathbf{U}\mathbf{V}^c\mathbf{U}^\top) \quad (19)$$

Similarly, as $\mathbf{X} \in \text{Sym}_+^d$, $\log(\mathbf{X}^c) = \mathbf{U} \log(\mathbf{V}^c) \mathbf{U}^\top$, where $\log(\mathbf{V}^c)$ is the diagonal matrix of the eigenvalue logarithm [35]. Hence we have:

$$\begin{aligned} \log(\mathbf{X}^c) &= \mathbf{U} \log(\mathbf{V}^c) \mathbf{U}^\top \\ &= c \mathbf{U} \log(\mathbf{V}) \mathbf{U}^\top \\ &= c \times \log(\mathbf{X}) \end{aligned}$$

Acknowledgements

This research was partly funded by Sullivan Nicolaides Pathology (Australia), and the Australian Research Council Linkage Projects Grant LP130100230.

NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

References

- [1] D. Achlioptas. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *Journal of Computer and System Sciences*, 66(4):671–687, 2003.
- [2] A. Alavi, M. T. Harandi, and C. Sanderson. Relational divergence based classification on Riemannian manifolds. In *IEEE Workshop on Applications of Computer Vision (WACV)*, pages 111–116, 2013.
- [3] C. Anoop, M. Vassilios, and P. Nikolaos. Dirichlet process mixture models on symmetric positive definite matrices for appearance clustering in video surveillance applications. In *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 3417–3424, 2011.
- [4] L. Bazzani, M. Cristani, A. Perina, M. Farenzena, and V. Murino. Multiple-shot person re-identification by HPE signature. In *Int. Conf. Pattern Recognition (ICPR)*, pages 1413–1416, 2010.
- [5] M. S. Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings of ACM Symposium on Theory of Computing*, pages 380–388, 2002.
- [6] R. Chaudhry and Y. Ivanov. Fast approximate nearest neighbor methods for non-Euclidean manifolds with applications to human activity analysis in videos. In *European Conference on Computer Vision*, pages 735–748, 2010.
- [7] A. Cherian, S. Sra, A. Banerjee, and N. Papanikolopoulos. Efficient similarity search for covariance matrices via the Jensen-Bregman LogDet divergence. In *Int. Conf. Computer Vision (ICCV)*, pages 2399–2406, 2011.
- [8] M. Datar, N. Immorlica, P. Indyk, and V. S. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of Symposium on Computational Geometry*, pages 253–262. ACM, 2004.
- [9] A. Ess, B. Leibe, and L. Van Gool. Depth and appearance for mobile scene analysis. In *Int. Conf. Computer Vision (ICCV)*, pages 1–8, 2007.
- [10] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008.
- [11] M. Farenzena, L. Bazzani, A. Perina, V. Murino, and M. Cristani. Person re-identification by symmetry-driven accumulation of local features. *IEEE Conf. Computer Vision and Pattern Recognition*, pages 2360–2367, 2010.
- [12] Y. Gong, S. Lazebnik, A. Gordo, and F. Perronnin. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 35(12):2916–2929, 2013.
- [13] K. Guo, P. Ishwar, and J. Konrad. Action recognition using sparse representation on covariance manifolds of optical flow. In *IEEE Conf. Advanced Video and Signal Based Surveillance (AVSS)*, pages 188–195, 2010.
- [14] J. Hamm and D. D. Lee. Extended Grassmann kernels for subspace-based learning. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems (NIPS)*, pages 601–608, 2009.
- [15] M. Harandi, C. Sanderson, R. Hartley, and B. C. Lovell. Sparse coding and dictionary learning for symmetric positive definite matrices: a kernel approach. In *European Conference on Computer Vision (ECCV)*, volume 7573 of *Lecture Notes in Computer Science*, pages 216–229. Springer, 2012.
- [16] M. Harandi, C. Sanderson, C. Shen, and B. C. Lovell. Dictionary learning and sparse coding on Grassmann manifolds: An extrinsic solution. In *IEEE Int. Conf. Computer Vision (ICCV)*, 2013.
- [17] M. Harandi, C. Sanderson, A. Wiliem, and B. C. Lovell. Kernel analysis over Riemannian manifolds for visual recognition of actions, pedestrians and textures. *IEEE Workshop on the Applications of Computer Vision (WACV)*, pages 433–439, 2012.
- [18] P. Jain, B. Kulis, and K. Grauman. Fast image search for learned metrics. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [19] B. Kulis and K. Grauman. Kernelized locality-sensitive hashing for scalable image search. In *International Conference on Computer Vision*, pages 2130–2137, 2009.
- [20] S. Lang. *Fundamentals of differential geometry*, volume 160. Springer Verlag, 1999.
- [21] Y. Lui. Tangent bundles on special manifolds for action recognition. *IEEE Trans. Circuits and Systems for Video Technology*, 22(6):930–942, 2011.
- [22] X. Pennec. Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, 2006.
- [23] P. Phillips, H. Moon, S. Rizvi, and P. Rauss. The FERET evaluation methodology for face-recognition algorithms. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22(10):1090–1104, 2000.
- [24] F. Porikli, O. Tuzel, and P. Meer. Covariance tracking using model update based on Lie algebra. In *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 728–735, 2006.
- [25] T. Randen and J. Husoy. Filtering for texture classification: A comparative study. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(4):291–310, 1999.
- [26] J. A. Rice. *Mathematical statistics and data analysis*. Cengage Learning, 2007.
- [27] A. Sanin, C. Sanderson, M. T. Harandi, and B. C. Lovell. K-tangent spaces on Riemannian manifolds for improved pedestrian detection. In *IEEE Int. Conf. Image Processing (ICIP)*, pages 473–476, 2012.
- [28] A. Sanin, C. Sanderson, M. T. Harandi, and B. C. Lovell. Spatio-temporal covariance descriptors for action and gesture recognition. In *IEEE Workshop on Applications of Computer Vision (WACV)*, pages 103–110, 2013.
- [29] W. R. Schwartz and L. S. Davis. Learning discriminative appearance-based models using partial least squares. In *SIBGRAPI*, pages 322–329, 2009.
- [30] J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- [31] Q. Shi, C. Shen, R. Hill, and A. v. d. Hengel. Is margin preserved after random projection? *ICML*, 2012.
- [32] R. Sivalingham, D. Boley, V. Morellas, and N. Papanikolopoulos. Tensor sparse coding for region covariances. *European Conference on Computer Vision (ECCV)*, pages 722–735, 2010.
- [33] S. Sra. Positive definite matrices and the symmetric Stein divergence. *Preprint: [arXiv:1110.1773]*, 2012.
- [34] S. Sra and A. Cherian. Generalized dictionary learning for symmetric positive definite matrices with application to nearest neighbor retrieval. In *Machine Learning and Knowledge Discovery in Databases*, pages 318–332. Springer, 2011.
- [35] O. Tuzel, F. Porikli, and P. Meer. Pedestrian detection via classification on Riemannian manifolds. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 30(10):1713–1727, 2008.
- [36] A. Veeraraghavan, A. Roy-Chowdhury, and R. Chellappa. Matching shape sequences in video with applications in human movement analysis. *Pattern Analysis and Machine Intelligence*, 27(12):1896–1909, 2005.
- [37] C. Yuan, W. Hu, X. Li, S. Maybank, and G. Luo. Human action recognition under log-Euclidean Riemannian metric. In *Asian Conference on Computer Vision (ACCV)*, volume 5994 of *Lecture Notes in Computer Science*, pages 343–353, 2010.