Published in final edited form as:

Comput Vis ECCV. 2014 ; 8690: 251–267. doi:10.1007/978-3-319-10605-2_17.

Canonical Correlation Analysis on Riemannian Manifolds and Its Applications

Hyunwoo J. Kim¹, Nagesh Adluru¹, Barbara B. Bendlin¹, Sterling C. Johnson¹, Baba C. Vemuri², and Vikas Singh¹

¹University of Wisconsin–Madison

²University of Florida

Abstract

Canonical correlation analysis (CCA) is a widely used statistical technique to capture correlations between two sets of multi-variate random variables and has found a multitude of applications in computer vision, medical imaging and machine learning. The classical formulation assumes that the data live in a pair of *vector spaces* which makes its use in certain important scientific domains problematic. For instance, the set of symmetric positive definite matrices (SPD), rotations and probability distributions, all belong to certain curved Riemannian manifolds where vector-space operations are in general not applicable. Analyzing the space of such data via the classical versions of inference models is rather sub-optimal. But perhaps more importantly, since the algorithms do not respect the underlying geometry of the data space, it is hard to provide statistical guarantees (if any) on the results. Using the space of SPD matrices as a concrete example, this paper gives a principled generalization of the well known CCA to the Riemannian setting. Our CCA algorithm operates on the product Riemannian manifold representing SPD matrix-valued fields to identify meaningful statistical relationships on the product Riemannian manifold. As a proof of principle, we present results on an Alzheimer’s disease (AD) study where the analysis task involves identifying correlations across diffusion tensor images (DTI) and Cauchy deformation tensor fields derived from T1-weighted magnetic resonance (MR) images.

1 Introduction

Canonical correlation analysis (CCA) is a powerful statistical technique to extract linear components that capture correlations between two multi-variate random variables [15]. CCA provides an answer to the following question: suppose we are given data of the form,

$(\mathbf{x}_i \in \mathcal{X}, \mathbf{y}_i \in \mathcal{Y})_{i=1}^N \subset \mathcal{X} \times \mathcal{Y}$ where $\mathbf{x}_i \in \mathbb{R}^m$ and $\mathbf{y}_i \in \mathbb{R}^n$, find a model that explains *both* of these observations. More precisely, CCA provides an answer to this question by identifying a pair of directions where the projections (namely, u and v) of the random variables, $\mathbf{x}$ and $\mathbf{y}$ yield maximum correlation $\rho_{u,v} = \text{COV}(u, v) / \sigma_u \sigma_v$. Here, $\text{COV}(u, v)$ denotes the covariance function and σ gives the standard deviation. During the last decade, the CCA formulation has been broadly applied to various unsupervised learning problems in computer vision and machine learning including image retrieval [11], face/gait recognition [38], super-resolution [19] and action classification [24].

Beyond the applications described above, a number of works have recently investigated the use of CCA in analyzing neuroimaging data [3], which is a main focus of this paper. Here, for each participant in a clinical study, we acquire different types of images such as Magnetic Resonance (MRI), Computed Tomography (CT) and functional MRI. It is expected that each imaging modality captures a unique aspect of the underlying disease pathology. Therefore, given a group of N subjects and their corresponding brain images, we may want to identify strong relationships (e.g., anatomical/functional correlations) across different image types. When performed across different diseases, such an analysis will reveal insights into what is similar and what is different across diseases even when their symptomatic presentation may be similar. Alternatively, CCA may serve a feature extraction role. That is, the brain regions found to be strongly correlated can be used directly in downstream statistical analysis. In a study of a large number of subjects, rather than performing a hypothesis test on *all* brain voxels independently for each imaging modality, restricting the number of tests only to the set of ‘relevant’ voxels (found via CCA) is known to improve statistical power (since the False Discovery Rate correction will be less severe).

The classical version of CCA described above concurrently seeks two linear subspaces (straight lines) in *vector spaces* $\mathbf{R}^m$ and $\mathbf{R}^n$ for the two multi-variate random variables $\mathbf{x}$ and $\mathbf{y}$. The projection on to the straight line (linear subspace) is obtained by an inner product. This formulation is broadly applicable but encounters problems for manifold-valued data that are becoming increasingly important in present day research. For example, diffusion tensor magnetic resonance images (DTI) allow one to infer the diffusion tensor characterizing the anisotropy of water diffusion at each voxel in an image volume. This tensorial feature can be visualized as an ellipsoid and represented by a 3×3 symmetric positive definite (SPD) matrix at each voxel in the acquired image volume. Neither the individual SPD matrices nor the field of these SPD matrices lie in a vector space but instead are elements of a negatively curved Riemannian manifold where standard vector space operations are not valid. Hence, classical CCA is not applicable in this setting. For T1-weighted Magnetic resonance images (MRIs), we are frequently interested in analyzing not just the 3D intensity image on its own, but rather a quantity that captures the deformation field between each image and a *population template*. A registration between the image and the template yields the deformation field required to align the image pairs and the determinant of the Jacobian J of this deformation at each voxel is a commonly used feature that captures local volume changes [6,17]. Quantities such as the Cauchy deformation tensor defined as $\sqrt{J^T J}$ have been reported in literature for use in morphometric analysis [18]. The input to the statistical analysis is a 3D image of voxels, where each voxel corresponds to a matrix $\sqrt{J^T J} \succ 0$ (the Cauchy deformation tensor). Another example of manifold-valued fields is derived from high angular resolution diffusion images (HARDI) and can be used to compute the ensemble average propagators (EAPs) at each voxel of the given HARDI data. The EAP is a probability density function that is related to the diffusion sensitized MR signal via the Fourier transform [5]. Since an EAP is a probability density function, by using a square root parameterization of this density function, it is possible to identify it with a point on the unit Hilbert Sphere. Once again, to perform any statistical analysis of these data derived features, we cannot apply standard vector-space operations since the unit Hilbert sphere is a positively curved manifold. When analyzing real brain imaging data, it is entirely

possible that no meaningful correlations exist in the data. The key difficulty is that we do not know whether the experiment (i.e., inference) failed because there is in fact no statistically meaningful signal in the dataset or if the algorithms being used are sub-optimal.

Related Work

There are two somewhat distinct bodies of work that are related to and motivate this work. The first one relates to the extensive study of the classical CCA and its non-linear variants. These include various interesting results based on kernelization [1,4,12], neural networks [25,16], and deep architectures [2]. Most, if not all of these strategies extend CCA to arbitrary nonlinear spaces. However, this flexibility brings with it the associated issues of model selection (and thereby, regularization), controlling the complexity of the neural network structure, choosing an appropriate activation function and so on. It is an interesting question though not completely clear to us what type of a regularizer should be used if one were to explicitly impose a Riemannian structure on the objectives described in the works above. As opposed to regularization, the second line of work incorporates the specific geometry of the data directly within the estimation problem. Various statistical constructs have been generalized to Riemannian manifolds: these include regression [39,31], classification [36], kernel methods [21], margin-based and boosting classifiers [26], interpolation, convolution, filtering [10] and dictionary learning [14,27]. Among the most closely related are ideas related to projective dimensionality reduction methods. For instance, the generalization of Principal Components analysis (PCA) via the so-called Principal Geodesic Analysis (PGA) [9], Geodesic PCA [20], Exact PGA [33], Horizontal Dimension Reduction [32] with frame bundles, and an extension of PGA to the product space of Riemannian manifolds, namely, tensor fields [36]. It is important to note that except the non-parametric method of [34], most of these strategies focus on one rather than two sets of random variables (as is the case in CCA). Even in this setting, the first results on successful generalization of parametric regression models to Riemannian manifolds is relatively recent: geodesic regression [8,29] and polynomial regression [13] (note that the adaptive CCA formulation in [37] seems related to our work but is not designed for manifold-valued data).

This paper provides a parametric model between two different tensor fields on a Riemannian manifold, which is a significant step beyond these recent works. The CCA formulation we present requires the optimization of functions over either a single product manifold or a pair of product manifolds (of different dimensions) concurrently. The latter problem involving product manifolds of different dimensions will not be addressed in this paper. Note that in general, on manifolds the projection operation does not have a nice closed form solution. So, we need to perform projections via an optimization scheme on the two manifolds and find the best pair of geodesic subspaces. We provide a precise solution to this problem. To our knowledge, this is the first extension of CCA to Riemannian manifolds. Our approach has two advantages relative to other non-linear extensions of CCA. The first advantage is that no model selection is required. Also our method incorporates the known geometry of data space. Our **key contributions** are: **a)** A principled generalization of CCA for Riemannian manifolds; **b)** First, a numerical optimization scheme for identifying the subspaces and later, single path algorithms with approximate projections (both these ideas may be applicable

beyond the CCA formulation). c) Providing experimental evidence how the Riemannian CCA formulation expands the operating range of statistical analysis of neuroimaging data.

2 Canonical Correlation in Euclidean Space

First, we will briefly review the classical CCA in Euclidean space to motivate the rest of our presentation. Recall that Pearson's product-moment correlation coefficient is a quantity to measure the relationship of two random variables, $x \in \mathbf{R}$ and $y \in \mathbf{R}$. For one dimensional random variables,

$$\rho_{x,y} = \frac{\text{COV}(x,y)}{\sigma_x \sigma_y} = \frac{\mathbb{E}[(x - \mu_x)(y - \mu_y)]}{\sigma_x \sigma_y} = \frac{\sum_{i=1}^N (x_i - \mu_x)(y_i - \mu_y)}{\sqrt{\sum_{i=1}^N (x_i - \mu_x)^2} \sqrt{\sum_{i=1}^N (y_i - \mu_y)^2}} \quad (1)$$

For high dimensional data, $\mathbf{x} \in \mathbb{R}^m$ and $\mathbf{y} \in \mathbb{R}^n$, we cannot however perform a direct calculation as above. So, we need to project each set of variables on to a special axis in each space $\mathcal{X}$ and $\mathcal{Y}$. CCA generalizes the concept of correlation to random vectors (potentially of different dimensions). It is convenient to think of CCA as a measure of correlation between two multivariate data based on the *best* projection which maximizes their mutual correlation.

Canonical Correlation for $\mathbf{x} \in \mathbb{R}^m$ and $\mathbf{y} \in \mathbb{R}^n$ is given by

$$\max_{\mathbf{w}_x, \mathbf{w}_y} \text{corr}(\pi_{\mathbf{w}_x}(\mathbf{x}), \pi_{\mathbf{w}_y}(\mathbf{y})) = \max_{\mathbf{w}_x, \mathbf{w}_y} \frac{\sum_{i=1}^N \mathbf{w}_x^T (\mathbf{x}_i - \mu_x) \mathbf{w}_y^T (\mathbf{y}_i - \mu_y)}{\sqrt{\sum_{i=1}^N (\mathbf{w}_x^T (\mathbf{x}_i - \mu_x))^2} \sqrt{\sum_{i=1}^N (\mathbf{w}_y^T (\mathbf{y}_i - \mu_y))^2}} \quad (2)$$

where $\pi_{\mathbf{w}_x}(\mathbf{x}) := \arg\min_{t \in \mathbb{R}} d(t\mathbf{w}_x, \mathbf{x})^2$. We will call $\pi_{\mathbf{w}_x}(\mathbf{x})$ the *projection coefficient* for $\mathbf{x}$ (similarly for $\mathbf{y}$). Define $S_{\mathbf{w}_x}$ as the subspace which is the span of $\mathbf{w}_x$. The projection of $\mathbf{x}$ on to $S_{\mathbf{w}_x}$ is given by $\Pi_{S_{\mathbf{w}_x}}(\mathbf{x})$. We can then verify that the relationship between the projection and the projection coefficient is,

$$\Pi_{S_{\mathbf{w}_x}}(\mathbf{x}) := \arg \min_{\mathbf{x}' \in S_{\mathbf{w}_x}} d(\mathbf{x}, \mathbf{x}')^2 = \frac{\mathbf{w}_x^T \mathbf{x}}{\|\mathbf{w}_x\|} \frac{\mathbf{w}_x}{\|\mathbf{w}_x\|} = \frac{\mathbf{w}_x^T \mathbf{x}}{\|\mathbf{w}_x\|^2} \mathbf{w}_x = \pi_{\mathbf{w}_x}(\mathbf{x}) \mathbf{w}_x \quad (3)$$

In the Euclidean space, $\Pi_{S_{\mathbf{w}_x}}(\mathbf{x})$ has a closed form solution. In fact, it is obtained by an inner product, $\mathbf{w}_x^T \mathbf{x}$. Hence, by replacing the projection coefficient $\pi_{\mathbf{w}_x}(\mathbf{x})$ with $\mathbf{w}_x^T \mathbf{x} / \|\mathbf{w}_x\|^2$ and after a simple calculation, one obtains the form in (2). Without loss of generality, assume that $\mathbf{x}, \mathbf{y}$ are centered. Then the optimization problem can be written as,

$$\max_{\mathbf{w}_x, \mathbf{w}_y} \mathbf{w}_x^T X^T Y \mathbf{w}_y \text{ subject to } \mathbf{w}_x^T X^T X \mathbf{w}_x = \mathbf{w}_y^T Y^T Y \mathbf{w}_y = 1 \quad (4)$$

where $\mathbf{x}, \mathbf{w}_x \in \mathbf{R}^m, \mathbf{y}, \mathbf{w}_y \in \mathbf{R}^n, X = [\mathbf{x}_1 \dots \mathbf{x}_N]^T$ and $Y = [\mathbf{y}_1 \dots \mathbf{y}_N]^T$. The only difference here is that we remove the denominator. Instead, we have two equality constraints (note that correlation is scale-invariant).

3 Mathematical Preliminaries

We now briefly summarize certain basic concepts [7] which we will use later.

Riemannian Manifolds

A *differentiable manifold* [7] of dimension n is a set $\mathcal{M}$ and a family of *injective* mappings $\phi_i : U_i \subset \mathbf{R}^n \rightarrow \mathcal{M}$ of open sets U_i of $\mathbf{R}^n$ into $\mathcal{M}$ such that: **(1)** $\cup_i \phi_i(U_i) = \mathcal{M}$; **(2)** for any pair i, j with $\phi_i(U_i) \cap \phi_j(U_j) = W \neq \emptyset$, the sets $\phi_i^{-1}(W)$ and $\phi_j^{-1}(W)$ are open sets in $\mathbf{R}^n$ and the mappings $\phi_j^{-1} \circ \phi_i$ are differentiable, where $\circ$ denotes function composition. In other words, a differentiable manifold $\mathcal{M}$ is a topological space that is locally similar to an Euclidean space and has a globally defined differential structure. The tangent space at a point p on the manifold, $T_p \mathcal{M}$, is a vector space that consists of the tangent vectors of *all* possible curves passing through p .

A Riemannian manifold is equipped with a smoothly varying inner product. The family of inner products on all tangent spaces is known as the *Riemannian metric* of the manifold. The *geodesic distance* between two points on $\mathcal{M}$ is the length of the shortest *geodesic* curve connecting the two points, analogous to straight lines in $\mathbf{R}^n$. The geodesic curve from x_i to x_j can be parameterized by a tangent vector in the tangent space at y_i with an exponential map $\text{Exp}(y_i, \cdot) : T_{y_i} \mathcal{M} \rightarrow \mathcal{M}$. The inverse of the exponential map is the logarithm map, $\text{Log}(y_i, \cdot) : \mathcal{M} \rightarrow T_{y_i} \mathcal{M}$. Separate from these notations, matrix exponential (and logarithm) are given as $\exp(\cdot)$ (and $\log(\cdot)$).

Intrinsic Mean

Let $d(\cdot, \cdot)$ define the geodesic distance between two points. The intrinsic (or Karcher) mean of a set of points $\{x_i\}$ with non-negative weights $\{w_i\}$ is the minimizer of,

$$\hat{\mu} = \arg \min_{y \in \mathcal{M}} \sum_{i=1}^N w_i d(y, y_i)^2, \quad (5)$$

which may be an arithmetic, geometric or harmonic mean depending on $d(\cdot, \cdot)$.

On manifolds, the Karcher mean with distance $d(y_i, y_j) = \|\text{Log}_{y_i} y_j\|$ is, $\sum_{i=1}^N \text{Log}_{\bar{y}} y_i = 0$. This identity implies that $\bar{y}$ is a local minimum which has a zero norm gradient [22], i.e., the sum of all tangent vectors corresponding to geodesic curves from mean $\bar{y}$ to all points y_i is zero in the tangent space $T_{\bar{y}} \mathcal{M}$. On manifolds, the existence and uniqueness of the Karcher mean is not guaranteed, unless we assume, for uniqueness, that the data is in a small neighborhood.

Geodesically Convex

A subset C of $\mathcal{M}$ is said to be a *geodesically convex set* if there is a minimizing geodesic curve in C between *any* two points in C . This assumption is commonly used [8] and essential to ensure that the Riemannian operations such as the exponential and logarithm maps are well-defined.

4 A Model for CCA on Riemannian Manifolds

We now present a step by step derivation of our Riemannian CCA model. Classical CCA finds the mean of each data modality. Then, it maximizes correlation between projected data on each subspace at the mean. Similarly, CCA on manifolds must first compute the intrinsic mean (i.e., Karcher mean) of each data set. It must then identify a ‘generalized’ version of a subspace at each Karcher mean to maximize the correlation of projected data. The generalized form of a subspace on Riemannian manifolds has been studied in the literature [33,26,20,9]. The so-called *geodesic submanifold* [9,36,23] which has been used for geodesic regression serves our purpose well and is defined as $S = \text{Exp}(\boldsymbol{\mu}, \text{span}(\{\mathbf{v}_i\}) \cap U)$, where $U \subset T_{\boldsymbol{\mu}}\mathcal{M}$, and $\mathbf{v}_i \in T_{\boldsymbol{\mu}}\mathcal{M}$ [9]. When S has only one tangent vector $\mathbf{v}$, then the geodesic submanifold is simply a geodesic curve, see Figure 1.

We can now proceed to formulate the precise form of projection on to a geodesic submanifold. Recall that when given a point, its projection on a set is the closest point in the set. So, the projection on to a geodesic submanifold (S) must be a function satisfying this behavior. This is given by,

$$\Pi_S(\mathbf{x}) = \arg \min_{\mathbf{x}' \in S} d(\mathbf{x}, \mathbf{x}')^2 \quad (6)$$

In Euclidean space, the projection on a convex set (e.g., subspace) is unique. It is also unique on some manifolds under special conditions, e.g., quaternion sphere [30]. However, the uniqueness of the projection on geodesic submanifolds in general conditions cannot be ensured. Like other methods, we assume that given the specific manifold and the data, the projection is well-posed.

Finally, the correlation of points (*after* projection) can be measured by the distance from the mean to the projected points. To be specific, the projection on a geodesic submanifold corresponding to $\mathbf{w}_x$ in classical CCA is given by

$$\Pi_{S_{\mathbf{w}_x}}(\mathbf{x}) := \arg \min_{\mathbf{x}' \in S_{\mathbf{w}_x}} \|\text{Log}(\mathbf{x}, \mathbf{x}')\|_{\mathbf{x}}^2 \quad (7)$$

$S_{\mathbf{w}_x} := \text{Exp}(\boldsymbol{\mu}_x, \text{span}\{\mathbf{w}_x\} \cap U)$ where $\mathbf{w}_x$ is a basis tangent vector and $U \subset T_{\boldsymbol{\mu}_x}\mathcal{M}_x$ is a small neighborhood of $\boldsymbol{\mu}_x$. The expression for *projection coefficients* can now be given as

$$t_i = \pi_{\mathbf{w}_x}(\mathbf{x}_i) := \arg \min_{t'_i \in (-\varepsilon, \varepsilon)} \|\text{Log}(\text{Exp}(\boldsymbol{\mu}_x, t'_i \mathbf{w}_x), \mathbf{x}_i)\|_{\boldsymbol{\mu}_x}^2 \quad (8)$$

where $\mathbf{x}_i, \boldsymbol{\mu}_x \in \mathcal{M}_x, \mathbf{w}_x \in T_{\boldsymbol{\mu}_x}\mathcal{M}_x, t_i \in \mathbf{R}$. The term, $u_i = \pi_{\mathbf{w}_y}(\mathbf{y})$ is defined analogously. t_i is a real value to obtain the point $\Pi_{S_{\mathbf{w}_x}}(\mathbf{x}) = \text{Exp}(\boldsymbol{\mu}_x, t_i \mathbf{w}_x)$. As mentioned above, $\mathbf{x}$ and $\mathbf{y}$ belong to the same manifold. Note that we are dealing with a single manifold, however, we use two different notations $\mathcal{M}_x$, and $\mathcal{M}_y$ to show that they are differently distributed for ease of discussion.

Notice that we have $d(\boldsymbol{\mu}_x, \Pi_{S_{\mathbf{w}_x}}(\mathbf{x}_i)) = \|\text{Log}(\boldsymbol{\mu}_x, \text{Exp}(\boldsymbol{\mu}_x, t_i \mathbf{w}_x))\|_{\boldsymbol{\mu}_x} = t_i \|\mathbf{w}_x\|_{\boldsymbol{\mu}_x}$. By inspection, this shows that the projection coefficient is proportional to the length of the geodesic curve

from the base point μ_x to the projection of $\mathbf{x}$, $\Pi_{S_{w_x}}(\mathbf{x})$. Correlation is scale invariant, as expected. Therefore, the correlation between projected points $\{\Pi_{S_{w_x}}(\mathbf{x}_i)\}_{i=1}^N$ and $\{\Pi_{S_{w_y}}(\mathbf{y}_i)\}_{i=1}^N$ reduces to the correlation between the quantities that serve as projection coefficients here, $\{t_i\}_{i=1}^N$ and $\{u_i\}_{i=1}^N$.

Putting these pieces together, we obtain our generalized formulation for CCA,

$$\rho_{\mathbf{x},\mathbf{y}} = \text{corr}(\pi_{w_x}(\mathbf{x}), \pi_{w_y}(\mathbf{y})) = \max_{w_x, w_y, \mathbf{t}, \mathbf{u}} \frac{\sum_{i=1}^N (t_i - \bar{t})(u_i - \bar{u})}{\sqrt{\sum_{i=1}^N (t_i - \bar{t})^2} \sqrt{\sum_{i=1}^N (u_i - \bar{u})^2}} \quad (9)$$

where $t_i = \pi_{w_x}(\mathbf{x}_i)$, $\mathbf{t} := \{t_i\}$, $u_i = \pi_{w_y}(\mathbf{y}_i)$, $\mathbf{u} := \{u_i\}$, $\bar{t} = \frac{1}{N} \sum_{i=1}^N t_i$ and $\bar{u} = \frac{1}{N} \sum_{i=1}^N u_i$. Expanding out components in (9) further, it takes the form,

$$\begin{aligned} \rho_{\mathbf{x},\mathbf{y}} = & \max_{w_x, w_y, \mathbf{t}, \mathbf{u}} \frac{\sum_{i=1}^N (t_i - \bar{t})(u_i - \bar{u})}{\sqrt{\sum_{i=1}^N (t_i - \bar{t})^2} \sqrt{\sum_{i=1}^N (u_i - \bar{u})^2}} \\ & s.t. t_i = \arg \min_{t_i \in (-\varepsilon, \varepsilon)} \|\text{Log}(\text{Exp}(\mu_x, t_i w_x), \mathbf{x}_i)\|^2, \forall i \in \{1, \dots, N\} \\ & u_i = \arg \min_{u_i \in (-\varepsilon, \varepsilon)} \|\text{Log}(\text{Exp}(\mu_y, u_i w_y), \mathbf{y}_i)\|^2, \forall i \in \{1, \dots, N\} \end{aligned} \quad (10)$$

Directly, we see that (10) is a multilevel optimization and solutions from nested sub-optimization problems may be needed to solve the higher level problem. It turns out that deriving the first order optimality conditions suggests a cleaner approach.

Define $f(\mathbf{t}, \mathbf{u}) := \frac{\sum_{i=1}^N (t_i - \bar{t})(u_i - \bar{u})}{\sqrt{\sum_{i=1}^N (t_i - \bar{t})^2} \sqrt{\sum_{i=1}^N (u_i - \bar{u})^2}}$, $g(t_i, w_x) := \|\text{Log}(\text{Exp}(\mu_x, t_i w_x), \mathbf{x}_i)\|^2$, and $g(u_i, w_y) := \|\text{Log}(\text{Exp}(\mu_y, u_i w_y), \mathbf{y}_i)\|^2$. Then, we may replace the equality constraints in (10) with optimality conditions rather than another optimization problem for each i . Using this idea, we have

$$\rho(w_x, w_y) = \max_{w_x, w_y, \mathbf{t}, \mathbf{u}} f(\mathbf{t}, \mathbf{u}) \text{ s.t. } \nabla_{t_i} g(t_i, w_x) = 0, \nabla_{u_i} g(u_i, w_y) = 0, \forall i \in \{1, \dots, N\} \quad (11)$$

5 Optimization Schemes

We present two different algorithms to solve the problem of computing CCA on Riemannian manifolds. The first algorithm is based on a numerical optimization for (11). We only summarize the main model here and provide all technical details in the extended version for space reasons. Subsequently, we present the second approach which is based on an approximation for a more efficient algorithm.

5.1 An Augmented Lagrangian Method

The augmented Lagrangian technique is a well known variation of the penalty method for constrained optimization problems. Given a constrained optimization problem $\max f(\mathbf{x})$ s.t.

$c_i(\mathbf{x}) = 0, \forall i$, the augmented Lagrangian method solves a sequence of the following models while increasing ν_k .

$$\max_{\mathbf{x}} f(\mathbf{x}) + \sum_i \lambda_i c_i(\mathbf{x}) - \nu^k \sum_i c_i(\mathbf{x})^2 \quad (12)$$

The augmented Lagrangian formulation for our CCA formulation is given by

$$\begin{aligned} & \max_{\mathbf{w}_x, \mathbf{w}_y, \mathbf{t}, \mathbf{u}} \mathcal{L}_A(\mathbf{w}_x, \mathbf{w}_y, \mathbf{t}, \mathbf{u}, \boldsymbol{\lambda}^k; \nu^k) \\ &= \max_{\mathbf{w}_x, \mathbf{w}_y, \mathbf{t}, \mathbf{u}} f(\mathbf{t}, \mathbf{u}) \\ &+ \sum_i^N \lambda_{t_i}^k \nabla_{t_i} g(t_i, \mathbf{w}_x) \\ &+ \sum_i^N \lambda_{u_i}^k \nabla_{u_i} g(u_i, \mathbf{w}_y) \\ &- \frac{\nu^k}{2} \left(\sum_{i=1}^N \nabla_{t_i} g(t_i, \mathbf{w}_x)^2 \right. \\ &\quad \left. + \nabla_{u_i} g(u_i, \mathbf{w}_y)^2 \right) \end{aligned} \quad (13)$$

The pseudocode for our algorithm is summarized in Algorithm 1.

Remarks—Note that for Algorithm 1, we need the second derivative of g , in particular, for $\frac{d^2}{dw dt} g, \frac{d^2}{dt^2} g$. The literature does not provide a great deal of guidance on second derivatives of functions involving $\text{Log}(\cdot)$ and $\text{Exp}(\cdot)$ maps on general Riemannian manifolds. However, depending on the manifold, it can be obtained analytically or numerically (see extended version of the paper).

Approximate strategies—It is clear that the core difficulty in deriving the algorithm above was the lack of a closed form solution to projections on to geodesic submanifolds. If however, an approximate form of the projection can lead to significant gains in computational efficiency with little sacrifice in accuracy, it is worthy of consideration. The simplest approximation is to use a Log-Euclidean model. But it is well known that the Log-Euclidean is reasonable for data that are tightly clustered on the manifold and not otherwise. Further, the Log-Euclidean metric lacks the important property of affine invariance. We can obtain a more

Algorithm 1

Riemannian CCA based on the Augmented Lagrangian method

-
- 1: $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathcal{M}_x, \mathbf{y}_1, \dots, \mathbf{y}_N \in \mathcal{M}_y$
 - 2: Given $\nu^0 > 0, \tau^0 > 0$, starting points $(\mathbf{w}_x^0, \mathbf{w}_y^0, \mathbf{t}^0, \mathbf{u}^0)$ and $\boldsymbol{\lambda}^0$

```

3: for  $k = 0, 1, 2 \dots$  do
4:   Start at  $(\mathbf{w}_x^k, \mathbf{w}_y^k, \mathbf{t}^k, \mathbf{u}^k)$ 
5:   Find an approximate minimizer  $(\mathbf{w}_x^k, \mathbf{w}_y^k, \mathbf{t}^k, \mathbf{u}^k)$  of  $\mathcal{L}_A(\cdot, \lambda^k; \mathbf{v}^k)$ , and terminate when  $\|\nabla \mathcal{L}_A(\mathbf{w}_x^k, \mathbf{w}_y^k, \mathbf{t}^k, \mathbf{u}^k, \lambda^k; \mathbf{v}^k)\| \leq \tau^k$ 
6:   if a convergence test for (11) is satisfied then
7:     Stop with approximate feasible solution
8:   end if
9:    $\lambda_{t_i}^{k+1} = \lambda_{t_i}^k - \nu^k \nabla_{t_i} g(t_i, \mathbf{w}_x), \forall i$ 
10:   $\lambda_{u_i}^{k+1} = \lambda_{u_i}^k - \nu^k \nabla_{u_i} g(u_i, \mathbf{w}_y), \forall i$ 
11:  Choose new penalty parameter  $\mathbf{v}^{k+1} = \mathbf{v}^k$ 
12:  Set starting point for the next iteration
13:  Select tolerance  $\tau^{k+1}$ 
14: end for

```

Algorithm 2

CCA with approximate projection

```

1: Input  $X_1, \dots, X_N \in \mathcal{M}_x, Y_1, \dots, Y_N \in \mathcal{M}_y$ 
2: Compute intrinsic mean  $\mu_x, \mu_y$  of  $\{X_i\}, \{Y_i\}$ 
3: Compute  $X_i^\lambda = \text{Log}(\mu_x, X_i), Y_i^\lambda = \text{Log}(\mu_y, Y_i)$ 
4: Transform (using group action)  $\{X_i^\lambda\}, \{Y_i^\lambda\}$  to the  $T_I \mathcal{M}_x, T_I \mathcal{M}_y$ 
5: Perform CCA between  $T_I \mathcal{M}_x, T_I \mathcal{M}_y$  and get axes  $W_a \in T_I \mathcal{M}_x, W_b \in T_I \mathcal{M}_y$ 
6: Transform (using group action)  $W_a, W_b$  to  $T_{\mu_x} \mathcal{M}_x, T_{\mu_y} \mathcal{M}_y$ 

```

accurate projection using the submanifold expression given in [36]. The form of projection is,

$$\Pi_S(\mathbf{x}) \approx \text{Exp}(\mu, \sum_{i=1}^d \mathbf{v}_i \langle \mathbf{v}_i, \text{Log}(\mu, \mathbf{x}) \rangle_\mu) \quad (14)$$

where $\{\mathbf{v}_i\}$ are *orthonormal basis* at $T_\mu \mathcal{M}$. The CCA algorithm with this approximation for the projection is summarized as Algorithm 2.

Finally, we provide a brief remark on one remaining issue. This relates to the question why we use group action rather than other transformations such as parallel transport. Observe that Algorithm 2 sends the data from the tangent space at the Karcher mean of the samples to the tangent space at Identity I . The purpose of the transformation is to put all samples at the Identity of the SPD manifold, to obtain a more accurate projection, which can be understood using (14). The projection and inner product depend on the anchor point μ . If μ is

Identity, then there is no discrepancy between the Euclidean and the Riemannian inner products. Of course, one may use a parallel transport. However, group action may be substantially more efficient than parallel transport since the former does not require computing a geodesic curve (which is needed for parallel transport). Interestingly, it turns out that on SPD manifolds with a GL-invariant metric, parallel transport from an arbitrary point p to Identity I is *equivalent* to the transform using a group action. So, one can parallel transport tangent vectors from p to I using the group action more efficiently. The proof of Theorem 1 is available in the extended version.

Theorem 1: *On SPD manifold, let $\Gamma_{p \rightarrow I}(w)$ denote the parallel transport of $w \in T_p \mathcal{M}$ along the geodesic from $p \in \mathcal{M}$ to $I \in \mathcal{M}$. The parallel transport is equivalent to group action by $p^{-1/2}wp^{-T/2}$, where the inner product $\langle u, v \rangle_p = \text{tr}(p^{-1/2}up^{-1/2}v)$*

5.2 Extensions to the Product Riemannian Manifold

In the types of imaging datasets of interest in this paper, we seek to perform an analysis on an entire population of images (of multiple types). For such data, each image must be treated as a single entity, which necessitates extending the formulation above to a Riemannian product space.

Let us define a Riemannian metric on the product space $\mathcal{M} = \mathcal{M}_1 \times \dots \times \mathcal{M}_m$. A natural choice is the following idea from [36].

$$\langle \mathbf{X}_1, \mathbf{X}_2 \rangle_P = \sum_{j=1}^m \langle X_1^j, X_2^j \rangle_{P^j} \quad (15)$$

where $\mathbf{X}_1 = (X_1^1, \dots, X_1^m) \in \mathcal{M}$, and $\mathbf{X}_2 = (X_2^1, \dots, X_2^m) \in \mathcal{M}$ and $\mathbf{P} = (P^1, \dots, P^m) \in \mathcal{M}$. Once we have the exponential and logarithm maps, CCA on a Riemannian product space can be directly performed by Algorithm 2. The exponential map $\text{Exp}(\mathbf{P}, \mathbf{V})$ and logarithm map $\text{Log}(\mathbf{P}, \mathbf{X})$ are given by

$$(\text{Exp}(P^1, V^1), \dots, \text{Exp}(P^m, V^m)) \text{ and } (\text{Log}(P^1, X^1), \dots, \text{Log}(P^m, X^m)) \quad (16)$$

respectively, where $\mathbf{V} = (V^1, \dots, V^m) \in T_P \mathcal{M}$. The length of tangent vector is

$\|\mathbf{V}\| = \sqrt{\|V^1\|_{P^1}^2 + \dots + \|V^m\|_{P^m}^2}$, where $V^i \in T_{P^i} \mathcal{M}_i$. The geodesic distance between two points $d(\mathbf{X}_1, \mathbf{X}_2)$ on Riemannian product space is also measured by the length of tangent vector from one point to the other. So we have

$$d(\mu_x, \mathbf{X}) = \sqrt{d(\mu_x^1, X^1)^2 + \dots + d(\mu_x^m, X^m)^2} \quad (17)$$

From our previous discussion of the relationship between *projection coefficients* and distance from the mean to points (after *projection*) in Section 4, we have $t_i = d(\mu_x, \Pi_{S_{W_x}}(X_i)) / \|\mathbf{W}_x\|_{\mu_x}$ and $t_i^j = d(\mu_x^j, \Pi_{S_{W_x^j}}(X_i^j)) / \|W_x^j\|_{\mu_x^j}$. By substitution, the *projection coefficients* on Riemannian product space are given by

$$t_i = d(\mu_x, \Pi_{S_{W_x}}(X_i)) / \|W_x\|_{\mu_x} = \sqrt{\sum_j^m (t_i^j \|W_x^j\|_{\mu_x})^2 / \sum_{j=1}^m \|W_x^j\|_{\mu_x}^2} \quad (18)$$

We can now mechanically substitute these “product space” versions of the terms in (18) to derive a CCA on Riemannian product space. The full model is provided in the extended version.

6 Experiments

6.1 CCA on SPD Manifolds

Diffusion tensors are symmetric positive definite matrices at each voxel in DTI. Let $\text{SPD}(n)$ be a manifold for symmetric positive definite matrices of size $n \times n$. This forms a quotient space $GL(n)/O(n)$, where $GL(n)$ denotes the general linear group and $O(n)$ is the orthogonal group. The inner product of two tangent vectors $u, v \in T_p \mathcal{M}$ is given by $\langle u, v \rangle_p = \text{tr}(p^{-1/2} u p^{-1/2} v p^{-1/2})$. Here, $T_p \mathcal{M}$ is a tangent space at p (which is a vector space) is the space of symmetric matrices of dimension $(n+1)n/2$. The geodesic distance is $d(p, q)^2 = \text{tr}(\log^2(p^{-1/2} q p^{-1/2}))$.

Here, the exponential map and logarithm map are defined as,

$$\text{Exp}(p, v) = p^{1/2} \exp(p^{-1/2} v p^{-1/2}) p^{1/2}, \text{Log}(p, q) = p^{1/2} \log(p^{-1/2} q p^{-1/2}) p^{1/2} \quad (19)$$

and the first derivative of g in equation (11) on $\text{SPD}(n)$ is given by

$$\begin{aligned} \frac{d}{dt_i} g(t_i, w_x) &= \frac{d}{dt_i} \|\text{Log}(\text{Exp}(\mu_x, t_i W_x), X_i)\|^2 \\ &= \frac{d}{dt_i} \text{tr}[\log^2(X_i^{-1} S(t_i))] = 2 \text{tr}[\log(X_i^{-1} S(t_i)) S(t_i)^{-1} \dot{S}(t_i)], \text{ according to prop. 2.1 in [28]} \end{aligned} \quad (20)$$

where $S(t_i) = \text{Exp}(\mu_x, t_i W_x) = \mu_x^{1/2} \exp^{t_i A} \mu_x^{1/2}$, and $\dot{S}(t_i) = \mu_x^{1/2} A \exp^{t_i A} \mu_x^{1/2}$ and

$A = \mu_x^{-1/2} W_x \mu_x^{-1/2}$. The derivative of equality constraints, namely $\frac{d^2}{dW dt} g$, $\frac{d^2}{dt^2} g$ are calculated by numerical derivatives. Embedding the tangent vectors in the $n(n+1)/2$ dimensional space with orthonormal basis in the tangent space enables one to compute numerical differentiation. Details are provided in the extended paper.

6.2 Synthetic Experiments

In this section we provide experimental results using a synthetic dataset to evaluate the performance of Riemannian CCA. The samples are generated to be spread far apart on the manifold $\mathcal{M} (\equiv \text{SPD}(3))$ so that the curvature of the manifold plays a key role in the maximization of the correlation function. In order to sample data from different regions of the manifold, we generate data around two well separated means $\mu_{x1}, \mu_{x2} \in \mathcal{X}$, $\mu_{y1}, \mu_{y2} \in \mathcal{Y}$

by perturbing the data randomly (see the extended version) in the corresponding tangent spaces. Fig. 2 shows the CCA results obtained by Riemannian and Euclidean methods. We can clearly see the improvements from the manifold approach by inspecting the correlation coefficients $\rho_{x,y}$ on the respective titles.

6.3 CCA for Multi-modal Risk Analysis

Motivation—We collected multi-modal magnetic resonance imaging (MRI) data to investigate the effects of risk for Alzheimer’s disease (AD) on the white and gray matter in the brain. One of the central goals in analyzing this rich dataset is to find statistically significant AD risk $\leftrightarrow$ brain relationships. We can adopt many different ways of modeling these relationships but a potentially useful way is to analyze multi modality imaging data simultaneously, using CCA.

Risk for AD is characterized by their familial history (FH) status as well as APOE genotype risk factor. In the current experiments, we include a subset of 343 subjects and first investigate the effects of age and gender in a multimodal fashion since these variables are also important factors in healthy aging.

Brain structure is characterized by diffusion weighted images (DWI) for white matter and T1-weighted (T1W) image data for the gray matter. DWI data provides us information about the microstructure of the white matter. We use diffusion tensor ($\mathcal{D} \in \text{SPD}(3)$) model to represent the diffusivity in the microstructure. T1W data can be used to obtain volumetric properties of the gray-matter. The volumetric information is obtained from Jacobian matrices (J) of the diffeomorphic mapping to a population specific template. These Jacobian matrices can be used to obtain the Cauchy deformation tensors which also belong to $\text{SPD}(3)$.

Hippocampus and cingulum bundle (shown in Fig. 3) are two important regions in the brain. They are *a priori* believed to be significant in AD $\leftrightarrow$ brain structure relationships, primarily due to the role of hippocampus in memory function and the projections of cingulum onto the hippocampus. However, detecting *risk*-brain relationships *before* the memory/cognitive function is impaired is difficult due to several factors (such as noise in the data, small sample and effect sizes, type I error due to multiple comparisons.). One approach to improve the statistical power in such a setting would be to perform tests on average properties in regions of interest (ROI) in the brain. This procedure reduces both noise and the number of comparisons/tests. However, taking averages will also dampen the signal of interest which is already weak in such pre-clinical studies. CCA can take the multi-modal information from the imaging data and project the voxels into a space where the signal of interest is likely to be stronger.

Experimental Design—The key multimodal linear relations we examine are

$$Y_{\text{DTI}} = \beta_0 + \beta_1 \text{Gender} + \beta_2 X_{\text{T1W}} + \beta_3 X_{\text{T1W}} \cdot \text{Gender} + \varepsilon, \quad Y_{\text{DTI}} = \beta'_0 + \beta'_1 \text{AgeGroup} + \beta'_2 X_{\text{T1W}} + \beta'_3 X_{\text{T1W}} \cdot \text{AgeGroup} + \varepsilon,$$

where the AgeGroup is defined as a categorical variable with 0 (middle aged) if the age of the subject ≥ 65 and 1 (old) otherwise. The sample under investigation is between 43 and 75

years of age. The statistical tests ask if we can reject the Null hypotheses $\beta_3 = 0$ and $\beta'_3 = 0$ using our data at $\alpha = 0.05$. We report the results from the following four sets of analyses: **(i)** Classical ROI-average analysis: This is a standard type of setting where the brain measurements in an ROI are averaged. Here $Y_{DTI} = \overline{MD}$ i.e., the average mean diffusivity in the cingulum bundle. $X_{T1W} = \log |J|$ i.e., the average volumetric change (relative to the population template) in the hippocampus. **(ii)** Euclidean CCA using scalar measures (MD and $\log |J|$) in the ROIs: Here, the voxel data is projected using the classical CCA approach [35] i.e., $Y_{DTI} = w_{MD}^T MD$ and $X_{T1W} = w_{\log |J|}^T \log |J|$. **(iii)** Euclidean CCA using $\mathcal{D}$ and $\mathcal{J}$ in the ROIs: This setting is an improvement to the setting above in that the projections are performed using the full tensor data [35]. Here $Y_{DTI} = w_{\mathcal{D}}^T \mathcal{D}$ and $X_{T1W} = w_{\mathcal{J}}^T \mathcal{J}$. **(iv)** Riemannian CCA using $\mathcal{D}$ and $\mathcal{J}$ in the ROIs: Here $Y_{DTI} = \langle w_{\mathcal{D}}, \mathcal{D} \rangle_{\mu_{\mathcal{D}}}$ and $X_{T1W} = \langle w_{\mathcal{J}}, \mathcal{J} \rangle_{\mu_{\mathcal{J}}}$.

The findings are shown in Fig. 4. We can see that the performance of CCA using the full tensor information improves the statistical significance for both Euclidean and Riemannian approaches. The weight vectors in the different settings for both Euclidean and Riemannian CCA are shown in Fig. 5 top row. We would like to note that there are several different approaches of using the data from CCA and we performed experiments with full gray matter and white matter regions in the brain whose results are included in the extended version. We show the representative weight vectors (in Fig. 5 bottom row) obtained using the full brain analyses. Interestingly, the weight vectors are spatially cohesive even without enforcing any spatial constraints. What is even more remarkable is that the regions picked between the DTI and T1W modalities are complimentary in a biological sense. Specifically, when performing our CCA on the ROIs, although the cingulum bundle extends into the superior mid-brain regions the weights are non-zero in its hippocampal projections. In the case of entire white and gray matter regions, the volumetric difference (from the population template) in the inferior part of the corpus callosum seem to be highly cross-correlated to the diffusivity in the corpus callosum. Our CCA finds these projections without any a priori constraints in the optimization suggesting that performing CCA on the intrinsic nature of the data can reveal biologically meaningful patterns. Due to space constraints, we refer the interested reader to the extended version of the paper for additional details.

7 Conclusion

The classical CCA assumes that data live in a pair of vector spaces. However, many modern scientific disciplines require the analysis of data which belong to *curved* spaces where classical CCA is no longer applicable. Motivated by the properties of imaging data from neuroimaging studies, we generalize CCA to Riemannian manifolds. We employ differential geometry tools to extend operations in CCA to the manifold setting. Such a formulation results in a multi-level optimization problem. We derive solutions using the first order condition of projection and an augmented Lagrangian method. In addition, we also develop an efficient single path algorithm with approximate projections. Finally, we propose a generalization to the product space of $SPD(n)$, namely, tensor fields allowing us to treat a full brain image as a point on the product manifold. On the experimental side, we presented

neuroimaging findings using our proposed CCA on DTI and T1W imaging modalities on an Alzheimer's disease (AD) dataset focused on risk factors for this disease. Here, the proposed methods perform well and yield scientifically meaningful results. In closing, we note that our core optimization methods can be readily applied when maximizing correlation between data from two *different* types of Riemannian manifolds — this may open the doors to various other types of analysis not explicitly investigated in this paper.

Acknowledgments

This work was supported in part by NIH grants AG040396 (VS), AG037639 (BBB), AG021155 (SCJ), AG027161 (SCJ), NS066340 (BCV) and NSF CAREER award 1252725 (VS). Partial support was also provided by UW ADRC, UW ICTR, and Waisman Core grant P30 HD003352-45. The contents do not represent views of the Dept. of Veterans Affairs or the United States Government.

References

1. Akaho, S. International Meeting on Psychometric Society; 2001. A kernel method for canonical correlation analysis.
2. Andrew, G.; Arora, R.; Bilmes, J.; Livescu, K. Deep canonical correlation analysis. ICML; 2013.
3. Avants BB, Cook PA, Ungar L, Gee JC, Grossman M. Dementia induces correlated reductions in white matter integrity and cortical thickness: a multivariate neuroimaging study with SCCA. *Neuroimage*. 2010; 50(3):1004–1016. [PubMed: 20083207]
4. Bach FR, Jordan MI. Kernel independent component analysis. *The Journal of Machine Learning Research*. 2003; 3:1–48.
5. Callaghan, PT. Principles of nuclear magnetic resonance microscopy. Oxford University Press; 1991.
6. Chung M, Worsley K, Paus T, Cherif C, Collins D, Giedd J, Rapoport J, Evans A. A unified statistical approach to deformation-based morphometry. *NeuroImage*. 2001; 14(3):595–606. [PubMed: 11506533]
7. Do Carmo MP. Riemannian geometry. 1992
8. Fletcher PT. Geodesic regression and the theory of least squares on riemannian manifolds. *International Journal of Computer Vision*. 2013; 105(2):171–185.
9. Fletcher PT, Lu C, Pizer SM, Joshi S. Principal geodesic analysis for the study of nonlinear statistics of shape. *Medical Imaging*. 2004; 23(8):995–1005. [PubMed: 15338733]
10. Goh A, Lenglet C, Thompson PM, Vidal R. A nonparametric Riemannian framework for processing high angular resolution diffusion images (HARDI). 2009:2496–2503.
11. Hardoon DR, Szedmak S, Shawe-Taylor J. CCA: An overview with application to learning methods. *Neural Computation*. 2004; 16(12):2639–2664. [PubMed: 15516276]
12. Hardoon DR, et al. Unsupervised analysis of fMRI data using kernel canonical correlation. *NeuroImage*. 2007; 37(4):1250–1259. [PubMed: 17686634]
13. Hinkle J, Fletcher PT, Joshi S. Intrinsic polynomials for regression on Riemannian manifolds. *Journal of Mathematical Imaging and Vision*. 2014:1–21.
14. Ho, J.; Xie, Y.; Vemuri, B. On a nonlinear generalization of sparse coding and dictionary learning. ICML; 2013. p. 1480-1488.
15. Hotelling H. Relations between two sets of variates. *Biometrika*. 1936; 28(3/4):321–377.
16. Hsieh WW. Nonlinear canonical correlation analysis by neural networks. *Neural Networks*. 2000; 13(10):1095–1105. [PubMed: 11156190]
17. Hua X, Gutman B, et al. Accurate measurement of brain changes in longitudinal MRI scans using tensor-based morphometry. *Neuroimage*. 2011; 57(1):5–14. [PubMed: 21320612]
18. Hua X, Leow AD, Parikshak N, Lee S, Chiang MC, Toga AW, Jack CR Jr, Weiner MW, Thompson PM. Tensor-based morphometry as a neuroimaging biomarker for Alzheimer's disease:

- an MRI study of 676 AD, MCI, and normal subjects. *Neuroimage*. 2008; 43(3):458–469. [PubMed: 18691658]
19. Huang H, He H, Fan X, Zhang J. Super-resolution of human face image using canonical correlation analysis. *Pattern Recognition*. 2010; 43(7):2532–2543.
 20. Huckemann S, Hotz T, Munk A. Intrinsic shape analysis: Geodesic PCA for Riemannian manifolds modulo isometric Lie group actions. *Statistica Sinica*. 2010; 20:1–100.
 21. Jayasumana, S.; Hartley, R.; Salzmann, M.; Li, H.; Harandi, M. Kernel methods on the Riemannian manifold of symmetric positive definite matrices. *CVPR*; 2013. p. 73-80.
 22. Karcher H. Riemannian center of mass and mollifier smoothing. *Communications on Pure and Applied Mathematics*. 1977; 30(5):509–541.
 23. Kim, HJ.; Adluru, N.; Collins, MD.; Chung, MK.; Bendlin, BB.; Johnson, SC.; Davidson, RJ.; Singh, V. Multivariate general linear models (MGLM) on Riemannian manifolds with applications to statistical analysis of diffusion weighted images. *CVPR*; 2014.
 24. Kim TK, Cipolla R. CCA of video volume tensors for action categorization and detection. *PAMI*. 2009; 31(8):1415–1428.
 25. Lai PL, Fyfe C. A neural implementation of canonical correlation analysis. *Neural Networks*. 1999; 12(10):1391–1397. [PubMed: 12662622]
 26. Lebanon, G., et al. Ph.D. thesis. Carnegie Mellon University, Language Technologies Institute, School of Computer Science; 2005. Riemannian geometry and statistical machine learning.
 27. Li, P.; Wang, Q.; Zuo, W.; Zhang, L. Log-Euclidean kernels for sparse representation and dictionary learning. *ICCV*; 2013. p. 1601-1608.
 28. Moakher M. A differential geometric approach to the geometric mean of symmetric positive-definite matrices. *SIAM Journal on Matrix Analysis and Applications*. 2005; 26(3):735–747.
 29. Niethammer, M.; Huang, Y.; Vialard, FX. Geodesic regression for image timeseries. *MICCAI*; 2011. p. 655-662.
 30. Said, S.; Courty, N.; Le Bihan, N.; Sangwine, SJ., et al. Exact principal geodesic analysis for data on $SO(3)$; *Proceedings of the 15th European Signal Processing Conference*; 2007. p. 1700-1705.
 31. Shi, X.; Styner, M.; Lieberman, J.; Ibrahim, JG.; Lin, W.; Zhu, H. Intrinsic regression models for manifold-valued data. In: Yang, G-Z.; Hawkes, D.; Rueckert, D.; Noble, A.; Taylor, C., editors. *MICCAI 2009, Part II. LNCS. Vol. 5762*. Springer, Heidelberg; 2009. p. 192-199.
 32. Sommer, S. Horizontal dimensionality reduction and iterated frame bundle development. In: Nielsen, F.; Barbaresco, F., editors. *GSI 2013. LNCS. Vol. 8085*. Springer, Heidelberg; 2013. p. 76-83.
 33. Sommer S, Lauze F, Nielsen M. Optimization over geodesics for exact principal geodesic analysis. *Advances in Computational Mathematics*. :1–31.
 34. Steinke F, Hein M, Schölkopf B. Nonparametric regression between general Riemannian manifolds. *SIAM Journal on Imaging Sciences*. 2010; 3(3):527–563.
 35. Witten DM, Tibshirani R, Hastie T. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*. 2009; 10(3):515–534. [PubMed: 19377034]
 36. Xie, Y.; Vemuri, BC.; Ho, J. Statistical analysis of tensor fields. In: Jiang, T.; Navab, N.; Pluim, JPW.; Viergever, MA., editors. *MICCAI 2010, Part I. LNCS. Vol. 6361*. Springer, Heidelberg; 2010. p. 682-689.
 37. Yger, F.; Berar, M.; Gasso, G.; Rakotomamonjy, A. Adaptive canonical correlation analysis based on matrix manifolds. *ICML*; 2012.
 38. Yu S, Tan T, Huang K, Jia K, Wu X. A study on gait-based gender classification. *IEEE Transactions on Image Processing*. 2009; 18(8):1905–1910. [PubMed: 19447706]
 39. Zhu H, Chen Y, Ibrahim JG, Li Y, Hall C, Lin W. Intrinsic regression models for positive-definite matrices with applications to diffusion tensor imaging. *Journal of the American Statistical Association*. 2009; 104(487)

Fig. 1.

CCA on Riemannian manifolds. CCA searches geodesic submanifolds (subspaces), S_{w_x} and S_{w_y} at the Karcher mean of data on each manifold. Correlation between projected points

$\{\Pi_{S_{w_x}}(x_i)\}_{i=1}^N$ and $\{\Pi_{S_{w_y}}(y_i)\}_{i=1}^N$ is equivalent to the correlation between *projection coefficients* $\{t_i\}_{i=1}^N$ and $\{u_i\}_{i=1}^N$. Although $\mathbf{x}$ and $\mathbf{y}$ belong to the same manifold we show them in different plots for ease of explanation.

Fig. 2.

Synthetic experiments showing the benefits of Riemannian CCA. The top row shows the projected data using the Euclidean CCA and the bottom using Riemannian CCA. P_X and P_Y denote the projected axes. Each column represents a synthetic experiment with a specific set of $\{\mu_{x_j}, \varepsilon_{x_j}; \mu_{y_j}, \varepsilon_{y_j}\}$. The first column presents results with 100 samples while the three columns on the right show with 1000 samples. The improvements in the correlation coefficients $\rho_{X,Y}$ can be clearly seen from the corresponding titles.

Fig. 3. Shown on the left are the bilateral cingulum bundles (green) inside a brain surface obtained from a population DTI template. Similarly on the right are the bilateral hippocampi. The gray and white matter ROIs are also shown on the right.

Fig. 4.

Experimental evidence showing the improvements in statistical significance of finding the multi-modal risk-brain interaction effects. Top row shows the gender, volume and diffusivity interactions. Second row shows the interaction effects of the middle/old age groups.

Fig. 5.

Weight vectors (in red-yellow color) obtained from our Riemannian CCA approach. The weights are in arbitrary units. The top row is from applying Riemannian CCA on data from the cingulum and hippocampus ROIs (Fig. 3) while the bottom row is obtained using data from the entire white and gray matter regions of the brain. On the left (three columns) block we show the results in orthogonal view for DTI and on the right for T1W. The corresponding underlays are the population averages of the fractional anisotropy and T1W contrast images respectively.