

On Characterizing and Analyzing Diffusion Tensor Images by Learning their Underlying Manifold Structure

Parmeshwar Khurd

Ragini Verma

Christos Davatzikos

Dept. of Radiology, University of Pennsylvania, Philadelphia, PA 19104

khurdp@uphs.upenn.edu

Abstract

The growing importance of diffusion tensor imaging (DTI) in studying the white matter architecture in normal and pathologic states necessitates the development of tools for comprehensive analysis of diffusion tensor data. Operations such as multivariate statistical analysis and hypothesis testing, interpolation and filtering, must now be performed on tensor data, and must overcome challenges introduced by the non-linearity and high dimensionality of the tensors. In this paper, we present a novel approach to performing these computations by modeling the underlying manifold structure of the tensors, using a combination of two manifold learning techniques, isometric mapping (ISOMAP) and local tangent space alignment (LTSA). While ISOMAP identifies the dimensionality of the manifold of the tensors and embeds the tensors into a linear space, facilitating statistical computations therein, operations like interpolation and filtering, integral to the process of normalization, require the reconstruction of the tensor in the tensor domain. To obtain this reverse mapping from the linear space to the tensor domain, i.e. to the domain of the original tensor data, we use LTSA. The modeling of the underlying manifold structure renders our approach better applicable to tensor data than existing methods that may not always be able to capture the non-linearity present in the tensors under consideration. In various simulations with known ground truth, we demonstrate the effectiveness of our framework based on ISOMAP and LTSA in performing a comprehensive analysis of DTI data.

1. Introduction

Diffusion tensor imaging (DTI) has become an important modality to study brain white matter architecture in normal and pathological states, for diseases such as schizophrenia [9]. Such studies involve a multivariate statistical analysis of the anatomical structures in diffusion tensor images of

diseased and normal brains. As in the group-wise analysis of structural MR images [12], studies of this nature need to employ a multi-faceted analysis paradigm requiring spatial normalization of all DTI images to a brain template (which in turn requires tensor interpolation and filtering) and a statistical analysis involving hypothesis testing on tensors. Since diffusion tensors are symmetric positive definite matrices, the tensor domain is non-linear and conventional vector space techniques cannot be used to perform operations such as interpolation, filtering and hypothesis testing on tensors. Early studies [9] circumvented this problem by using scalar measures such as fractional anisotropy (FA) to quantify white matter structure. However, such scalar measures do not account for the full information present in the tensors and require *a priori* knowledge of how pathology affects these measures. A vector space approach which embeds the six components of the diffusion tensor into \mathbb{R}^6 and uses linear operations such as principal component analysis (PCA) also gets stymied by the nonlinear nature of tensors. Hence several groups have recently derived methods using a variety of tensor metrics [7, 10, 2, 1, 13], such as those based upon Riemannian symmetric spaces [7, 10, 2], to perform the requisite computations on tensors. But these approaches may not always be able to capture the specific non-linearities present in the tensors under consideration.

In this paper, we present a novel approach to performing these tensor computations by modeling the underlying manifold structure of the tensors using manifold learning techniques [4]. We use a combination of two manifold learning techniques, isometric mapping (ISOMAP) [11] and local tangent space alignment (LTSA) [15] to learn the manifold structure present in the tensors and then use this manifold structure to perform the desired tensor computations. We use ISOMAP to determine the dimensionality of the underlying manifold structure present in the tensors and to embed them into a linear space. Most statistical computations can then be readily performed in this linear space. However, certain computations such as interpolation and filtering require us to return to the tensor domain, and hence an in-




Figure 1. In the top figure, tensors (denoted by ellipses) lie on a non-linear manifold embedded in \mathbb{R}^6 . The geodesic distance is shown by a red curve. The green dotted line denotes the linear distance between tensors treated as Euclidean entities. The linear distance is not along the manifold and hence not an appropriate measure of comparison among tensors. In the bottom figure, a flattened vector-space embedding of the tensors is shown along with new elements created in the lower dimensional space. This embedding may be obtained via ISOMAP or LTSA. LTSA needs the manifold dimensionality d from ISOMAP to obtain the embedding. The dotted arrows indicate that LTSA can also map an arbitrary point in the embedded space back into the tensor domain.

verse of the mapping found in the embedding operation is needed. Since ISOMAP does not provide such an inverse mapping, we use a different manifold learning technique - LTSA. However, in order to obtain an embedding, LTSA itself requires the manifold dimensionality, as identified by ISOMAP, as one of its input parameters. In Sec. 2, we will discuss the manifold learning techniques, ISOMAP and LTSA, and in Sec. 3, we will show how they can be used to comprehensively analyze DTI data.

The modeling of the underlying manifold structure renders our approach superior to existing methods that employ various tensor metrics [7, 10, 1, 2]. In Sec. 4, we have conducted various simulations with known ground truth through which we demonstrate how ISOMAP and LTSA can be effectively used to perform a comprehensive analysis of diffusion tensor images. We have also conducted a limited comparison with the approaches in [7, 10, 1, 2].

2. Tensor Manifold Learning

To motivate our methods, we shall first consider some simple intuitive examples of how tensor non-linearity can affect the calculation of tensor distances, tensor averages,




Figure 2. (a) Collection of tensors whose statistical average is to be determined (b) The underlying manifold structure of these tensors is identified as one dimensional (marked M) as these are created by rotating the principal eigenvector of the tensor A, from 0 to 180 degrees. The ellipse C is the linear average and the ellipse B is the average on the manifold.

etc. Consider the tensors depicted at the top of Fig. 1. Distances between tensors have to be measured via geodesic distances (along the manifold) rather than Euclidean distances (outside the manifold), as shown in Fig. 1. (The red curve along the manifold represents the geodesic distance while the green line represents the Euclidean distance). Fig. 2 shows a synthesized example that emphasizes this aspect. Although tensor data might appear extremely complex and randomly oriented (Fig 2(a)), it may have an underlying low dimensional structure, as can be seen in Fig. 2(b). The tensors in Fig. 2(a) are actually generated by rotating the principal eigen vector of tensor A (Fig. 2(b)). These six-dimensional tensors, therefore, lie on a one dimensional sub-manifold depicted by the dotted arc M generated by the variation of angle of rotation of the principal eigen vector. If we were to form the mean of these tensors via linear averaging, which is effectively based on Euclidean distances, we would find a mean tensor that lies outside of the manifold generated by the angle of rotation (ellipse C in Fig. 2(b)). The average along the manifold is the tensor B in Fig. 2(b) which lies along this manifold and is obtained based on geodesic distances along the manifold. Therefore, the primary issue here is fitting a manifold to the tensors and determining its dimensionality.

In subsections 2.1 and 2.2, we shall describe how ISOMAP and LTSA can be applied to obtain the manifold structure present in a set of tensors. Fig. 1 graphically illustrates the relationships between the underlying manifold, ISOMAP and LTSA. In section 3, we later describe how ISOMAP and LTSA can be used to solve specific problems that arise in the analysis of diffusion tensor images. We shall then see that ISOMAP is especially useful in group-wise analysis of DTI data and that LTSA is especially useful in interpolation and filtering of DT images.

A brief note on mathematical notation: We shall denote vectors by bold-faced lower case letters, e.g. \mathbf{x} , and matrices by upper-case letters, e.g. A . We use \mathbf{e} to denote the vector of all 1's and I to denote the identity matrix. We will use the superscripts T and \dagger to denote the matrix transpose

and the matrix pseudo-inverse respectively. We denote the sample mean of a set of vectors $\{\mathbf{x}_i, i = 1, \dots, K\}$ by $\bar{\mathbf{x}}$.

Let $D_i, i = 1, \dots, N$ be the dataset of tensors. We represent each tensor as a six dimensional vector \mathbf{x}_i , with the six components being the six upper triangular elements of D_i . The aim is to be able to find the underlying manifold for the set $\mathbf{x}_i, i = 1, \dots, N$ and produce the corresponding embedded set of vectors $\mathbf{y}_i, i = 1, \dots, N$.

2.1. ISOMAP

ISOMAP learns the dimensionality and embedding for the manifold by following three steps:

1. *Construction of a tensor neighborhood graph:* This use a tensor-based distance metric to determine neighboring points in the tensor dataset and creates a graph with the tensors as nodes. By “neighbors” here we dont mean spatial neighbors, but neighboring samples in \mathbb{R}^6 . For any two tensors D_i and D_j , we use the Frobenius norm of $D_i - D_j$ as the tensor-based distance between them. (We found that using the Frobenius norm of $D_i - D_j$ yields better performance than using the Euclidean distance between \mathbf{x}_i and \mathbf{x}_j .) The size of the neighborhood at each point is determined from user-specified parameters.
2. *Calculating the geodesic distance and creation of Distance Matrix:* The geodesic distance between two far-away tensors on the graph is estimated by the graph based shortest distance between nodes, which is computed using the Floyd-Warshall algorithm or Dijkstra’s algorithm. A distance matrix relating all tensors is then readily obtained.
3. *Fitting the manifold:* Multidimensional scaling (MDS) is used to extract the manifold fitted to the tensors. In the eigen decomposition of the distance matrix, let $\lambda_i, i = 1, \dots, N$ be the eigen values of the $N \times N$ distance matrix in descending order of magnitude and $\mathbf{v}_i, i = 1, \dots, N$ be the corresponding eigen vectors. The dimension d of the underlying manifold is equal to the rank of the centered distance matrix. The d eigen vectors corresponding to the d largest eigenvalues are then chosen to obtain the representation of these tensors in the d - dimensional space. Let λ_p be the p -th eigen value (in decreasing order) of the distance matrix and let v_p^i be the i -th component of the p -th eigen vector. Then for the i -th sample point, we set the p -th component of the d -dimensional coordinate vector equal to $\sqrt{\lambda_p} v_p^i$ and thus obtain the representation on the fitted manifold, that is $\mathbf{y}_i = [\sqrt{\lambda_1} v_1^i, \dots, \sqrt{\lambda_d} v_d^i]^T$.

Rather than estimate the dimensionality of the manifold using the rank of the distance matrix, the following simpler and more robust method is usually used in conjunction with

ISOMAP: Let the residual variance or the fraction of the “unexplained residuals” as captured by a set of d eigen vectors is defined as : $RV(d) = \frac{\sum_{i=d+1}^N \lambda_i}{\sum_{i=1}^N \lambda_i}$. This provides a measure of the information captured by the data. Visually, the dimension can be found by plotting the number of vectors used to create the flattened manifold versus the residual variance and identifying the value of d at which the residual variance levels off. From a visual point of view, this is the point at which an elbow appears in the graph and this value of d is the dimensionality of the data. If no elbow is present, then the manifold dimensionality is equal to 1, as shown in Fig. 3(d).

While ISOMAP embeds the tensors into a linear space, on which statistical operations can be performed, some operations like filetring require a way of reconstructing a tensor corresponding to a point in the lower dimensional linear space, that is, a mapping from the embedding space back to \mathbb{R}^6 . We use LTSA to provide this capability.

2.2. LTSA

LTSA requires that the dimensionality d of the underlying embedding be known in advance. We shall use ISOMAP to determine the manifold dimensionality, but we note that it can also be obtained or by a host of other techniques [5]. LTSA obtains the embedding in two steps:

1. *Estimating the local tangent spaces:* Based on the Euclidean distance between \mathbf{x}_i and \mathbf{x}_j , a set of neighboring points is determined around each point \mathbf{x}_i . Let this neighborhood be denoted by the matrix $X_i = [\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_k}]$, where as in ISOMAP, the neighborhood size k is determined from user-specified parameters. Using local PCA, we may estimate the tangent space at \mathbf{x}_i such that

$$\mathbf{x}_{i_j} = \bar{\mathbf{x}}_i + Q_i \mathbf{z}_j^i + \boldsymbol{\zeta}_j^i \quad (1)$$

where the matrix Q_i consists of the d left singular eigenvectors of $X_i(I - \frac{\mathbf{e}\mathbf{e}^T}{k})$ corresponding to the d largest singular values and the local coordinate vectors \mathbf{z}_j^i are given by $\mathbf{z}_j^i = Q_i^T(\mathbf{x}_{i_j} - \bar{\mathbf{x}}_i)$. The term $\boldsymbol{\zeta}_j^i$ accounts for residual error.

2. *Aligning the local tangent spaces:* In order to align the tangent spaces, two simplifying assumptions are now made. First it is assumed that the manifold is a d - dimensional surface in \mathbb{R}^6 . Therefore, it can be represented with a 1-1 mapping $f : \Omega \subset \mathbb{R}^d \rightarrow \mathbb{R}^6$. Secondly, it is assumed that the global coordinates \mathbf{y}_i in Ω and the local coordinates in the tangent space are related by an affine mapping:

$$\mathbf{y}_{i_j} = \bar{\mathbf{y}}_i + L_i \mathbf{z}_j^i + \boldsymbol{\epsilon}_j^i \quad (2)$$

where ϵ_j^i accounts for residual error. Let $Z_i = [\mathbf{z}_1^i, \dots, \mathbf{z}_k^i]$, $Y_i = [\mathbf{y}_{i1}, \dots, \mathbf{y}_{ik}]$ and $Y = [\mathbf{y}_1, \dots, \mathbf{y}_N]$. The matrices L_i and the global coordinates \mathbf{y}_i are then determined so as to minimize the sum of all squared errors $\sum_i \sum_j (\epsilon_j^i)^T \epsilon_j^i$, subject to the constraint $YY^T = I_d$.

Let S_i be a 0-1 matrix such that $YS_i = Y_i$, let $W_i = (I - \frac{\mathbf{e}\mathbf{e}^T}{k})(I - Z_i^\dagger Z_i)$ and let $B = \sum_{i=1}^N S_i W_i W_i^T S_i^T$. It turns out that the optimal Y is given by the d eigenvectors of B corresponding to the 2nd to $(d+1)$ st smallest eigenvalues, and that the optimal L_i is given by $L_i = Y_i Z_i^\dagger$.

Let \mathbf{y} be an arbitrary point in the LTSA embedded space. In order to obtain its representation \mathbf{x} in \mathcal{R}^6 , we first determine the closest point \mathbf{y}_i in the current embedding. Then $\mathbf{x} = \bar{\mathbf{x}}_i + Q_i L_i^{-1}(\mathbf{y} - \bar{\mathbf{y}}_i)$. We cannot mathematically prove that the reconstructed tensor D corresponding to \mathbf{x} is positive-definite, but we expect this condition to hold true since LTSA should be able to learn the non-linearities introduced into the manifold by positive-definiteness.

3. Applications in Analysis of Diffusion Tensor Images

We shall now describe how the techniques in 2 can be applied to the problem of group-wise statistical analysis of tensor images and used to address the issues of interpolation and filtering of tensor images.

3.1. Group-wise statistical analysis

In doing group wise analysis, images from multiple individuals are spatially normalized [14] to a template, so that the subsequent statistical analysis can be performed in the same coordinate system. (In the spatial normalization approach in [14], non-rigid registration is performed on the FA maps and hence issues such as tensor interpolation do not arise. However, tensor interpolation is an important problem in more general approaches to registration and we shall address it in the next subsection.) The template can be the scan of a different individual or the same individual at a different time point depending on whether we are analyzing for group differences or studying the progression of disease. Once these images are spatially normalized, we may analyze these images by collecting tensors voxel-wise or region-wise.

Voxel-wise Analysis: We first describe a voxel-wise analysis. Let $D_i, D_i \in \text{Image}(i), i = 1, \dots, N$ be the tensors from N individuals at a given voxel location. The structure of the manifold on which these tensors lie is determined through the procedure of Section 2.1. The application of ISOMAP on a voxel-by-voxel basis flattens the high dimensional non-linear tensor data to a Euclidean submanifold, transforms the tensors to their linear counterparts

and defines a geodesic distance on this submanifold, as the shortest distance between points on the graph. This flattened dataset can now be used for statistical analysis. ISOMAP is therefore applied to capture the manifold structure on a voxel-by-voxel basis, collected across the whole population. Multivariate statistics can now be applied to the tensors embedded in to the Euclidean manifold. We apply the Hotelling T^2 -test for separation and equality of means of two tensor data sets which are embedded into manifolds of dimension greater than one. It may be noted that if the dimension of the lower dimensional manifold is 1, t-test can be used for significance. The application of this test requires the entries to be normally distributed, which will not be the case for the original tensors, which lie on a high dimensional non linear space. However on flattening to a Euclidean space, these generally tend to follow a normal distribution, for the kind of tensor data we work with. In case ISOMAP finds more than one graph component at a particular voxel, then we may apply the Hotelling T^2 -test to the 6-d vector formed by the symmetric matrix logarithm of the tensors at that voxel [1]. Neighborhood information can also be included for each voxel, to obtain a more robust manifold structure.

Region-based Analysis: In region-wise analysis, we start with a Region of Interest (ROI) where we expect to see significant differences between normal and abnormal subjects. Such an ROI may be known *a priori* or could be identified using voxel-wise analysis. Given this ROI, we may then embed all the tensors in this region across all subjects into a common ISOMAP embedding. A single Hotelling T^2 statistic can be computed to quantitatively summarize the differences between the two groups. In the case when the dimensionality of the ISOMAP embedding is ≤ 3 , we may visualize whether the tensors belonging to normal and abnormal subjects form two separate clusters. One advantage of having this common embedding is that given an additional normalized DTI of a subject, we may then select the tensors belonging to the ROI with significant differences and embed them using ISOMAP. The method in [3] can be used to efficiently obtain such an out-of-sample ISOMAP embedding. If all tensors belonging to the subject get embedded into a specific cluster (either from visual inspection, or as measured, for example, by the Hausdorff distance), then we can readily classify the new subject as normal or abnormal. The common lower dimensional ISOMAP embedding could also form the features for a classifier based upon supervised learning.

3.2. Interpolation and Filtering

We adopt the following approach to interpolation and smoothing:

1. Find the dimensionality of the manifold structure in all tensors belonging to the image using ISOMAP.

2. Find a lower-dimensional vector space embedding for all tensors using LTSA.
3. Interpolate or filter the representation in the vector space using a standard technique. For instance, we may perform trilinear interpolation on each component of the vector embedding.
4. Return to the tensor domain using the LTSA mapping.

A potential difficulty with this approach is that the manifold learning approaches will be overwhelmed by the large number of tensors (tensors from all voxels) presented to them. However, for methods such as ISOMAP, efficient methods based on the Nystrom approximation have been devised to deal with this problem [4, 3, 8]. To our knowledge, such efficient methods have not been used in conjunction with LTSA. We plan to investigate their applicability to LTSA in future work.

One could also adopt an alternative atlas-based approach in conjunction with manifold learning. We could non-rigidly register the given tensor image with a segmented tensor atlas, divide the tensors into classes as per the labels in the atlas, and then separately apply the method described in the earlier paragraph to the tensors belonging to each class. This procedure has three advantages over the more straightforward approach: (1) It will better maintain the local structure of the tensors. (2) The manifold learning techniques will be presented with fewer tensors. (3) Since the tensors belonging to a particular class may lie on a manifold of lower dimensionality, the interpolation or filtering in the embedded domain will also be faster.

Having described the methods for applying the manifold learning techniques in 2 comprehensively analyze DTI data, we shall now validate these techniques in a series of experiments with simulated and real data in Sec. 4.

4. Results

We shall illustrate a variety of ways in which manifold learning can be successfully used to analyze diffusion tensor images. In subsection 4.1, we describe some simple validation experiments where the manifold on which tensors reside is exactly known. In subsection 4.2, we perform statistical analysis on tensor images with simulated group differences and validate that our technique, as described in 3.1, agrees with the ground truth. In subsection 4.3, we compare the quality of our LTSA-based tensor interpolation approach described in 3.2 with alternative approaches.

4.1. Analysis on Tensors with Known Manifold Structure

In one experiment, we started with a typical elliptical diffusion tensor and increased the azimuthal angle corresponding to the principal direction in small increments so as to generate a set of tensors. (It may be noted that we

present results on 2D tensors here, for the ease of visualization. The experiment has been conducted on 3D tensors with the same degree of success). Thus the set of tensors reside on a known 1-d manifold (curve) as shown in Fig. 3(a). Through this experiment, we sought to answer the following questions: (1) Can ISOMAP correctly identify the dimensionality of the underlying manifold? (2) Do the ISOMAP and LTSA embeddings agree with the points lying on the original 1-dimensional manifold, i.e. the various azimuthal angles? (3) Can a linear technique such as principal component analysis (PCA) learn the underlying manifold structure? (4) Can LTSA-based interpolation be used to obtain a finer sampling of the original manifold?

We identified the manifold dimensionality using ISOMAP and obtained lower-dimensional embeddings using ISOMAP and LTSA. The lack of elbow in the ISOMAP residual variance versus dimensionality plot in Fig. 3(d) indicates a dimensionality of 1. This is further evident in Fig. 3(c) where we plot the angular change against the 1d embedding produced by ISOMAP and LTSA. We see that both the LTSA and ISOMAP embeddings (they overlap with each other in Fig. 3(c)) agree well with the angles used to generate these tensors. The standard linear approach of PCA fails to identify the manifold dimensionality, as evidenced in Fig. 3(e) by the fact that more than one PCA eigen value is significant, and the largest principal component does not correlate well with the azimuthal angles.

We sorted the LTSA 1-d embedding in increasing order and linearly interpolated it. We then mapped the interpolated LTSA embedding back to the tensor domain. Results are shown in Fig. 3(b). The LTSA interpolation also reproduces a fairly accurate finely sampled version of the original manifold. We note that the LTSA interpolation remained unchanged even if we randomly permuted the original tensors so that they no longer correspond to increasing azimuthal angles. Straightforward linear approaches to tensor interpolation or approaches based upon Riemannian symmetric spaces would yield similar interpolation results if the tensors are presented in the order of increasing (or decreasing) azimuthal angles, but given a random permutation of the tensors, they produce an incorrect interpolation result. This indicates that LTSA determines and incorporates the underlying manifold structure to produce the correct interpolation.

4.2. Groupwise Statistical Analysis

To validate our approach in 3.1, we constructed two sets of DTIs with known statistical changes in the following manner: We obtained a set of $256 \times 256 \times 56$ DTI images of 10 healthy subjects and normalized them to a common template (chosen as one of the subjects) [14]. We selected an ROI on the template and mapped it to the original 10 DTI images. We then introduced a known i.i.d. statistical

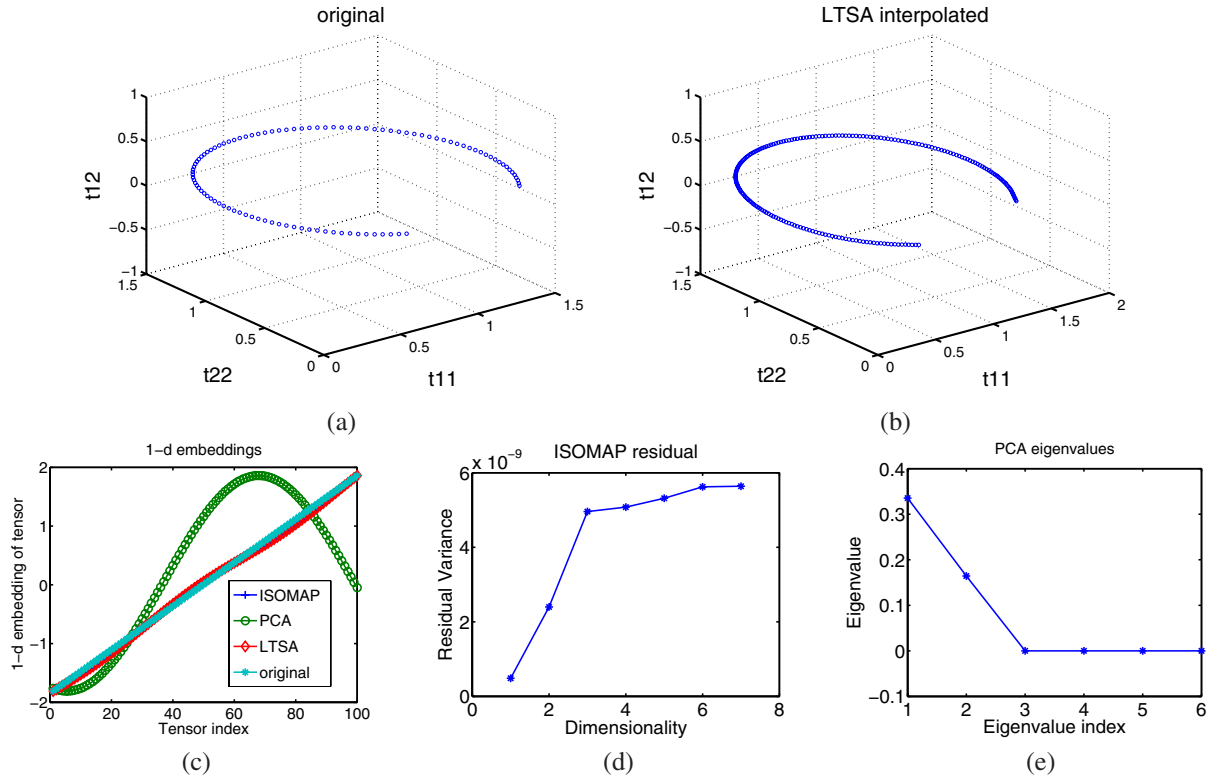


Figure 3. (a) The three variable components ($t11, t12, t22$) of the original tensors are shown as points in \mathbb{R}^3 (b) The three variable components of the interpolated tensors are shown as points in the same domain (c) The 1-d PCA (green), LTSA (red) and ISOMAP (blue) embeddings are compared with the original angles (cyan) used to generate the tensors. LTSA and ISOMAP embeddings overlap almost exactly with each other and match the original angles quite well, but the largest PCA component does not correlate well with the original angles. (d) ISOMAP residual variance vs. dimensionality plot indicates a dimensionality of 1 due to absence of an elbow. (e) PCA eigenvalues.

change in each tensor belonging to the 10 ROI's and then re-normalized the resulting abnormal DTIs to the template. We thus end up with a dataset with 20 normalized DTIs, 10 DTIs corresponding to the original healthy subjects and 10 abnormal DTIs. We simulated three types of statistical changes in the ROI: (1) We changed the three eigen values. (2) We applied a random rotation to the three eigenvectors. (3) We changed the three eigen values and also applied random rotations to the three eigenvectors.

We analyzed each of three datasets using ISOMAP in two ways. In the first approach, we analyzed each voxel point-wise using ISOMAP, i.e. we collected 20 tensors at each voxel corresponding to each subject, obtained their low-dimensional embedding using ISOMAP and then obtained a p-value using the Hotelling T^2 test. Fig. 4 shows the results obtained using this approach. Clearly the results agree very well with ground truth. In all cases, ISOMAP identified the dimensionality of the underlying manifold as 3.

In the second approach, we took a smaller ROI than the one determined using the first approach and embedded all tensors in that ROI across all images into a common linear space using ISOMAP. As in the voxel-wise analysis,

ISOMAP identified the dimensionality of the underlying manifold as 3 in all cases. Fig. 5(a)-(c) shows these 3-dimensional embeddings. Clearly, the normal and abnormal vectors are well separated in this common space. A Hotelling T^2 test applied to the two groups of normal and abnormal tensor embeddings in this linear space then gave us a single scalar measure to summarize group differences. Since the embeddings clearly separated the groups in all cases, the T^2 P-value was close to 0 in all cases. For computational reasons, we selected only 5 normal and 5 abnormal DTIs in this second approach to group analysis. We also validated the concept of using out-of-sample extension on tensors from the ROI of an additional subject. Of the earlier 5 subjects, we first obtained an ISOMAP embedding for 4 normal and 4 abnormal subjects and then computed out-of-sample extensions for the remaining normal and abnormal subject. Results are shown in fig. 5(d). The out-of-sample tensors from the new normal and abnormal subjects fall into their respective clusters.

4.3. Interpolation

We used the the first dataset from the validation experiments of Sec. 4.2 to design the following validation exper-

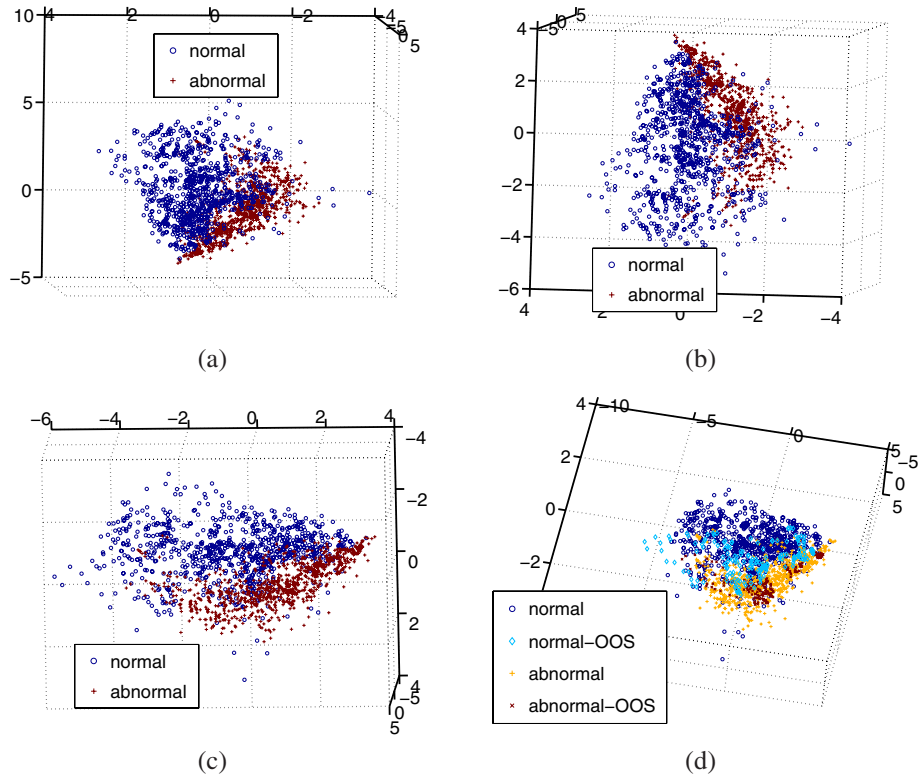


Figure 5. (a) ISOMAP embedding obtained for eigenvalue changes. (b) ISOMAP embedding obtained for eigenvector changes. (c) ISOMAP embedding for eigenvalue+eigenvector changes. (In (a)-(c), the normal tensors are indicated in dark blue and the abnormal tensors are indicated in brown.) (d) ISOMAP embedding obtained for eigenvalue changes with out-of-sample (OOS) points. (The normal tensors in the sample are indicated in dark blue and the abnormal tensors in the sample are indicated in yellow, whereas the out-of-sample normal tensors are indicated in light blue and the abnormal out-of-sample tensors are indicated in brown.)

iment for our approach to tensor interpolation: We selected an ROI on the FA map as shown in Fig. 6(a) and downsampled the tensors in this ROI by a factor of 2. We then interpolated the downsampled tensors back to their original size, i.e. we identified the dimensionality of the manifold structure present in the downsampled tensors via ISOMAP, embedded them into a lower dimensional space via LTSA, performed bilinear interpolation in the lower-dimensional embedding, and returned to the tensor domain using the LTSA mapping. Using the voxel-wise approach described in Sec. 4.2, we performed a group-wise statistical analysis on the interpolated DTIs. We also performed a similar analysis for linearly interpolated tensors and tensors interpolated using the Log-Euclidean (LE) framework [1]. Results are shown in Fig. 6(b)-(d) and may compare these results with the “ground truth” in Fig. 4(b). Here, LTSA interpolation does not perform as well as the other methods (as compared to the case of simulated tensors shown in Fig. 3). We are currently investigating the reasons for this.

5. Discussion and Conclusions

We have presented a framework for comprehensive analysis of DTI data using manifold learning techniques. Since

our approach takes advantage of the underlying manifold structure present in the tensors, it is superior to conventional analyses based upon scalar measures such as fractional anisotropy and presents an interesting alternative to approaches based upon Riemannian symmetric spaces and approaches based upon other tensor metrics. We have presented numerous simulations through which we have demonstrated the effectiveness of our approach.

A primary assumption of our work is that the manifold learning techniques used are successfully able to capture the manifold structure in the tensor data and are able to find informative vector space embeddings. However, most manifold learning techniques including ISOMAP are guaranteed to succeed only when certain conditions hold [6]. Although such conditions are difficult to directly verify in practical problems, we have found in our simulations that results obtained using manifold learning agree well with ground truth.

In future work, we plan to conduct a more thorough validation of our approach. We also plan to devise a more computationally efficient approach to tensor interpolation and plan to investigate the applicability of our approach to other problems such as PDE-based filtering and regularization of DT images [10].




Figure 4. (a) The region selected for introducing changes is marked on the normalized FA map of an abnormal subject. (b) P-values obtained for eigenvalue changes overlaid on the template FA map. (c) P-values obtained for eigenvector changes overlaid on the template FA map. (d) P-values obtained for eigenvalue+eigenvector changes overlaid on the template FA map. In all figures, only P-values < 0.001 are shown in red.




Figure 6. (a) The region selected for interpolation is marked on the template FA map. (b) P-values obtained on the LTSA-interpolated tensors overlaid on the template FA map. (c) P-values obtained on the LE-interpolated tensors overlaid on the template FA map. (d) P-values obtained on the linearly-interpolated tensors overlaid on the template FA map. (compare P-values with ground truth in Fig. 4(b)). In all figures, only P-values < 0.001 are shown in red.

References

- [1] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Fast and simple calculus on tensors in the Log-Euclidean framework. In *Lecture Notes in Computer Science: Med. Image Computing and Comp. Assist. Int.*, volume 3749, pages 115–122, 2005. 1, 2, 4, 7
- [2] P. Batchelor, M. Moakher, D. Atkinson, F. Calamante, and A. Connelly. A rigorous framework for diffusion tensor calculus. *Magnetic Resonance in Medicine*, 53:221–225, 2005. 1, 2
- [3] Y. Bengio, J.-F. Paiement, P. Vincent, O. D. N. L. Roux, and M. Ouimet. Out-of-sample extensions for LLE, isomap, MDS, eigenmaps, and spectral clustering. In *Neural Info. Proc. Sys.*, 2003. 4, 5
- [4] C. J. C. Burges. *Data Mining and Knowledge Discovery Handbook: A Complete Guide for Practitioners and Researchers*, chapter Geometric Methods for Feature Extraction and Dimensional Reduction. Kluwer Academic, 2005. 1, 5
- [5] J. Costa and A. Hero. Geodesic entropic graphs for dimension and entropy estimation in manifold learning. *IEEE Transactions on Signal Processing*, 52(8):2210–2221, 2004. 3
- [6] D. Donoho and C. Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100(10):5591–5596, 2003. 7
- [7] P. T. Fletcher and S. Joshi. Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. In *Lecture Notes in Computer Science: Computer Vision Approaches to Medical Image Analysis*, volume 3117, pages 87–98, 2004. 1, 2
- [8] M. Law and A. Jain. Incremental nonlinear dimensionality reduction by manifold learning. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 28(3):377–391, 2006. 5
- [9] D. LeBihan, J.-F. Mangin, C. Poupon, C. Clark, S. Pappata, N. Molko, and H. Chabriet. Diffusion tensor imaging: concepts and applications. *Journal of Magnetic Resonance Imaging*, 13:534–546, 2001. 1
- [10] X. Pennec, P. Fillard, and N. Ayache. A Riemannian framework for tensor computing. *Int. Journal of Computer Vision*, 66(1):41–66, 2006. 1, 2, 7
- [11] J. Tenenbaum, V. de Silva, and J. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2004. 1
- [12] P. Thompson and A. Toga. A framework for computational anatomy. *Computing and Visualization in Science*, 5:13–34, 2002. 1
- [13] Z. Wang and B. Vemuri. DTI segmentation using an information theoretic tensor dissimilarity measure. *IEEE Trans. Med. Imag.*, 24(10):1267–1277, 2005. 1
- [14] D. Xu, S. Mori, D. Shen, P. van Zijl, and C. Davatzikos. Spatial normalization of diffusion tensor fields. *Magnetic Resonance in Medicine*, 50(1):175–182, 2003. 4, 5
- [15] Z. Zhang and H. Zha. Principal manifolds and nonlinear dimensionality reduction via tangent space alignment. *SIAM Journal of Scientific Computing*, 26(1):313–338, 2004. 1