Sparse Coding and Dictionary Learning for Symmetric Positive Definite Matrices: A Kernel Approach

Mehrtash T. Harandi, Conrad Sanderson, Richard Hartley, Brian C. Lovell

NICTA, PO Box 6020, St Lucia, QLD 4067, Australia *
University of Queensland, School of ITEE, QLD 4072, Australia
NICTA, Locked Bag 8001, Canberra, ACT 2601, Australia
Australian National University, Canberra, ACT 0200, Australia

Abstract. Recent advances suggest that a wide range of computer vision problems can be addressed more appropriately by considering non-Euclidean geometry. This paper tackles the problem of sparse coding and dictionary learning in the space of symmetric positive definite matrices, which form a Riemannian manifold. With the aid of the recently introduced Stein kernel (related to a symmetric version of Bregman matrix divergence), we propose to perform sparse coding by embedding Riemannian manifolds into reproducing kernel Hilbert spaces. This leads to a convex and kernel version of the Lasso problem, which can be solved efficiently. We furthermore propose an algorithm for learning a Riemannian dictionary (used for sparse coding), closely tied to the Stein kernel. Experiments on several classification tasks (face recognition, texture classification, person re-identification) show that the proposed sparse coding approach achieves notable improvements in discrimination accuracy, in comparison to state-of-the-art methods such as tensor sparse coding, Riemannian locality preserving projection, and symmetry-driven accumulation of local features.

1 Introduction

Sparse representation (SR), the linear decomposition of a signal using a few atoms of a dictionary, has led to notable results for various image processing and computer vision tasks [1,2]. While significant steps have been taken towards expanding the theory of SR, such representations in non-Euclidean spaces have received comparatively little attention. This paper tackles the problem of sparse coding within the space of symmetric positive definite (SPD) matrices.

SPD matrices are fundamental building blocks in computer vision and machine learning. A notable example is the covariance descriptor [3], which offer a compact way of describing regions/cuboids in images/videos and fusion of multiple features. Covariance descriptors have been exploited in several applications, such as diffusion tensor imaging [4], action recognition [5,6,7], pedestrian detection [3], face recognition [8,9], texture classification [9,10], and tracking [11].

SPD matrices form a cone of zero curvature and can be analysed using the geometry of Euclidean space. However, several studies have shown that a Riemannian structure

* Published in: **Lecture Notes in Computer Science (LNCS)**, Vol. 7573, pp. 216–229, 2012.
http://dx.doi.org/10.1007/978-3-642-33709-3_16

of negative curvature is more suitable for analysing SPD matrices [4,12]. More specifically, Pennec et al. [4] introduced the Affine Invariant Riemannian Metric (AIRM) and showed that the induced Riemannian structure is invariant to inversion and similarity transforms. The AIRM is perhaps the most widely used similarity measure for SPD matrices. Nevertheless, efficiently and accurately handling the Riemannian structure is non-trivial as basic computations on Riemannian manifolds (such as similarities and distances) involve non-linear operators. This not only hinders the development of optimisation algorithms but also incurs a significant numerical burden.

To address the above drawbacks, in this paper we propose to perform the sparse coding of SPD matrices by embedding Riemannian manifolds into reproducing kernel Hilbert spaces (RKHS) [13]. This is in contrast to directly embedding into Euclidean spaces [7,6,14].

Related Work. Sra et al. [14] used the cone of SPD matrices and the Frobenius norm as a measure of similarity between SPD matrices. While this results in a regularised non-negative least-squares approach, it does not consider the Riemannian geometry induced by AIRM.

Guo et al. [6] and Yuan et al. [7] separately proposed to solve sparse representation by a log-Euclidean approach, where a Riemannian problem is converted to an Euclidean one by embedding manifolds into tangent spaces. While log-Euclidean approaches benefit from simplicity, the true geometry of the manifold is not taken into account. More specifically, on a tangent space only distances to the pole of space are true geodesic distances. As such, the pairwise distances between arbitrary points on the tangent space do not represent the structure of the manifold.

Sivalingam et al. [9] used Burg divergence [15] as a metric and reformulated the Riemannian¹ SR problem as a determinant maximisation problem. This has the advantage of avoiding the explicit manifold embedding, as well as resulting in a convex MAXDET problem [16] that can be solved by interior point methods. However, there are two downsides: the solution is computationally very expensive, and the relations between Burg divergence and the geometry of Riemannian manifolds were not well established.

Contributions. With the aid of the recently introduced Stein kernel [17], which is related to AIRM via a tight bound, we propose a Riemannian sparse solver by embedding Riemannian manifolds into RKHS. We show that the embedding leads to a convex and kernelised version of the Lasso problem [1], which can be solved efficiently. We furthermore propose a sparsity-maximising algorithm for dictionary learning within the space of SPD matrices, closely tied to the Stein kernel. Lastly, we show that the proposed sparse coding approach obtains superior performance on several visual classification tasks (face recognition, texture classification, person re-identification), in comparison to several state-of-the-art methods: tensor sparse coding [9], log-Euclidean sparse representation [6,7], Gabor feature based sparse representation [18], and Riemannian locality preserving projection [10].

We continue the paper as follows. Section 2 begins with an overview of Bregman divergence and the Stein kernel. Section 3 describes the proposed kernel solution of Riemannian sparse coding, followed by Section 4, which covers the problem of dictio-

¹ We loosely use ‘Riemannian’ to refer to the Riemannian manifold formed by SPD matrices.

nary learning on Riemannian manifolds. In Section 5 we compare the performance of the proposed method with previous approaches on several visual classification tasks. The main findings and possible future directions are summarised in Section 6.

2 Background

In this section we first overview the properties of Bregman matrix divergences, including a special case known as the symmetric Stein divergence. This leads to the Stein kernel, which can be used to embed Riemannian manifolds into RKHS.

2.1 Bregman Matrix Divergences

The *Bregman matrix divergence* for two symmetric matrices $\mathbf{X}$ and $\mathbf{Y}$ is defined as [15]:

$$D_\zeta(\mathbf{X}, \mathbf{Y}) \triangleq \zeta(\mathbf{X}) - \zeta(\mathbf{Y}) - \langle \nabla_\zeta(\mathbf{Y}), \mathbf{X} - \mathbf{Y} \rangle \quad (1)$$

where $\langle \mathbf{A}, \mathbf{B} \rangle = \text{Tr}(\mathbf{A}^T \mathbf{B})$ and ζ is a real valued, strictly convex function on symmetric matrices. Bregman divergences are non-negative, definite, and in general asymmetric. Among the several ways to symmetrise them, the *Jensen-Shannon* symmetrisation is often used [15]:

$$D_\zeta^{JS}(\mathbf{X}, \mathbf{Y}) \triangleq \frac{1}{2} D_\zeta\left(\mathbf{X}, \frac{\mathbf{X} + \mathbf{Y}}{2}\right) + \frac{1}{2} D_\zeta\left(\mathbf{Y}, \frac{\mathbf{X} + \mathbf{Y}}{2}\right) \quad (2)$$

If $\zeta = -\log(\det(\mathbf{X}))$, then the symmetric Stein divergence is obtained from (2):

$$S(\mathbf{X}, \mathbf{Y}) \triangleq \log\left(\det\left(\frac{\mathbf{X} + \mathbf{Y}}{2}\right)\right) - \frac{1}{2} \log(\det(\mathbf{X}\mathbf{Y})), \text{ for } \mathbf{X}, \mathbf{Y} \succ 0 \quad (3)$$

The space induced by AIRM on symmetric positive definite matrices of dimension d is a Riemannian manifold Sym_+^d of negative curvature. For two points $\mathbf{X}, \mathbf{Y} \in Sym_+^d$, the AIRM is defined as $d_g^2 = \|\log_{\mathbf{X}}(\mathbf{Y})\|_{\mathbf{X}}^2 = \text{Tr}\left\{\log^2\left(\mathbf{X}^{-\frac{1}{2}}\mathbf{Y}\mathbf{X}^{-\frac{1}{2}}\right)\right\}$, where $\log_{\mathbf{X}}(\mathbf{Y}) = \mathbf{X}^{\frac{1}{2}} \log\left(\mathbf{X}^{-\frac{1}{2}}\mathbf{Y}\mathbf{X}^{-\frac{1}{2}}\right) \mathbf{X}^{\frac{1}{2}}$. The symmetric Stein divergence and Riemannian metric over Sym_+^d manifolds are related in several aspects. Two important properties are summarised below.

Property 1. Let $\mathbf{X}, \mathbf{Y} \in Sym_+^d$, and $\delta_T(\mathbf{X}, \mathbf{Y}) = \max_{1 \leq i \leq d} \{|\log \Lambda(\mathbf{X}\mathbf{Y}^{-1})|\}$ be the Thompson metric [19] with $\Lambda(\mathbf{X}\mathbf{Y}^{-1})$ representing the vector of eigenvalues of $\mathbf{X}\mathbf{Y}^{-1}$. The following sandwiching inequality between the symmetric Stein divergence and Riemannian metric exists [17]:

$$S(\mathbf{X}, \mathbf{Y}) \leq \frac{1}{8} d_g^2(\mathbf{X}, \mathbf{Y}) \leq \frac{1}{4} \delta_T(\mathbf{X}, \mathbf{Y}) (S(\mathbf{X}, \mathbf{Y}) + d \log d) \quad (4)$$

Property 2. The curve $\gamma(p) \triangleq \mathbf{X}^{\frac{1}{2}} \left(\mathbf{X}^{-\frac{1}{2}}\mathbf{Y}\mathbf{X}^{-\frac{1}{2}}\right)^p \mathbf{X}^{\frac{1}{2}}$ parameterises the unique geodesic between the SPD matrices $\mathbf{X}$ and $\mathbf{Y}$. On this curve the Riemannian geodesic distance satisfies $d_g(\mathbf{X}, \gamma(p)) = p d_g(\mathbf{X}, \mathbf{Y})$; $p \in [0, 1]$ [12]. The symmetric Stein divergence satisfies a similar but slightly weaker result, $S(\mathbf{X}, \gamma(p)) \leq p S(\mathbf{X}, \mathbf{Y})$.

The first property establishes a bound between the geodesic distance and Stein divergence, providing motivation for addressing Riemannian problems via the divergence. The second property reinforces the motivation by explaining that the behaviour of Stein divergences along geodesic curves is similar to true Riemannian geometry.

2.2 Stein Kernel

Definition 1. Let $\Omega = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N\}$ be a non-empty set on Riemannian manifold Sym_+^d . A function $\varphi : \Omega \times \Omega \rightarrow \mathbb{R}_+$ is a Riemannian kernel if φ is symmetric for all $\mathbf{X}, \mathbf{Y} \in \Omega$, ie., $\varphi(\mathbf{X}, \mathbf{Y}) = \varphi(\mathbf{Y}, \mathbf{X})$, and the following inequality is satisfied for all $a_i \in \mathbb{R}$:

$$\sum_{i,j} a_i a_j \varphi(\mathbf{X}_i, \mathbf{X}_j) \geq 0$$

Under a mild condition (explained afterwards), the following function forms a Riemannian kernel [17]:

$$k(\mathbf{X}, \mathbf{Y}) = e^{-\sigma S(\mathbf{X}, \mathbf{Y})} = 2^{d\sigma} \frac{\sqrt{\det(\mathbf{X})^\sigma \det(\mathbf{Y})^\sigma}}{\det(\mathbf{X} + \mathbf{Y})^\sigma} \quad (5)$$

We shall refer to this kernel as the *Stein kernel* from here on. The following theorem states the condition under which Stein kernel is positive definite.

Theorem 1. Let $\Omega = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N\}; \mathbf{X}_i \in Sym_+^d$ be a set of Riemannian points. The $N \times N$ matrix $K_\sigma = [k_\sigma(i, j)]; 1 \leq i, j \leq N$, with $k_\sigma(i, j) = k(\mathbf{X}_i, \mathbf{X}_j)$, defined in (5), is positive definite iff:

$$\sigma \in \left\{ \frac{1}{2}, \frac{2}{2}, \dots, \frac{d-1}{2} \right\} \cup \left\{ \tau \in \mathbb{R} : \tau > \frac{1}{2}(d-1) \right\} \quad (6)$$

Interested readers can follow the proof in [17]. For values of σ outside of the above set, it is possible to convert a pseudo kernel into a true kernel, as discussed for example in [20]. The determinant of an $d \times d$ SPD matrix can be efficiently computed by Cholesky decomposition in $O(d^3)$. As such, the complexity of computing Stein kernel is $O(3d^3 + 3\sigma)$.

3 Kernel Sparse Coding

Sparse coding on Riemannian manifolds in general means that a given query point on a manifold can be expressed as a sparse “combination” of dictionary elements. Our idea here is to embed the manifold into RKHS and replace the idea of “combination” on manifolds with the general concept of linear combination in Hilbert spaces. More specifically, given a Riemannian dictionary $\mathbb{D} = \{\mathbf{D}_1, \mathbf{D}_2, \dots, \mathbf{D}_N\}; \mathbf{D}_i \in Sym_+^d$, and an embedding function $\phi : Sym_+^d \rightarrow \mathbb{H}$, for a Riemannian point $\mathbf{X}$ we seek for a sparse vector $\mathbf{v} \in \mathbb{R}^N$ such that $\phi(\mathbf{X})$ admits the sparse representation $\mathbf{v}$ over $\{\phi(\mathbf{D}_1), \phi(\mathbf{D}_2), \dots, \phi(\mathbf{D}_N)\}$. In other words, we are interested in solving the following kernelised version of the Lasso problem [1]:

$$\min_{\mathbf{v} \in \mathbb{R}^N} \left(\left\| \phi(\mathbf{X}) - \sum_{i=1}^N v_i \phi(\mathbf{D}_i) \right\|^2 + \lambda \|\mathbf{v}\|_1 \right) \quad (7)$$

The first term in (7) can be expanded as:

$$\begin{aligned}
& \left\| \phi(\mathbf{X}) - \sum_{i=1}^N v_i \phi(\mathbf{D}_i) \right\|^2 \\
&= k(\mathbf{X}, \mathbf{X}) - 2 \sum_{i=1}^N v_i k(\mathbf{X}, \mathbf{D}_i) + \sum_{i=1}^N \sum_{j=1}^N v_i v_j k(\mathbf{D}_j, \mathbf{D}_i) \\
&= k(\mathbf{X}, \mathbf{X}) - 2\mathbf{v}^T \mathcal{K}(\mathbf{X}, \mathbb{D}) + \mathbf{v}^T \mathbb{K}(\mathbb{D}, \mathbb{D}) \mathbf{v}
\end{aligned} \tag{8}$$

where $\mathcal{K} = [a_i]_{N \times 1}$; $a_i = k(\mathbf{X}, \mathbf{D}_i)$ and $\mathbb{K} = [a_{ij}]_{N \times N}$; $a_{ij} = k(\mathbf{D}_i, \mathbf{D}_j)$. This reveals that the optimisation problem in (7) is convex and similar to its counterpart in Euclidean space, except for the definition of $\mathcal{K}$ and $\mathbb{K}$. Consequently, greedy or relaxation solutions can be adapted to obtain the sparse codes [1]. To solve problem (7) we used CVX [21], a package for specifying and solving convex programs².

3.1 Classification Using Sparse Codes

There are two main approaches for classification based on the obtained sparse codes (vectors) for a given query sample: (i) directly, and (ii) indirectly, with the aid of an Euclidean-based classifier. We elucidate the two approaches below.

(i) If the atoms in sparse dictionary $\mathbb{D}$ have associated class labels (ie. each atom in the dictionary is a training sample), the sparse codes can be directly used for classification. This approach is applicable only to closed-set identification tasks. Let $\mathbf{v}_i = [v_{i,1}\delta(l(1)-i), v_{i,2}\delta(l(2)-i), \dots, v_{i,N}\delta(l(N)-i)]^T$ be the class-specific sparse codes, where $l(j)$ is the class label of atom $\mathbf{D}_j$ and $\delta(x)$ is the discrete Dirac function [22]. An efficient way of using class-specific sparse codes is through computing residual errors [2]. In this case, the residual error of query sample $\mathbf{X}$ for class i is defined as:

$$\varepsilon_i(\mathbf{X}) = \left\| \phi(\mathbf{X}) - \sum_{j=1}^N v_j \phi(\mathbf{D}_j) \delta(l(j) - i) \right\|^2 \tag{9}$$

which can be computed via the use of a Riemannian kernel in a similar manner to (8). The class with the minimum residual error is deemed to represent the query. Alternatively, the similarity between query sample $\mathbf{X}$ to class i can be defined as $S_i(\mathbf{X}) = h_i(\mathbf{v})$. The function $h_i(\mathbf{v})$ can be linear like $\sum_{j=1}^N v_j \delta(l(j) - i)$ or even non-linear like $\max(v_j \delta(l(j) - i))$.

(ii) If the atoms in the sparse dictionary $\mathbb{D}$ are not labelled (eg. $\mathbb{D}$ is a generic dictionary not tied to any particular class [23]), the generated sparse codes (vectors) for both training and query data can be fed to Euclidean-based classifiers, such as support vector machines [22]. The sparse code is hence interpreted as a feature vector, which in essence means that a classification problem on a Riemannian manifold is converted to an Euclidean classification problem. This approach is applicable to both closed-set and open-set classification tasks.

4 Learning Riemannian Dictionaries

If the indirect classification of sparse codes is required (as elucidated in the preceding section) a Riemannian dictionary is first required. Given a set of Riemannian points $\Omega =$

² The SPAMS package can also be used: <http://spams-devel.gforge.inria.fr/>

$\{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_m\}; \mathbf{X}_i \in \text{Sym}_+^d$, learning a dictionary $\mathbb{D} = \{\mathbf{D}_1, \mathbf{D}_2, \dots, \mathbf{D}_N\}; \mathbf{D}_i \in \text{Sym}_+^d$ can be formulated as jointly minimising the energy function

$$J = \left\{ \sum_{j=1}^m \left(\left\| \phi(\mathbf{X}_j) - \sum_{i=1}^N v_{j,i} \phi(\mathbf{D}_i) \right\|^2 + \lambda \|\mathbf{v}_j\|_1 \right) \right\} \quad (10)$$

over the dictionary and the sparse codes $\mathbb{V} = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m\}; \mathbf{v}_i \in \mathbb{R}^N$, ie., $\min_{\mathbb{D}, \mathbb{V}}(J)$.

Among the various solutions to the problem of dictionary learning in Euclidean spaces, iterative methods like K-SVD have received much attention [1]. Borrowing the idea from Euclidean spaces, we propose to minimise the energy in (10) iteratively. After initialising the dictionary $\mathbb{D}$, for example by Riemannian clustering using the Karcher mean [4], we iterate between a sparse coding step and a dictionary update step. In the sparse coding step, $\mathbb{D}$ is fixed and $\mathbb{V}$ is computed. In the dictionary update step, $\mathbb{V}$ is fixed while $\mathbb{D}$ is updated, with each dictionary atom updated independently.

The derivative of (10) with respect to $\mathbf{D}_r$, while $\mathbb{V}$ and other atoms are fixed, is:

$$\frac{\partial J}{\partial \mathbf{D}_r} = \sum_{j=1}^m \left(-2v_{j,r} \frac{\partial k(\mathbf{X}_j, \mathbf{D}_r)}{\partial \mathbf{D}_r} + \sum_{i=1}^N v_{j,i} v_{j,r} \frac{\partial k(\mathbf{D}_i, \mathbf{D}_r)}{\partial \mathbf{D}_r} \right) \quad (11)$$

As $\nabla_{\mathbf{X}} S(\mathbf{X}, \mathbf{Y}) = (\mathbf{X} + \mathbf{Y})^{-1} - \frac{1}{2} \mathbf{X}^{-1}$, (11) can be further simplified to:

$$\begin{aligned} \frac{\partial J}{\partial \mathbf{D}_r} = & 2\beta \sum_{j=1}^m v_{j,r} k(\mathbf{X}_j, \mathbf{D}_r) \left((\mathbf{X}_j + \mathbf{D}_r)^{-1} - \frac{1}{2} \mathbf{D}_r^{-1} \right) \\ & - \beta \sum_{j=1}^m \sum_{i=1}^N v_{j,i} v_{j,r} k(\mathbf{D}_i, \mathbf{D}_r) \left((\mathbf{D}_i + \mathbf{D}_r)^{-1} - \frac{1}{2} \mathbf{D}_r^{-1} \right) \end{aligned} \quad (12)$$

Since (12) contains linear and non-linear terms of $\mathbf{D}_r$ (eg. inverse and kernel terms), a closed-form solution for computing its root, ie., $\mathbf{D}_r$, cannot be sought. As such, we propose an alternative solution by exploiting previous values of $k(\cdot, \mathbf{D}_r)$ and $(\mathbf{D}_i - \mathbf{D}_r)^{-1}$ in the updating step. More specifically, rearranging (12) and estimating $k(\cdot, \mathbf{D}_r)$ as well as $(\mathbf{D}_i - \mathbf{D}_r)^{-1}$ by their previous values, atom $\mathbf{D}_r$ at iteration $t+1$ is updated using:

$$\mathbf{D}_r^{(t+1)} = \frac{2}{\sum_{j=1}^m v_{j,r} (\mathbf{v}_j^T \mathbf{k}(\mathbb{D}, \mathbf{D}_r) - 2k(\mathbf{X}_j, \mathbf{D}_r))} \left(F^{(t)}(r) + G^{(t)}(r) \right)^{-1} \quad (13)$$

where

$$F^{(t)}(r) = \sum_{j=1}^m 2v_{j,r} k^{(t)}(\mathbf{X}_j, \mathbf{D}_r) \left(\mathbf{X}_j + \mathbf{D}_r^{(t)} \right)^{-1} \quad (14)$$

$$G^{(t)}(r) = \sum_{j=1}^m \sum_{i=1}^N v_{j,r} v_{j,i} k^{(t)}(\mathbf{D}_i, \mathbf{D}_r) \left(\mathbf{D}_i^{(t)} + \mathbf{D}_r^{(t)} \right)^{-1} \quad (15)$$

To avoid the degenerative case (due to numerical inconsistency), atoms are normalised by their second norms at the end of each iteration. Algorithm 1 assembles all the above details into pseudo-code for dictionary learning.

5 Experiments

Two sets of experiments³ are presented in this section. In the first set, we evaluate the performance of the proposed Riemannian SR (RSR) method (as described in Section 3)

³ Matlab/Octave source code is available at <http://itee.uq.edu.au/~uqmhara1>

Algorithm 1 : Dictionary learning over Sym_+^d using the Stein kernel

Input:

- training set $\mathbb{X} = \{\mathbf{X}_i\}_{i=1}^m$ from the underlying Riemannian manifold, where each $\mathbf{X}_i \in Sym_+^d$ is a SPD matrix
- Stein kernel function $k(\mathbf{X}, \mathbf{Y})$, as defined in Eqn. (5)
- $nIter$, the number of iterations

Output:

- Riemannian dictionary $\mathbb{D} = \{\mathbf{D}_i\}_{i=1}^N$
- 1: Initialise the dictionary $\mathbb{D}^{(1)} = \{\mathbf{D}_i^{(1)}\}_{i=1}^N$ by selecting N samples from $\mathbb{X}$ randomly, or by clustering $\mathbb{X}$ on the manifold [24]
 - 2: **for** $t = 1 \rightarrow nIter$ **do**
 - 3: Compute $k^{(t)}(\mathbf{X}_i, \mathbf{D}_j^{(t)})$, $1 \leq i \leq m$, $1 \leq j \leq N$
 - 4: Compute $k^{(t)}(\mathbf{D}_i^{(t)}, \mathbf{D}_j^{(t)})$, $1 \leq i, j \leq N$
 - 5: Solve $\min_{\mathbf{v}_j \in \mathbb{R}^N} -2\mathbf{v}_j^T \mathcal{K}(\mathbf{X}_j, \mathbb{D}^{(t)}) + \mathbf{v}^T \mathbb{K}(\mathbb{D}^{(t)}, \mathbb{D}^{(t)}) \mathbf{v} + \lambda \|\mathbf{v}_j\|_1$, $\forall j, \mathbf{X}_j \in \mathbb{X}$
 - 6: **for** $r = 1 \rightarrow N$ **do**
 - 7: Compute $G^{(t)}(r) = \sum_{j=1}^m \sum_{i=1}^N v_{j,r} v_{j,i} k^{(t)}(\mathbf{D}_i^{(t)}, \mathbf{D}_r^{(t)}) (\mathbf{D}_i^{(t)} + \mathbf{D}_r^{(t)})^{-1}$
 - 8: Compute $F^{(t)}(r) = \sum_{j=1}^m 2v_{j,r} k^{(t)}(\mathbf{X}_j, \mathbf{D}_r^{(t)}) (\mathbf{X}_j + \mathbf{D}_r^{(t)})^{-1}$
 - 9: $\mathbf{D}_r^{\text{temp}} \leftarrow \frac{2}{\sum_{j=1}^m v_{j,r} (\mathbf{v}_j^T \mathcal{K}(\mathbb{D}^{(t)}, \mathbf{D}_r^{(t)}) - 2k(\mathbf{X}_j, \mathbf{D}_r^{(t)})} (F^{(t)}(r) + G^{(t)}(r))^{-1}$
 - 10: $\mathbf{D}_r^{(t+1)} \leftarrow \frac{\mathbf{D}_r^{\text{temp}}}{\|\mathbf{D}_r^{\text{temp}}\|_2}$
 - 11: **end for**
 - 12: **end for**
-

without dictionary learning. Each atom in the dictionary is a training sample. This is to contrast RSR to previous state-of-the-art methods on several popular closed-set classification tasks. We use the residual error approach for classification, as described in Eqn. (9).

In the second set, the performance of the RSR method is evaluated in conjunction with dictionaries learned via three methods: random, Riemannian k -means, and the proposed dictionary learning technique (as described in Section 4).

5.1 Riemannian Sparse Representation

Synthetic Data. We first consider a multi-class classification problem over Sym_+^3 using synthetic data. We compared the proposed RSR against Tensor Sparse Coding (TSC) [9] and log-Euclidean Sparse Representation (logE-SR) [6,7]. The data used in the experiments constitutes 512 random samples from 4 classes. Half of the samples were used for training and the rest were used as test data.

To create a Riemannian manifold, samples were generated over a particular tangent space and then mapped back to the manifold using the exponential map [12]. The positions of tangent spaces were chosen randomly and samples in each class obeyed a

		logE-SR [6,7]	TSC [9]	RSR (proposed)
recognition accuracy	easy	68.08 ± 2.5	60.04 ± 6.8	83.05 ± 3.0
	hard	53.67 ± 3.2	50.35 ± 4.9	66.72 ± 2.7
run-time	easy+hard	6 sec	11107 sec	41 sec

Table 1. Average recognition accuracy (in %) and wall-clock time for the synthetic classification tasks using log-Euclidean sparse representation [6,7], tensor sparse coding [9] and the proposed RSR approach. Run time is represented by combining the times for the easy and hard tasks.

normal distribution. By fixing the mean of each class and increasing the class variance we created two classification problems: ‘easy’ and ‘hard’. To draw useful statistics, the data creation process was repeated 100 times.

Table 1 shows the average recognition accuracy and the total running time (in seconds). All algorithms were implemented in Matlab and executed on a 3 GHz Intel CPU. In terms of recognition accuracy, RSR obtains superior performance when compared with previous state-of-the-art approaches. We note that by increasing the class variance, samples from the four classes are intertwined, leading to a decrease in recognition accuracy. The performance of logE-SR is higher than TSC, which might be due to the fact that the generated data can be modelled by Gaussian distribution over tangent space, hence favouring the tangent-based solution.

Focusing on run time, Table 1 suggests that logE-SR has the lowest complexity while TSC has the highest. The proposed RSR method is substantially faster than TSC, while delivering the highest recognition accuracy.

Face Recognition. We used the ‘b’ subset of the FERET dataset [25], which includes 1400 images from 198 subjects. The images were closely cropped around the face and downsampled to 64×64 . Examples are shown in Figure 1.

We performed four tests with various pose angles. Training data was composed of images marked ‘ba’, ‘bj’ and ‘bk’ (ie., frontal faces with expression and illumination variations). Images with ‘bd’, ‘be’, ‘bf’ and ‘bg’ labels (ie., non-frontal faces) were used as test data. For Riemannian-based methods, a 43×43 covariance descriptor described a face image, using the following features:

$$F_{x,y} = [I(x,y), x, y, |G_{0,0}(x,y)|, \dots, |G_{0,7}(x,y)|, |G_{1,0}(x,y)|, \dots, |G_{4,7}(x,y)|]$$

where $I(x,y)$ is the intensity value at position x,y and $G_{u,v}(x,y)$ is the response of a 2D Gabor wavelet [26] centered at x,y with orientation u and scale v :

$$G_{u,v}(x,y) = \frac{k_v^2}{4\pi^2} \sum_{t,s} e^{-\frac{k_v^2}{8\pi^2}((x-s)^2 + (y-t)^2)} \left(e^{ik_v((x-t)\cos(\theta_u) + (y-s)\sin(\theta_u))} - e^{-2\pi^2} \right)$$

with $k_v = \frac{1}{\sqrt{2^{v-1}}}$ and $\theta_u = \frac{\pi u}{8}$.

	PCA-SR [2]	GSR [18]	logE-SR [6,7]	TSC [9]	RSR (proposed)
bg	26.0	79.0	46.5	44.5	86.0
bf	61.0	97.0	91.0	73.5	97.5
be	55.5	93.5	81.0	73.0	96.5
bd	27.5	77.0	34.5	36.0	79.5
average	42.50	86.63	63.25	56.75	89.88

Table 2. Recognition accuracy (in %) for the face recognition task using PCA-SRC [2], Gabor SR (GSR) [18], log-Euclidean sparse representation (logE-SR) [6,7], Tensor Sparse Coding (TSC) [9], and the proposed RSR approach.

Table 2 shows a comparison of RSR against logE-SR [6,7], TSC [9], and two purely Euclidean sparse representations, PCA-SRC [2] and Gabor SR (GSR) [18]. In all cases the proposed RSR method obtains the highest accuracy. Furthermore, the proposed approach significantly outperforms state-of-the-art Euclidean solutions, especially for test images with label ‘bg’.

Texture Classification. We performed a classification task using the Brodatz texture dataset [27]. Examples are shown in Fig. 2. We followed the test protocol devised in [9] and generated nine test scenarios with various number of classes. This includes 5-texture (‘5c’, ‘5m’, ‘5v’, ‘5v2’, ‘5v3’), 10-texture (‘10’, ‘10v’) and 16-texture (‘16c’, ‘16v’) mosaics. To create a Riemannian manifold, each image was first downsampled to 256×256 and then split into 64 regions of size 32×32 . The feature vector for any pixel $I(x, y)$ is $F(x, y) = [I(x, y), |\frac{\partial I}{\partial x}|, |\frac{\partial I}{\partial y}|, |\frac{\partial^2 I}{\partial x^2}|, |\frac{\partial^2 I}{\partial y^2}|]$. Each region is described by a 5×5 covariance descriptor of these features. For each test scenario, five covariance matrices per class were randomly selected as training data and the rest was used for testing. The random selection of training/testing data was repeated 20 times.

Fig. 3 compares the proposed RSR method against logE-SR [6,7] and TSC [9]. The proposed RSR approach obtains the highest recognition accuracy on all test scenarios except for the ‘5c’ test, where it has slightly worse performance than TSC.

Person Re-identification. We used the modified ETHZ dataset [31]. The original ETHZ dataset was captured using a moving camera [28], providing a range of variations in the appearance of people. The dataset is structured into 3 sequences. Sequence 1 con-

Fig. 1. Examples of closely-cropped faces from the FERET ‘b’ subset.

Fig. 2. Examples from the Brodatz texture dataset [27].

Fig. 3. Performance on the Bordatz texture dataset [27] using log-Euclidean sparse representation (logE-SR) [6,7], Tensor Sparse Coding (TSC) [9] and the proposed RSR approach. The black bars indicate standard deviations.

Fig. 4. Examples of pedestrians in the ETHZ dataset [28].

Fig. 5. Performance on Sequences 1 and 2 of the ETHZ dataset (left and right panels, respectively), in terms of Cumulative Matching Characteristic curves. The proposed RSR method is compared with Histogram Plus Epitome (HPE) [29], Symmetry-Driven Accumulation of Local Features (SDALF) [30] and Riemannian Locality Preserving Projection (RLPP) [10].

tains 83 pedestrians (4,857 images), Sequence 2 contains 35 pedestrians (1,936 images), and Sequence 3 contains 28 pedestrians (1,762 images). See Fig. 4 for examples.

We downsampled all images to 64×32 pixels. For each subject we randomly selected 10 images for training and used the rest for testing. Random selection of training and testing data was repeated 20 times to obtain reliable statistics. To describe each image, the covariance descriptor was computed using the following features:

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

where (x, y) is the position of a pixel, while $R_{x,y}$, $G_{x,y}$ and $B_{x,y}$ represent the corresponding colour information. The gradient and Laplacian for colour C are represented by $C'_{x,y} = [|\partial C / \partial x|, |\partial C / \partial y|]$ and $C''_{x,y} = [|\partial^2 C / \partial x^2|, |\partial^2 C / \partial y^2|]$, respectively.

We compared the proposed RSR method with several techniques previously used for pedestrian detection: Histogram Plus Epitome (HPE) [29], Symmetry-Driven Accumulation of Local Features (SDALF) [30], and Riemannian Locality Preserving Projection (RLPP) [10]. The performance of logE-SR was below HPE method and is not shown. The results for TSC could not be generated in a timely manner, due to the heavy computational load of the algorithm.

Results for Sequence 1 and 2 are shown in Fig. 5, in terms of cumulative matching characteristic (CMC) curves. The CMC curve represents the expectation of finding the correct match in the top n matches. The proposed method obtains the highest accuracy. For Sequence 3 (not shown), very similar performance is obtained by SDALF, RLPP and the proposed RSR, with HPE having the lowest performance.

5.2 Dictionary Learning

Here we compare the performance of the proposed Riemannian dictionary learning technique (as described in Section 4), with the performances of dictionaries obtained by random sampling and Riemannian k -means. We first use synthetic data to show that the proposed method obtains a lower representation error in RKHS, followed by classification experiments on texture data.

Synthetic Data. We synthesised 512 Riemannian samples from a set of 32 source points in Sym_+^5 . The source points can be considered as a form of ground-truth. The synthesised samples were then used for dictionary creation by Riemannian k -means [24] and the proposed algorithm.

To generate each source point, an SPD matrix was created by computing the covariance of 100 random samples of a 5 dimensional normal distribution. The mean and variance of the distribution are different for each source point. To synthesise each of the 512 Riemannian samples, we uniformly selected $T = 4$ source points and combined them with random positive weights, where the weights obeyed a normal distribution with zero mean and unit variance.

The performance is measured in terms of representation error in RKHS, ie. Eqn. (10). Fig. 6 shows the representation error as the algorithms iterate, with the proposed algorithm obtaining a lower error than Riemannian k -means.

Fig. 6. Representation error of learned dictionaries in RKHS, Eqn. (10), for synthetic data. The proposed method (Section 4) is compared with Riemannian k -means [24]. The source points can be interpreted as a form of ground-truth.

random	k -means	learning
46.09 ± 1.5	53.20 ± 1.1	60.65 ± 0.9

Table 3. Recognition accuracy (in %) for the texture classification task with dictionary learning. In all cases the proposed RSR approach was used, coupled with a dictionary generated via three separate methods: random dictionary generation, Riemannian k -means [24], and the proposed learning algorithm (Section 4).

Texture Classification. Here we consider a multi-class classification problem, using 111 texture images of the Brodatz texture dataset [27]. From each image we randomly extracted 50 blocks of size 32×32 . To train the dictionary, 20 blocks from each image were randomly selected, resulting in a dictionary learning problem with 2200 samples. From the remaining blocks, 20 per image were used as probe data and 10 as gallery samples. The process of random block creation and dictionary generation was repeated twenty times. The average recognition accuracies over probe data are reported here. In the same manner as in Section 5.1, we used the feature vector $F(x, y) = \left[I(x, y), \left| \frac{\partial I}{\partial x} \right|, \left| \frac{\partial I}{\partial y} \right|, \left| \frac{\partial^2 I}{\partial x^2} \right|, \left| \frac{\partial^2 I}{\partial y^2} \right| \right]$ to create the covariance, where the first dimension is the grayscale intensity, and the remaining dimensions capture first and second order gradients.

We used the proposed RSR approach to obtain the sparse codes, coupled with a dictionary generated via three separate methods: random dictionary generation, Riemannian k -means algorithm [24], and the proposed learning algorithm (Section 4). The sparse codes were then classified using a nearest-neighbour classifier.

For the randomly generated dictionary case, the classification rates are averaged over 10 runs, with each run using a different random dictionary. For all methods, dictionaries of size $k = \{8, 16, 24, \dots, 128\}$ were trained. The best results for each approach (ie., the results for the dictionary size that obtained the highest recognition accuracy) are reported in Table 3. For the random dictionary, $k = 64$; for the k -means dictionary, $k = 96$; for the proposed dictionary learning algorithm, $k = 24$. The results show that the proposed algorithm leads to a considerable gain in accuracy.

6 Main Findings and Future Directions

With the aim of addressing sparse representation on Riemannian manifolds, proposed to seek the solution through embedding the manifolds into RKHS, with the aid of the recently introduced Stein kernel. This led to a relaxed and extended version of the Lasso problem [1] on Riemannian manifolds.

Experiments on several classification tasks (face recognition, texture classification, person re-identification) show that the proposed approach achieves notable improvements in discrimination accuracy, in comparison to state-of-the-art methods such as tensor sparse coding, Riemannian locality preserving projection, and symmetry-driven accumulation of local features. We conjecture that this stems from better exploitation of Riemannian geometry, as the Stein kernel is related to geodesic distances via a tight bound. The proposed sparse coding method is also considerably faster than the state-of-the-art MAXDET reformulation used by Tensor Sparse Coding [9].

We have furthermore proposed an algorithm for learning a Riemannian dictionary, closely tied to the Stein kernel. In comparison to Riemannian k -means [24], the proposed algorithm obtains a lower representation error in RKHS and leads to improved classification accuracies.

Future directions include using the Stein kernel for solving large margin classification problems on Riemannian manifolds. This translates to designing a machinery that maximises a margin on SPD matrices based on Stein divergence, which can be considered as an extension of support vector machines [13] to tensor spaces.

Acknowledgements. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy, as well as by the Australian Research Council through the ICT Centre of Excellence program.

References

1. Elad, M.: Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing. Springer (2010)
2. Wright, J., Yang, A., Ganesh, A., Sastry, S., Ma, Y.: Robust face recognition via sparse representation. *IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI)* **31**(2) (2009) 210–227
3. Tuzel, O., Porikli, F., Meer, P.: Pedestrian detection via classification on Riemannian manifolds. *IEEE Trans. Pattern Analysis and Machine Intelligence* **30**(10) (2008) 1713–1727
4. Pennec, X.: Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision* **25**(1) (2006) 127–154
5. Lui, Y.M.: Advances in matrix manifolds for computer vision. *Image and Vision Computing* **30**(6–7) (2012) 380–388
6. Guo, K., Ishwar, P., Konrad, J.: Action recognition using sparse representation on covariance manifolds of optical flow. In: *IEEE Conf. Advanced Video and Signal Based Surveillance*. (2010) 188–195
7. Yuan, C., Hu, W., Li, X., Maybank, S., Luo, G.: Human action recognition under log-Euclidean Riemannian metric. In: *Asian Conference on Computer Vision*. Volume 5994 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg (2010) 343–353
8. Pang, Y., Yuan, Y., Li, X.: Gabor-based region covariance matrices for face recognition. *IEEE Transactions on Circuits and Systems for Video Technology* **18**(7) (2008) 989–993

9. Sivalingam, R., Boley, D., Morellas, V., Papanikolopoulos, N.: Tensor sparse coding for region covariances. In: European Conference on Computer Vision (ECCV). Volume 6314 of Lecture Notes in Computer Science. (2010) 722–735
10. Harandi, M.T., Sanderson, C., Wiliem, A., Lovell, B.C.: Kernel analysis over Riemannian manifolds for visual recognition of actions, pedestrians and textures. In: IEEE Workshop on the Applications of Computer Vision (WACV). (2012) 433–439
11. Hu, W., Li, X., Luo, W., Zhang, X., Maybank, S., Zhang, Z.: Single and multiple object tracking using log-Euclidean Riemannian subspace and block-division appearance model. *IEEE Trans. Pattern Analysis and Machine Intelligence* (in press) DOI: 10.1109/TPAMI.2012.42.
12. Bhatia, R.: Positive Definite Matrices. Princeton University Press (2007)
13. Shawe-Taylor, J., Cristianini, N.: Kernel Methods for Pattern Analysis. Cambridge University Press (2004)
14. Sra, S., Cherian, A.: Generalized dictionary learning for symmetric positive definite matrices with application to nearest neighbor retrieval. In: Machine Learning and Knowledge Discovery in Databases. Volume 6913 of Lecture Notes in Computer Science. (2011) 318–332
15. Kulis, B., Sustik, M.A., Dhillon, I.S.: Low-rank kernel learning with Bregman matrix divergences. *Journal of Machine Learning Research* **10** (2009) 341–376
16. Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press, New York, NY, USA (2004)
17. Sra, S.: Positive definite matrices and the symmetric Stein divergence. Preprint: [arXiv:1110.1773] (2012)
18. Yang, M., Zhang, L.: Gabor feature based sparse representation for face recognition with gabor occlusion dictionary. In: European Conference on Computer Vision (ECCV). Volume 6316 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg (2010) 448–461
19. Thompson, A.C.: On certain contraction mappings in a partially ordered vector space. *Proceedings of the American Mathematical Society* **14** (1963) 438–443
20. Chen, Y., Garcia, E.K., Gupta, M.R., Rahimi, A., Cazzanti, L.: Similarity-based classification: Concepts and algorithms. *Journal of Machine Learning Research* **10** (2009) 747–776
21. Grant, M., Boyd, S.: CVX: Matlab software for disciplined convex programming, version 1.21. <http://cvxr.com/cvx/> (April 2011)
22. Bishop, C.M.: Pattern Recognition and Machine Learning. Springer (2006)
23. Wong, Y., Harandi, M.T., Sanderson, C., Lovell, B.C.: On robust biometric identity verification via sparse encoding of faces: holistic vs local approaches. In: IEEE International Joint Conference on Neural Networks. (2012) 1762–1769
24. Goh, A., Vidal, R.: Clustering and dimensionality reduction on Riemannian manifolds. In: IEEE Conf. Computer Vision and Pattern Recognition. (2008) 1–7
25. Phillips, P., Moon, H., Rizvi, S., Rauss, P.: The FERET evaluation methodology for face-recognition algorithms. *IEEE Trans. Pattern Analysis and Machine Intelligence* **22**(10) (2000) 1090–1104
26. Lee, T.S.: Image representation using 2d Gabor wavelets. *IEEE Trans. Pattern Analysis and Machine Intelligence* **18** (1996) 959–971
27. Randen, T., Husøy, J.H.: Filtering for texture classification: A comparative study. *IEEE Trans. Pattern Analysis and Machine Intelligence* **21**(4) (1999) 291–310
28. Ess, A., Leibe, B., Gool, L.V.: Depth and appearance for mobile scene analysis. *Int. Conf. Computer Vision (ICCV)* (2007) 1–8
29. Bazzani, L., Cristani, M., Perina, A., Farenzena, M., Murino, V.: Multiple-shot person re-identification by HPE signature. In: Int. Conf. Pattern Recognition. (2010) 1413–1416
30. Farenzena, M., Bazzani, L., Perina, A., Murino, V., Cristani, M.: Person re-identification by symmetry-driven accumulation of local features. *IEEE Conf. Computer Vision and Pattern Recognition* (2010) 2360–2367
31. Schwartz, W.R., Davis, L.S.: Learning discriminative appearance-based models using partial least squares. In: Brazilian Symposium on Computer Graphics and Image Processing. (2009) 322–329