

On the Averaging of Cardiac Diffusion Tensor MRI Data: The Effect of Distance Function Selection

Archontis Giannakidis^{a,b,c,*}, Gerd Melkus^{d,e}, Guang Yang^{b,c}, Grant T. Gullberg^{a,e}

^a*Life Sciences Division, Lawrence Berkeley National Laboratory–Berkeley, CA, 94720, USA.*

^b*Cardiovascular Biomedical Research Unit, Royal Brompton Hospital–London, SW3 6NP, UK.*

^c*National Heart & Lung Institute, Imperial College London–London, SW3 6NP, UK.*

^d*Department of Radiology, University of Ottawa, Ottawa, ON, K1H 6L8, Canada.*

^e*Department of Radiology and Biomedical Imaging, University of California San Francisco–San Francisco, CA, 94143, USA.*

Abstract

Diffusion tensor magnetic resonance imaging (DT-MRI) allows a unique insight into the microstructure of highly-directional tissues. The selection of the most proper distance function for the space of diffusion tensors is crucial in enhancing the clinical application of this imaging modality. Both linear and nonlinear metrics have been proposed in the literature over the years. The debate on the most appropriate DT-MRI distance function is still ongoing. In this paper, we presented a framework to compare the Euclidean, affine-invariant Riemannian and log-Euclidean metrics using actual high-resolution DT-MRI rat heart data. We employed temporal averaging at the diffusion tensor level of three consecutive and identically-acquired DT-MRI datasets from each of five rat hearts as a means to rectify the background noise-induced loss of myocyte directional regularity. This procedure is applied here for the first time in the context of tensor distance function selection. When compared with previous studies that used a different concrete application to juxtapose the various DT-MRI distance functions, this work is unique in that it combined the following: (i) Metrics were judged by quantitative –rather than qualitative– criteria, (ii) the comparison tools were non-biased, (iii) a longitudinal comparison operation was used on a same-voxel basis. The statistical analyses of the comparison showed that the three DT-MRI distance functions tend to provide equivalent results. Hence, we came to the conclusion that the tensor manifold for cardiac DT-MRI studies is a curved space of almost zero curvature.

*Corresponding author: tel: +44 20 7351 8819, fax: ++44 20 7351 8819.

Email addresses: A.Giannakidis@rbht.nhs.uk (Archontis Giannakidis), gmelkus@toh.ca (Gerd Melkus), G.Yang@rbht.nhs.uk (Guang Yang), GTGullberg@lbl.gov (Grant T. Gullberg)

The signal to noise ratio dependence of the operations was investigated through simulations. Finally, the “swelling effect” occurrence following Euclidean averaging was found to be too unimportant to be worth consideration.

Keywords: Diffusion tensor magnetic resonance imaging, distance function, geodesics, six-dimensional manifolds, tensor averaging, myocardium microstructure, primary cardiomyocyte orientation.

1. Introduction

1.1. Subject Description

Diffusion tensor magnetic resonance imaging (DT-MRI) has emerged (Basser et al., 1994) as a powerful tool for inferring the microstructure of anisotropic tissues. To do so, it elegantly relates the self-diffusion of water molecules that undergo Brownian motion to proton spin relaxation magnetic resonance (MR) signals. The clinical significance of DT-MRI derives from the fact that it provides such information in a noninvasive, nondestructive fashion. Therefore, it paves the way for whole-organ longitudinal *in vivo* studies needed to evaluate disease-related tissue microstructural changes.

Ever since the advent of this imaging modality, there has been a substantial interest in the development of a consistent and complete framework that would allow the various manipulations of the whole diffusion tensors (and not only the analysis of some scalar or vector tensor-derived parameters). This interest has increased significantly nowadays following the initial success that DT-MRI has seen in elucidating the composition and connectivity of various anisotropic tissues including brain (Pierpaoli et al., 1996), and myocardium (Rohmer et al., 2007).

The selection of a tensor distance function (or metric) resides in the foundation of the tensor-variate framework. This selection is in essence equivalent to choosing a geodesic curve (i.e. the shortest path joining two tensors in their space). The cruciality of this selection lies in the fact that it dictates many operations such as mean estimation, anisotropy estimation, interpolation, regularization, segmentation, registration, and principal component analysis to name a few. This may be easily seen by looking at the definitions of the above operations: The averaging procedure is a tensor distance dispersion minimization problem; anisotropy estimation refers to the calculation of the distance to the isotropic tensor; the interpolation task is a weighted computation where the weights depend on tensor distances; regularization is a minimization estimation problem that aims at reducing the amount of spatial variation by taking into account the distances between neighboring tensors; segmentation algorithms involve choosing a distance function which should

be small between tensors belonging to the same structure and large between tensors belonging to different structures; registration techniques rely on distance metrics to match tensors between regions of different datasets; in principal component analysis, the computed modes of variation are represented as flows along paths of minimum tensor distance. In general, the use of an appropriate distance function for diffusion tensors provides the following twofold benefit: one may fully exploit the potential of DT-MRI by endowing the diffusion tensor space with a reliable manifold structure, while, at the same time, interpretation pitfalls are avoided.

1.2. Previous Studies & Current Situation in the Literature

With the view to satisfying the demand for a proper distance function for diffusion tensors and accordingly assisting in the development of a consistent tensor-variate framework, various metrics have been proposed in the literature over the years. To start with, the Euclidean distance function was introduced (Alexander et al., 1999; Jones et al., 2002; Basser and Pajevic, 2003; Pajevic and Basser, 2003; Pasternak et al., 2008, 2010, 2012) as a generalization of the classical scalar distance to tensors or, simply, on empirical grounds (that is to say evidence acquired by concurrently analyzing: (i) Physical properties of the measured quantities, (ii) properties of the several different candidate metrics, and (iii) the accumulative effect of the miscellaneous sources that cause variability to the measurements). As per this conventional approach, the –quantification of the distance between two tensors– operation is applied linearly to the matrix elements. Next, the affine-invariant Riemannian (AIR) distance function was put forward (Moakher, 2005; Batchelor et al., 2005; Pennec et al., 2006; Lenglet et al., 2006; Fletcher and Joshi, 2007) for DTMRI analysis based on differential geometry considerations. In particular, it was identified that the data in hand are sample covariance matrices of normally distributed random variables and, as such, they represent points in the space of real symmetric positive-definite matrices. This distance function brings about nonlinear calculations. Finally, in light of the high computational burden imposed by the AIR metric, the similitude-invariant log-Euclidean metric, that deals with the matrix logarithms of the diffusion tensors, was also proposed (Arsigny et al., 2006). Its computational efficiency lies in the fact that the high complexity calculations are converted into simple Euclidean ones, once tensors are transformed into their matrix logarithms.¹

¹Numerous other DT-MRI distance functions have been introduced such as Cholesky (Wang et al., 2004), matrix square root (Dryden et al., 2009), Procrustes (Dryden et al., 2009), square-root of the J-divergence (Wang and Vemuri, 2005), geodesic loxodromes-based (Kindlmann et al., 2007), weighted component-based (Gahm et al., 2013), and the chordal metric in the space of quaternions (Collard et al., 2014) to name a few. However, in this study we deal only with the three distance

Following the admission of the various DT-MRI distance functions, a multifaceted dispute broke out between the researchers who support the Euclidean metric and those who are in favor of the nonlinear ones. On the one hand, the non-Euclidean side (Moakher, 2005; Batchelor et al., 2005; Pennec et al., 2006; Lenglet et al., 2006; Fletcher and Joshi, 2007; Arsigny et al., 2006), driven chiefly by mathematical considerations, argue: (i) Since the space of the symmetric positive-definite matrices constitutes a curved subset (namely, the interior of a convex half-cone) in the vector space of symmetric 3×3 matrices, the AIR metric, that takes this space topology (curvature) into account, is a more natural solution. (ii) By using tools for DT-MRI, that are obtained by generalizing tools for scalar-to matrix-valued data, the natural properties of the diffusion tensors are not preserved leading to nonphysical complications. For instance, by relying on the Euclidean metric to perform non-convex operations (like regularization or second-order statistics), the positive-definiteness is not always preserved. However, diffusion tensors with negative eigenvalues lack substance. By the same token, the linear averaging/interpolation may not preserve the determinant, resulting in a mean/interpolated tensor that has a determinant (or, equivalently, volume) that is larger than the determinants (volumes) of the individual tensors. This “swelling effect” occurrence suggests that the averaging/interpolating procedure by itself introduces additional diffusion which is intangible. The two limitations described above are remedied by employing a nonlinear metric, which diminishes the “swelling effect” and places tensors with null or negative eigenvalues at an infinite distance from any symmetric positive-definite matrix. (iii) Non-Euclidean distance functions should be preferred as they offer a better contrast when transitioning from a highly anisotropic region to a more isotropic one (Fletcher and Joshi, 2007).

On the other side of the fence, the Euclidean approach investigators (Alexander et al., 1999; Jones et al., 2002; Basser and Pajevic, 2003; Pajevic and Basser, 2003; Pasternak et al., 2008, 2010, 2012) put forth the following arguments: (i) The selection of a metric should not be prompted by mathematical considerations alone. Practical and physical considerations must also be appraised. (ii) Affine-invariance is not a desirable property for a metric that deals with diffusion tensors because it leads to a substantial bias in the calculations (especially when dealing with a highly anisotropic tissue). This property would be suitable for a metric that measures the distance between arbitrarily scaled quantities. (iii) The Euclidean distance function is consistent with the expected statistical properties of the diffusion tensor distribution. (iv) If one wants to perform statistical inference through hypothesis testing for tensor-valued data, it is preferred to do so by relying on the

functions that have been most often adopted by the DT-MRI community.

Euclidean distance (Whitcher et al., 2007). Moreover, statistical tests that are based on non-Euclidean distances may easily lead to wrong conclusions (Pasternak et al., 2010). (v) With regard to the preservation of natural properties: The topic of negative eigenvalues may be remedied by excluding tensors with negative eigenvalues, or correcting tensors with negative eigenvalues through averaging, or estimating the nearest positive-definite tensor of each tensor with negative eigenvalues. As to preserving the determinant, there is no physical reason to do so. Changes in the determinant are perfectly predicted by the background noise that introduces variability in the orientations and lengths of the diffusion ellipsoid axes. Instead, preservation of the trace (which is proportional to the orientationally averaged diffusivity) might be more desired, and the Euclidean distance function (as well as the nonlinear functions) preserves the trace. (vi) The anisotropy indices based on a nonlinear distance function are inappropriate for DT-MRI analysis because they consistently and unrealistically increase the difference between highly anisotropic tensors and other more isotropic ones (Pasternak et al., 2010). (vii) Another drawback of Riemannian metrics is their sensitivity to small shape changes close to the degenerate cases. Because of this, small shape variations (due to noise) of two such tensors may result in an infinite distance between them (Peeters et al., 2009; Zhou et al., 2015).

The debate on the most appropriate DT-MRI distance function is still ongoing as suggested by more recent publications (Bouchon et al., 2015). To make things worse, the situation in the related literature is inconsistent as papers that advocated the non-Euclidean framework, they employed, at the same time, tools that are based on the Euclidean distance to evaluate their analysis results. As an illustration, Arsigny et al. (2006), and Fillard et al. (2007) used plots of a Euclidean distance-based tool to demonstrate the performance of the proposed non-Euclidean regularization. Similarly, Corouge et al. (2006) claimed that it is inappropriate to perform tensor calculations by relying on the Euclidean distance; nevertheless they employed plots of a Euclidean distance-based tool to test the feasibility of their fiber-tract-oriented approach. Finally, Verma et al. (2007) relied on Euclidean distance-based maps to simulate pathology, even though they stated that differences between diffusion tensors need to be measured by nonlinear metrics. The observations outlined in this paragraph served as the main motivation for this paper.

1.3. Our Contribution

In this work we use actual high-resolution DT-MRI rat heart data to compare the Euclidean, AIR and log-Euclidean distance functions. We look at temporal averaging as a means to rectify the background noise-induced loss of myocyte directional regularity through realigning post-diagonalization diffusion tensor primary

eigenvectors² with atypical directions. This simple and effective noise compensation method is based upon the following reasonable assumptions: (i) The tissue has constant true diffusion properties over time (Pasternak et al., 2010). (ii) The background noise processes of repetitive and identically-acquired scans are not directionally correlated (Sun et al., 2001). (iii) The cardiac wall has a well-ordered and highly coherent microstructure (Streeter and Bassett, 1966). By performing averaging at the diffusion tensor level of three consecutive and identically-acquired DT-MRI datasets from each of five rat hearts and by using the Euclidean, AIR, and log-Euclidean distance functions, the objective of this study is to shed more light on the problem of choosing the most appropriate tensor distance function. To this end, we seek to ascertain in a quantitative manner the metric that (drives the averaging method that) achieves the greatest restoration of microstructural organization (when compared with the organization of the individual datasets before averaging) or, equivalently, the most effective outlier realignment. In other words, we set out to find the distance function that yields the most accurate representation of the reality. Then, this will be identified as the preferable metric upon which a consistent tensor-variate framework may be built.

Even though the compensation of directional coherence loss through the time averaging procedure is applied here for the first time in the context of tensor distance function selection, a number of previous studies also presented an application-oriented comparison of the several different DT-MRI distance functions. For example, in Fletcher and Joshi (2007), and Dryden et al. (2009) the comparison of metrics was carried out with respect to the anisotropy estimation performance, in Pennec et al. (2006), Arsigny et al. (2006), and Frindel et al. (2009) the various distance functions were rated according to their regularization efficiency, while in Arsigny et al. (2006), Gahm et al. (2011) and Yang et al. (2012) the metric that best characterizes the distance between diffusion tensors was decided by the interpolation performance.

Our work is unique in that it combines the following: (i) The majority of the previous tensor distance function comparison studies, that were based on actual DT-MRI data, were qualitative. In order to validate their results, the authors of these papers relied on optical inspection of the volume size and color uniformity of neighboring ellipsoids³ (or other visually appealing structural renderings) combined with micro-anatomical knowledge. For instance, they labelled as the most

²The primary eigenvectors of the diffusion tensors are generally taken (Hsu et al., 1998) to coincide with the local primary cardiomyocyte orientation.

³It is common to use ellipsoids to describe diffusion tensors, where the three principal axes of the ellipsoid point at the diffusion tensor eigenvectors' directions and each axis length is proportional to the corresponding eigenvalue.

proper DT-MRI distance function the one that appeared to suppress the outliers to the greatest extent (Pennec et al., 2006; Frindel et al., 2009), or the one that seemed to yield the best contrast (Fletcher and Joshi, 2007; Dryden et al., 2009), or the one that manifested the lowest degree of blurring (Arsigny et al., 2006), or the one that presented the highest color (orientation) and size uniformity among neighboring ellipsoid volumes (Yang et al., 2012). In contrast to these comparison studies that were judged by unquantifiable criteria, our work provides quantification of the relative performance of each distance function. To this end, we relied on the inter-voxel diffusion coherence (IVDC) index (Wang et al., 2008). The validity of the specific eigenvector homogeneity criterion for assessing the perturbation of cardiac tissue directional regularity due to disease has been tested elsewhere (Giannakidis et al., 2012). (ii) There exist few studies (Frindel et al., 2009; Yang et al., 2012; Collard et al., 2014; Zhou et al., 2015) that used fractional anisotropy (FA) to compare the various diffusion tensor distance functions. However, such a criterion is unfairly prejudiced, as it represents the Euclidean distance of each tensor from the fully isotropic tensor. Unlike these studies, our comparison tool is not biased. (iii) In contrast to all previous studies, where the operation (that yielded the comparison results) was performed in the neighborhood of the examined voxel, our study involves a longitudinal comparison operation that is carried out on the individual voxel. In the final analysis, this difference in the comparison circumstances might play a role in the conclusive choice according to the idea proposed in Pasternak et al. (2008) about context-dependent DT-MRI metrics.

2. Methods

2.1. Research Animal Model & Heart Preparation

We performed *ex vivo*⁴ studies of five Wistar Kyoto (WKY) rats bought from the Charles River Laboratories International, Inc, Wilmington, MA, USA. WKY rat is a well-established animal model of normal nonpathologic cardiac behavior. Under deep isoflurane-inhalation anaesthesia, each intact heart was rapidly removed from the chest and flushed with warmed isotonic saline. Once the heart was rinsed, it was weighed and instantaneously placed in 60cc of 10% buffered formalin for fixation. The age of the rats during the heart excision was one month old. The time period between excision and imaging was approximately one week.

⁴The large (compared to the diffusion-induced displacement length) movement of a pumping heart poses a significant challenge to *in vivo* cardiac DT-MRI (Hsu et al., 2010). As a result, at present, cardiac DT-MRI is mainly conducted *ex vivo* using excised fixed hearts, especially when three-dimensional (3D) high spatial resolution images are required.

All animal procedures conformed to the guidelines set forth by the Animal Welfare and Research Committee of Lawrence Berkeley National Laboratory.

2.2. Imaging & DT-MRI Post-processing

Three back-to-back DT-MRI acquisitions of each of the five rat hearts were carried out at the UCSF Imaging Center at China Basin using a 7 T (310 mm bore size) superconducting magnet equipped with actively shielded imaging gradients (400 mT/m maximum gradient strength, 120 mm inner bore size) (Agilent Technologies, Palo Alto, CA, USA). A 20 mm inner diameter linear ^1H birdcage resonator was used for radio frequency pulse transmission and signal reception. For the imaging, the excised rat hearts were suspended in a 15 mm diameter cylinder filled with Fomblin (Sigma Aldrich Corp., St. Louis, MO, USA). Fomblin does not have a visible proton MR signal and reduces susceptibility artifacts near the boundary of the heart. Hearts were secured inside the containers using gauze to prevent the specimen from floating in the container. We did not observe any effects from movement or vibrations. The long axis of the heart was aligned with the direction of the main magnetic field (z -axis). The field of view (FOV) for the DT-MRI sequence was placed to cover the full short axis and the central part (mid-ventricular region) of the long axis of the heart. The DT-MRI imaging protocol parameters (that were common for all acquisitions) are summarized in Table 1.

Table 1: The imaging parameters used for each scan.

Sequence type	3D spin echo
Repetition time (T_R)	500ms
Echo time (T_E)	20ms
Field of view	$20.5 \times 20.5 \times 5.1\text{mm}^3$
Matrix	$128 \times 128 \times 32$
Resolution	0.160mm (isotropic)
Diffusion gradient duration (δ)	4ms
Diffusion gradient separation (Δ)	10ms
Receiver bandwidth	30.5 Hz
Number of gradient directions and b -value	12 at $b=1000\frac{\text{s}}{\text{mm}^2}$, 1 at $b=0\frac{\text{s}}{\text{mm}^2}$ (Papadakis et al., 1999)
Number of averages	1
Total scan time	7.4h

The signal to noise ratio (SNR) of each dataset was measured as the mean signal of the myocardium in the central slice of the dataset divided by the standard

deviation of the noise (measured outside the myocardium), multiplied by the factor 0.655

$$SNR_{(b=0)} = 0.655 * \frac{Mean(Myocardium)}{StandardDeviation(BackgroundNoise)} \quad (1)$$

where the factor 0.655 comes from the fact that the SNR was estimated on the magnitude data and the mean of the noise is not 0 (Rician Noise) (Henkelman, 1985).

For DT-MRI analysis, the raw diffusion weighted datasets were loaded into the FSL software package. The diffusion tensor fitting and the calculations of eigenvectors, eigenvalues and FA were performed using the FSL diffusion toolbox (FDT) (Behrens et al., 2003).

2.3. Regions of Interest - Segmentation

We analyzed the lateral wall (Fig. 1) from the short-axis slice with the largest area (in the mid left ventricular region) for each heart of this study. The region of interest (ROI) was segmented semi-automatically in the B0 image using cubic splines (de Boor, 1978). The papillary muscles of the heart were excluded from our analysis. In addition, any voxel, at which the diffusion tensor was estimated to have at least one negative eigenvalue, was excluded on thermodynamic grounds.

Left Ventricular Segmentation


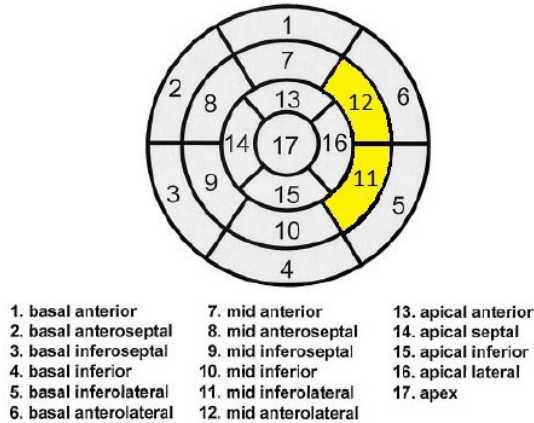


Figure 1: The region of interest (in yellow color) for the current comparison study.

2.4. From Geodesic Curves to Averaging

Consider that we have a set of N ($N \geq 2$) symmetric positive-definite 3×3 matrices D_1, \dots, D_N each of them representing a diffusion tensor. Below we

briefly outline the definitions and properties of the metrics for the three studied frameworks. We also provide the mean estimators.

2.4.1. The Euclidean Tensor-Variate Framework

By placing the set of N tensors in a Euclidean space, the geodesic (i.e. curve of minimum length) $\gamma_E(t)$ going from \mathbf{D}_1 (at $t = 0$) to \mathbf{D}_2 (at $t = 1$) is a straight line given by:

$$\gamma_E(t) = (1 - t)\mathbf{D}_1 + t\mathbf{D}_2 \quad (2)$$

The distance between tensors \mathbf{D}_1 and \mathbf{D}_2 is

$$\text{dist}_E = \|\mathbf{D}_1 - \mathbf{D}_2\|_F \quad (3)$$

where $\|\cdot\|_F$ denotes the Frobenius norm, and $\|\mathbf{D}\|_F = \sqrt{\text{tr}(\mathbf{D}^T \mathbf{D})}$ where “tr” is the matrix trace. The metric of Eq. (3) is reflection and rotation invariant⁵, and is defined over the entire space of symmetric matrices. The average of the set of N tensors that is associated with the Euclidean distance function is unique (Dryden et al., 2009) given by:

$$\mu_E(\mathbf{D}_1, \dots, \mathbf{D}_N) \equiv \arg \min_{\mathbf{D}} \sum_{i=1}^N \|\mathbf{D}_i - \mathbf{D}\|_F^2 = \frac{1}{N} \sum_{i=1}^N \mathbf{D}_i \quad (4)$$

2.4.2. The Affine-Invariant Riemannian (AIR) Tensor-Variate Framework

By placing the set of N tensors in an AIR manifold, the geodesic $\gamma_{AIR}(t)$ going from \mathbf{D}_1 (at $t = 0$) to \mathbf{D}_2 at ($t = 1$) is given (Arsigny et al., 2006) by

$$\gamma_{AIR}(t) = \mathbf{D}_1^{\frac{1}{2}} \exp \left(t \log \left(\mathbf{D}_1^{-\frac{1}{2}} \mathbf{D}_2 \mathbf{D}_1^{-\frac{1}{2}} \right) \right) \mathbf{D}_1^{\frac{1}{2}} \quad (5)$$

where $\mathbf{D}_1^{-\frac{1}{2}} \mathbf{D}_2 \mathbf{D}_1^{-\frac{1}{2}}$ belongs to the space of real symmetric positive-definite matrices (Lenglet et al., 2006), and the square root, the exponential and the logarithm of a matrix are defined in Appendix B. The distance between these two tensors is represented by

$$\text{dist}_{AIR} = \|\log(\mathbf{D}_1^{-1} \mathbf{D}_2)\|_F = \sqrt{\text{tr}(\log^2(\mathbf{D}_1^{-1} \mathbf{D}_2))} = \sqrt{\sum_{i=1}^3 \log^2 \eta_i} \quad (6)$$

⁵The invariance properties of the DT-MRI distance functions are explained in Appendix A.

where $\eta_i \{i = 1, 2, 3\}$ are the three real positive eigenvalues of matrix $\mathbf{D}_1^{-1}\mathbf{D}_2$ (Moakher, 2005). The metric of Eq. (6) is affine (i.e. reflection, rotation, scale, shear, and inversion) invariant and pertains only to symmetric positive-definite matrices. The mean estimator in the AIR sense of the N tensors is obtained by solving the following nonlinear minimization problem:

$$\mu_{AIR} \equiv \arg \min_{\mathbf{D}} \sum_{i=1}^N \|\log(\mathbf{D}_i^{-1}\mathbf{D})\|_F^2 \quad (7)$$

Even though this problem always has a unique solution (given the negative curvature of the space of symmetric positive-definite matrices (Dryden et al., 2009)), there is a closed-form solution only for the case of $N = 2$. In this case, the average estimator is given explicitly (Moakher, 2005) by:

$$\mu_{AIR} = \mathbf{D}_1(\mathbf{D}_1^{-1}\mathbf{D}_2)^{\frac{1}{2}} \quad (8)$$

For the general case of $N \geq 3$ tensors, there is no closed-form solution and one may obtain the mean via an unconstrained optimization algorithm. In our study (where $N = 3$), μ_{AIR} was estimated using the well-known quasi-Newton method (in particular the BFGS algorithm) (Nocedal and Wright, 2006). Some of the advantages of the particular quasi-Newton method are: (i) It achieves rapid superlinear convergence. (ii) Unlike other Newton methods, it is designed to prevent the analytical calculation of the Hessian matrix or its inverse at each iteration. (iii) Due to the efficiency of the quasi-Newton methods, it facilitates the realistic development and solution of large-scale optimization problems.

2.4.3. The Log-Euclidean Tensor-Variate Framework

In the log-Euclidean case, the geodesic $\gamma_{LE}(t)$ going from \mathbf{D}_1 (at $t = 0$) to \mathbf{D}_2 (at $t = 1$) is a straight line in the domain of matrix logarithms described by the following equation:

$$\gamma_{LE}(t) = \exp((1-t)\log(\mathbf{D}_1) + t\log(\mathbf{D}_2)) \quad (9)$$

In an analogous manner to the Euclidean space, the distance between \mathbf{D}_1 and \mathbf{D}_2 is given by:

$$\text{dist}_{LE} = \|\log(\mathbf{D}_1) - \log(\mathbf{D}_2)\|_F \quad (10)$$

The metric of Eq. (10) is reflection, rotation, scale, and inversion invariant. The mean estimator of the set of N tensors that is associated with the log-Euclidean distance function is given (Dryden et al., 2009) by:

$$\mu_{LE} \equiv \arg \min_{\mathbf{D}} \sum_{i=1}^N \|\log(\mathbf{D}_i) - \log(\mathbf{D})\|_F^2 = \frac{1}{N} \exp\left(\sum_{i=1}^N \log(\mathbf{D}_i)\right) \quad (11)$$

2.5. Metric Comparison Tool Based on Microstructural Orientation Coherence

After performing averaging at the diffusion tensor level by relying on the Euclidean, AIR and log-Euclidean distance functions (as detailed in Section 2.4), the goal was to juxtapose the primary cardiomyocyte orientation coherence restoration performance of the three metrics. To gauge the performance of each distance function, it was necessary to quantify the primary cardiomyocyte orientation coherence before and after each type of averaging. To quantitatively map the orientation integrity of the principal water diffusion eigenvectors within a given voxel neighborhood, we relied on the intervoxel diffusion coherence (IVDC) index (Wang et al., 2008). IVDC offers a quantitative assessment and allows a more objective directional analyses than the visually appealing red-green-blue (RGB) color maps (Pajevic and Pierpaoli, 1999) and cardiac tractography methods (Rohmer et al., 2007). In addition and unlike other measures of primary eigenvector dispersion applied in the DT-MRI literature such as the coherence index (CI) (Klingberg et al., 1999), IVDC has the supplementary favorable quality that it is insensitive to the inherent eigenvector sign ambiguity (Beg et al., 2007). Below the definition of the IVDC index is briefly provided.

IVDC is a scatter matrix-based tool (Wang et al., 2008), so let us first recall what is a scatter matrix. Consider an arbitrary voxel (i, j, k) , of the cardiac wall, and a local ROI that consists of that voxel and its 26 nearest neighbors that span in the same short-axis slice and the adjacent slices above and below. Then, the scatter matrix $\mathbf{T}(i, j, k)$ (with respect to the primary diffusion eigenvectors) at this voxel and for this ROI is a symmetric positive semidefinite second-order dyadic tensor of size 3×3 given by

$$\mathbf{T}(i, j, k) \equiv \frac{1}{27} \sum_{l=i-1}^{i+1} \sum_{m=i-1}^{i+1} \sum_{n=i-1}^{i+1} \epsilon_1(l, m, n) \epsilon_1^T(l, m, n) \quad (12)$$

where $\epsilon_1(l, m, n)$ is the major eigenvector at the (l, m, n) voxel. Finally, the IVDC index is given by

$$\text{IVDC}(i, j, k) \equiv \frac{\sqrt{(t_1 - \bar{t})^2 + (t_2 - \bar{t})^2 + (t_3 - \bar{t})^2}}{\bar{t}\sqrt{6}} \quad (13)$$

where $t_1 \equiv t_1(i, j, k)$, $t_2 \equiv t_2(i, j, k)$, $t_3 \equiv t_3(i, j, k)$ are the eigenvalues of $\mathbf{T}(i, j, k)$ and $\bar{t} \equiv \bar{t}(i, j, k)$ is the mean of these eigenvalues. IVDC is normalized between 0 and 1. A large value of this index at a specific voxel indicates that there is uniform primary cardiomyocyte orientation distribution in this voxel's neighborhood (that is the case in tissues with highly coherent cardiomyocyte organization such as the myocardium). On the other hand, smaller values of IVDC reflect the

regional loss of primary cardiomyocyte orientation integrity [that could be due to the presence of noise-induced outliers or disease (disarray)]. The foundation work for the IVDC definition as a measure of vector angular uniformity may be found in Bingham (1974).

To assess the extent of the directional regularity restoration achieved by each distance function, we estimated the relative IVDC changes before and after each of the three types of averaging.

2.6. *Signal to Noise Ratio Dependence of the Operations*

We performed simulations to investigate the SNR dependence of the operations. We synthesized noisy replicates of the original datasets through adding noise of Rician distribution in the magnitude images (Henkelman, 1985). Four simulated realizations of various levels of noise were used. We performed the comparison of the three tensor distance functions for each noisy realization.

2.7. *Preservation of Diffusion Tensor Natural Properties by Euclidean Averaging*

Linear algebra-based tensor averaging might not preserve the determinant giving rise to a mean diffusion tensor that: (i) Has a determinant larger than the determinants of the individual tensors (Corouge et al., 2006), and (ii) is substantially more isotropic than the nonlinear average tensors (Arsigny et al., 2006). This “swelling-effect” occurs because (contrary to the vector-case, where the linear averaging of a set of vectors gives the mean event) the linear averaging of a set of tensors provides both the mean event and the range of present events (Westin et al., 2002). Even though the determinant preservation requirement is in question (Pasternak et al., 2010), the extensive discussions about this topic provided us with a reason to calculate the “swelling-effect” occurrence as it was observed in the averaging procedure of our real cardiac DT-MRI data.

In addition, it has been argued that by relying on the Euclidean metric to perform operations, the positive-definiteness is not always preserved. Motivated by this statement, we also calculated the number of voxels for which the Euclidean averaging operation generated negative eigenvalues.

2.8. *Statistical Analysis*

To test the statistical significance of the differences in the primary cardiomyocyte orientation coherence recovery performance (as well as determinant and degree of isotropy) among the three studied distance functions, we employed the nonparametric Kruskal-Wallis test (Breslow, 1970). A value of $p < 0.05$ was considered to be statistically significant.

All computations described in Section 2 were performed using in-house code written in Matlab (Mathworks, Natick, MA, USA).

3. Results

The measured SNR values for the datasets of this study are provided in Table 2. To visualize the extent of the directional regularity restoration achieved by each distance function, we plotted the logarithmic relative change in IVDC before and after each of the three types of averaging (Fig. 2 - top). This plot shows that the three different DT-MRI distance functions tend to provide equivalent results.

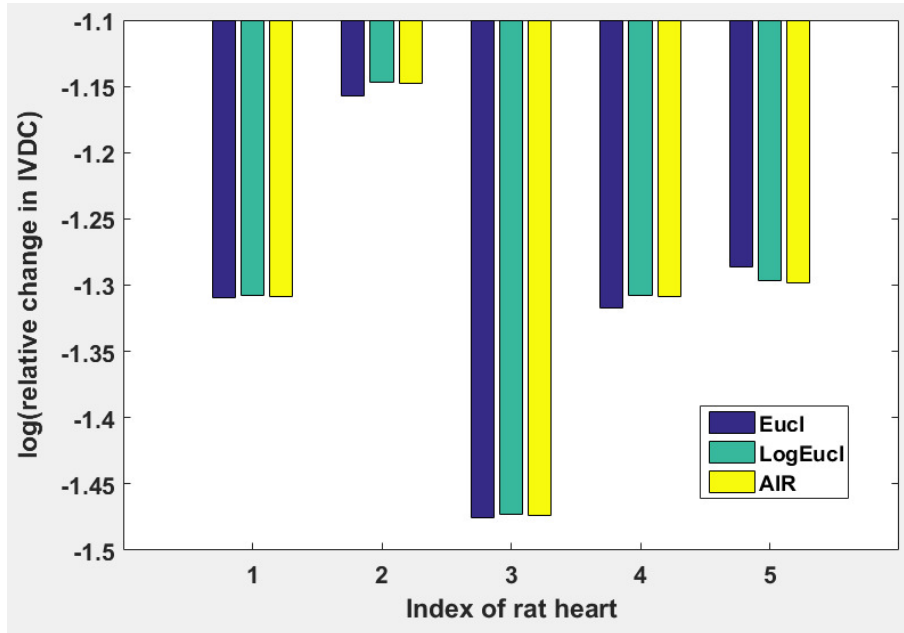
Table 2: The measured signal to noise ratio (SNR) values for the datasets of our study.

	SNR value
Rat heart 1	80
Rat heart 2	75
Rat heart 3	78
Rat heart 4	85
Rat heart 5	84

To determine whether the primary cardiomyocyte orientation coherence recovery performance is indeed the same for the three metrics, we performed the Kruskal-Wallis test for the IVDC-based comparison results. The outcomes of this nonparametric test are summarized in the table shown at the bottom of Fig. 2. This table has columns for the sum of squares (SS), degrees of freedom (df), mean squares ($MS=SS/df$), chi-square statistic, and p -value. It is this chi-square statistic that was actually used to test the null hypothesis that all three metrics perform the same. The extremely high p -value that was obtained consists a strong indication that we cannot reject the null hypothesis. Hence, there is no important statistical difference in the three studied metrics.

To assess whether the results are sensitive to the cardiac wall region, we compared the performance of the Euclidean, AIR, and log-Euclidean tensor distance functions on 20 short-axis slices of one rat heart of this study. The 20 slices spanned the whole 3D acquisition volume. The comparison results, shown in Fig. 3, indicate that the three distance functions tend to provide equivalent results, irrespective of cardiac region.

Regarding the SNR dependence investigation, and taking into account that the SNR of the acquired datasets was found to be approximately 80, the synthesized noisy replicates corresponded to SNR values of 28, 41, 54, and 67. The results (p -values) of the tensor distance function comparison for each noisy realization are given in Table 3. It is obvious from these results that the three distance functions tend to provide equivalent results, irrespective of the noise level.



Kruskal-Wallis ANOVA Table					
Source	SS	df	MS	Chi-sq	Prob>Chi-sq
Columns	11.2	2	5.6	0.56	0.7554
Error	268.3	12	22.3583		
Total	279.5	14			

Figure 2: Top: The directional regularity restoration performance of the Euclidean, log-Euclidean and AIR distance functions as assessed by the logarithmic relative change in the IVDC index before and after averaging for the five hearts. Bottom: The results of the Kruskal-Wallis test for the null hypothesis that the primary cardiomyocyte orientation coherence recovery performance, as assessed by relying on the IVDC index, is the same for the three distance functions.

Table 3: The results (p -values) of the tensor distance function comparison for each noisy realization.

SNR value	p -value
28	0.8781
41	0.8106
54	0.8781
67	0.8781

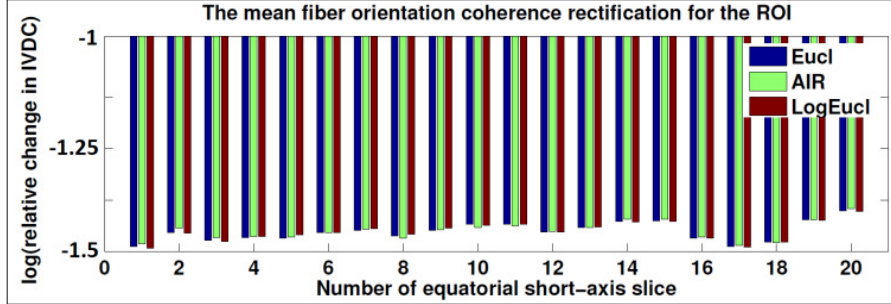


Figure 3: The directional regularity restoration performance of the Euclidean (Eucl), log-Euclidean (logEucl) and affine-invariant Riemannian (AIR) distance functions as assessed by the logarithmic relative change in the IVDC index before and after averaging for the five hearts. Results are presented over 20 short-axis slices of one rat heart of this study. The 20 slices spanned the whole 3D acquisition volume.

The results of our investigation about the preservation of diffusion tensor natural properties by Euclidean averaging are described in Table 4, and Figs. 4, 5. It may be seen there that: (i) Euclidean averaging did not generate any negative eigenvalues (Table 4). (ii) The percentage of ROI voxels that did not preserve the determinant of the Euclidean mean ranged between 7.47% and 33.33% for the five hearts (Table 4). (iii) The determinant of the Euclidean mean at a voxel was almost always (>99.9%) larger than the determinant of the other two nonlinear means at the same voxel (Table 4). (iv) The differences in the average determinant among the three resultant tensor average datasets was insignificant ($p=0.7788$) (Fig. 4). (v) The differences in the degree of isotropy among the three resultant tensor average datasets was insignificant ($p=0.5117$) (Fig. 5). (vi) The increased determinant of the Euclidean mean tensor (when compared with the determinants of the two nonlinear mean tensors) at a voxel did not always translate to a more isotropic tensor, though it happened more often than not (>72%) (Table 4).

4. Discussion

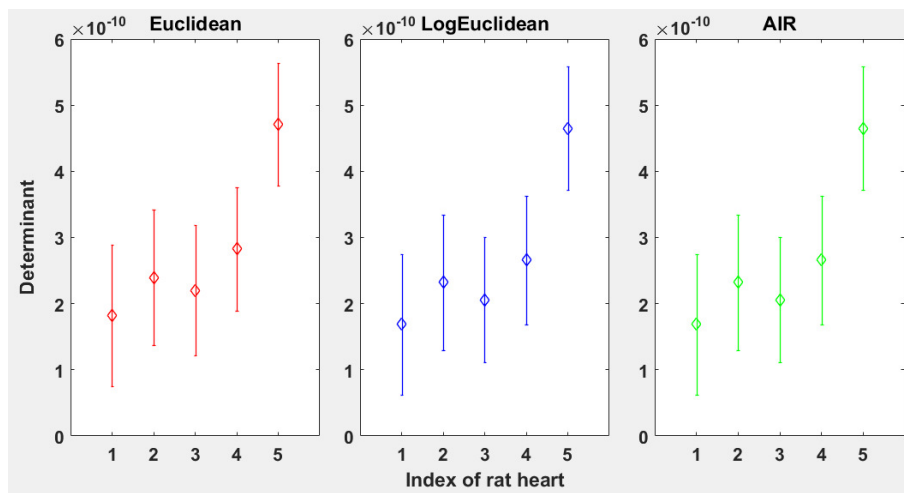
4.1. The Comparison Results

We presented a framework to compare the Euclidean, AIR and log-Euclidean distance functions using actual high-resolution DT-MRI rat heart data. By performing averaging (which is an operation that is dictated by the distance function) at the diffusion tensor level of three DT-MRI datasets (that were back-to-back and identically acquired) from each of five excised and fixed rat hearts, and by applying the three distance functions, this study sought to shed some more light on the

Table 4: The results of our investigation about the preservation of diffusion tensor natural properties by Euclidean averaging.

	Heart 1	Heart 2	Heart 3	Heart 4	Heart 5
Number of voxels where the Euclidean averaging operation generated negative eigenvalues	0	0	0	0	0
Percentage (%) of voxels at which $\det_{Eucl.mean}$ was not preserved	8.83	7.47	2.20	18.97	33.33
Percentage (%) of voxels at which $\det_{Eucl.mean} > \det_{AIR.mean}$	99.90	100	100	100	100
Percentage (%) of voxels at which $\det_{Eucl.mean} > \det_{LogEucl.mean}$	99.90	100	100	100	100
Percentage (%) of voxels at which $FA_{Eucl.mean} < FA_{AIR.mean}$	91.78	72.51	97.09	97.38	81.19
Percentage (%) of voxels at which $FA_{Eucl.mean} < FA_{LogEucl.mean}$	92.89	75.97	97.60	98.15	84.05

selection of the most appropriate tensor distance function. To this end, we sought to ascertain in a quantitative manner the metric that achieved the greatest rectification of background noise-induced misaligned principal diffusion eigenvectors. Primary orientation is a piece of central information carried by the diffusion tensor data, as well as a main measure affected by disease (Giannakidis et al., 2012). The comparison performed in this paper showed that the linear and non-linear DT-MRI functions tend to provide equivalent results (Fig. 2). Hence, we conclude that the diffusion tensor manifold for cardiac studies is a curved space of almost zero curva-

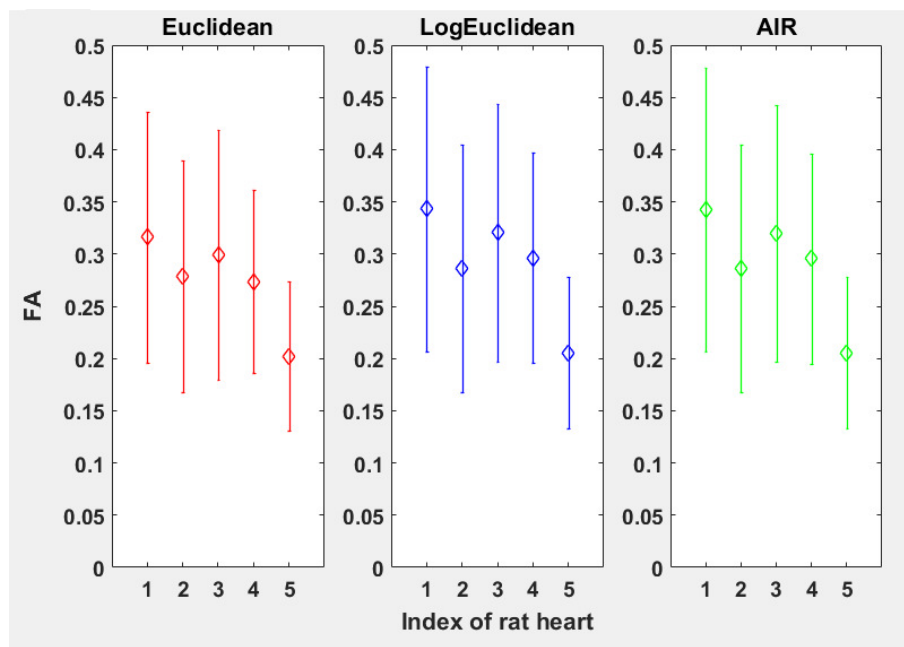


Kruskal-Wallis ANOVA Table					
Source	SS	df	MS	Chi-sq	Prob>Chi-sq
Columns	10	2	5	0.5	0.7788
Error	270	12	22.5		
Total	280	14			

Figure 4: Top: The determinant of the Euclidean (left), log-Euclidean (middle) and AIR (right) resultant tensor average datasets for the five hearts. ROI averages for each heart appear as symmetric error bars that are centered in the mean value and two standard deviations long. Bottom: The results of the Kruskal-Wallis test for the null hypothesis that the average determinant is the same for the three ways of tensor averaging.

ture (Kay, 1988). The results of our study are different from the findings of several works in the related literature that suggest the superiority of one of the three metrics. Moreover, it is worth noting that despite the solid mathematics that underlie the Riemannian framework, its superiority over the uncomplicated Euclidean and log-Euclidean frameworks could not be demonstrated here. For equivalent performances, the Euclidean and log-Euclidean frameworks offer the advantage of being much more (~ 13 times here) computationally inexpensive.

When compared with previous DT-MRI metric comparison studies that used a different concrete application to assess suitability, this work is unique in that it combines the following: (i) We provide quantification (by using IVDC) of the relative performance of each distance function. Thus, our strict formalism allows the objective comparison of metrics instead of subjective assessment of displays. (ii) Our comparison tool (IVDC) is not biased. (iii) Our study involved a longitudinal comparison operation that was carried out on a same-voxel-basis.



Kruskal-Wallis ANOVA Table					
Source	SS	df	MS	Chi-sq	Prob>Chi-sq
Columns	26.8	2	13.4	1.34	0.5117
Error	253.2	12	21.1		
Total	280	14			

Figure 5: Top: The fractional anisotropy (FA) of the Euclidean (left), log-Euclidean (middle) and AIR (right) resultant tensor average datasets for the five hearts. ROI averages for each heart appear as symmetric error bars that are centered in the mean value and two standard deviations long. Bottom: The results of the Kruskal-Wallis test for the null hypothesis that FA is the same for the three ways of tensor averaging.

There has been a lot of discussion in the literature about the “swelling effect” occurrence when performing linear averaging. In this paper, we investigated the relevance of the Euclidean metric in terms of stability. We found that maybe this topic is more an academic rather than a practical one (Table 4, Fig. 4, Fig. 5). Also it is worth noting that the increased determinants of the Euclidean mean tensor dataset (when compared with the determinants of the log-Euclidean and AIR mean tensor datasets) did not always translate to a more isotropic tensor distribution, though it happened more often than not (>72%) (Table 4). It would be reasonable

to accept that the “swelling effect” occurrence will be more substantial when averaging: (i) Tensors that are different realizations of a tensor in the same voxel and the individual datasets were obtained at experiments of great variability (for example, in the imaging parameters). (ii) Tensors that belong to areas of significantly different degree of anisotropy. However, great caution should be exerted relative to both cases described above: (i) One should be very careful when interrelating tensors obtained at various experiments with different imaging parameters. (ii) One should avoid performing operations between regions of very different anisotropy degree. For example, Arsigny et al. (2006) should have refrained from the interpolation between tensors that belong to the corpus callosum and the adjacent ventricles. It would be more prudent to apply some pre-segmentation.

This paper did not investigate the effect of experimental parameters such as the number of diffusion gradients. Instead, we used the gradient scheme proposed by Papadakis et al. (1999). Papadakis et al. (1999) performed a comparison of gradient schemes, and the scheme used in our work was that selected as optimal in the comparison. While there is still some debate (Lebel et al., 2012) as to whether it is more beneficial to use imaging time to acquire more diffusion directions or signal averages, we do not believe that the parameter selection would have a large effect on the relative performance of the tensor distance metrics. In fact, the error caused by the limited number of gradients is not pertinent to this study, as it is systematic error (Hsu et al., 2010); the implemented temporal averaging has an impact only on non-systematic errors, such as the one caused by the background noise.

In this paper, averaging was performed at the level of diffusion tensors. The mean estimation procedure may be also applied at the level of diffusion weighted images, in k -space data, or at the level of diffusion eigenvector fields. However, the subject of this paper is to compare metrics/distance functions for the 6-dimensional space of diffusion tensors; the scalar/vector averaging topic and the error propagation study (Koay et al., 2007) are beyond the scope of this work.

4.2. *The Comparison Method*

To ensure the proper quantitative comparison of the performances of the three DT-MRI distance functions, through the results obtained by a primary cardiomyocyte orientation coherence restoration study, a noise-free “gold standard” was needed. As a rule, it is difficult for one to have the objective reality of the situation when dealing with real tissue samples. Traditionally, non-destructive DT-MRI has been validated against histologically reconstructed cardiomyocyte orientations (Hsu et al., 1998), and high resolution non-destructive 3D fast low angle shot (FLASH) MRI (Bernus et al., 2015). However, the histological validation is difficult and imperfect. It requires cutting and mounting the histological sections

which can easily distort the tissue. Subsequent registration of the DT-MRI and histological data is also difficult. FLASH MRI is also a surrogate measure of cardiomyocyte orientation. In fact, it can only infer the cardiomyocyte orientation based on the direction of minimal change in image gradient. Therefore, the ground truth for this study was accommodated by the related literature. In particular, invasive morphological studies at the cellular level showed (Streeter and Bassett, 1966) that myocardium is a well-ordered and highly coherent medium. As a matter of fact, by examining locally the myocardium composition, it was found by Streeter and Bassett (1966) that the long axes of the myocytes are parallel, thus providing a systematic and relatively homogeneous primary structure for the aggregates. In addition, it was revealed by Streeter and Bassett (1966) that the rate, with which the cardiomyocyte inclination angles change as a function of transmural location, is continuous and smooth. The conformation of the heart wall cells described above is a key determinant of the cardiac electromechanics (Roberts et al., 1979).

In this paper, we analyzed the mid-lateral cardiac wall of five rats which were at an early stage of their life when the heart excision took place. By staying clear of the posterior and anterior junction areas of the left ventricle wall and the right ventricle free wall (where cardiomyocyte disarray has been reported (Kuribayashi, 1987)) and also by avoiding noncollinearities in the cardiomyocyte distribution due to aging (Kuribayashi, 1987), this further enhances the truthfulness of our expectation that the noise-free “gold-standard” of the studied tissue orientation distribution is a collinear one.

DT-MRI is highly susceptible to noise with the background signal of the diffusion weighted volumes being the main interference component (Chen and Hsu, 2005; Bao et al., 2009; Hsu et al., 2010).⁶ The manifestation of this background signal has its roots in the signal attenuation nature of the diffusion encoding and the prolonged data acquisition requirement (i.e. a minimum of seven MRI readings is needed) (Chen and Hsu, 2005; Hsu et al., 2010). If one attempts to improve the spatial resolution (which is equivalent to obtaining larger datasets), this gives a further boost to the background noise effects due to the unfortunate trade-off (that DT-MRI is subject to) between spatial resolution and signal to noise ratio (SNR). The diffusion tensor at each voxel (which is derived by the noisy diffusion weighted volumes) is typically diagonalized to yield eigenvectors (that indicate the principal directions of diffusion) and eigenvalues (that are measures of the amount of diffusion along the directions pointed by the respective eigenvectors). Therefore, it

⁶Other stochasticity sources in DT-MRI entail: (i) The eddy currents following the high demands placed on the gradient system, and (ii) other sources of systematic error such as deficiencies in the diffusion encoding scheme and poor diffusion modeling.

is quite justifiable that the presence of the background noise discussed above has been accounted (Bao et al., 2009; Basser and Pajevic, 2000; Martin et al., 1999) for the miscalculation of the eigenvector directions and the magnitude (or, even worse, the sorting) of the eigenvalues. With respect to the primary eigenvectors (that are of particular interest in this study), the background noise in the diffusion weighted volumes causes some of them to have atypical directions. In the case of cardiac DT-MRI, the background signal results in the microstructure of a few myocardium regions erroneously appearing as discontinuous (Yang et al., 2011; Frindel et al., 2009). In fact, there have been proposed few DT-MRI regularization methods (Frindel et al., 2009; Zhou et al., 2015) that used priors to favor smoothness of local diffusion in highly homogeneous tissues.

In our study, averaging of repetitive DT-MRI acquisitions was employed with a view to correcting misaligned principal diffusion eigenvectors. In fact, it has been a common practice (Haacke et al., 1999) (due to simplicity and effectiveness) in diffusion imaging to increase the accuracy of the estimation of various parameters by performing multiple measurements and then averaging. As an illustration, averaging of consecutive (and identically acquired) scans was carried out in Sun et al. (2001), and Skare et al. (2000) with a view to reducing the background noise influence in the estimation of anisotropy indices.

4.3. Quantitation of Averaging Impact on Tissue Orientation Integrity

We would like to emphasize that the rectification of misaligned (due to noise) primary eigenvectors in DT-MRI is a topic that has been researched at a fair extent in the past. A number of sophisticated (and much more time efficient than averaging) techniques based, for example, on partial differential equation anisotropic filtering (Chen and Hsu, 2005), sparse representation (Bao et al., 2009), and the variational method (Coulon et al., 2004) have been proposed. In this work, we do not propose averaging as a means to carry out denoising. We only bring this topic in the context of DT-MRI distance function selection. However, it has caught the authors' attention the fact that while, on one hand, the time averaging for DT-MRI background noise removal is a task that is carried out heavily by researchers and scanners (Haacke et al., 1999), on the other hand, the quantitative impact of this task in specific characteristics (such as tissue orientation integrity) remains largely uncharacterized. Motivated by this observation, a supplementary contribution of this paper is the quantification (see Table 5) of the improvement in cardiomyocyte direction coherence (amongst adjacent voxels) that is achieved by averaging three (and/or two) back-to-back DT-MRI scans of a rat heart over the individual scans. From this table it is clear that the IVDC value increases with the number of averages. A main reason why the averaging impact on primary cardiomyocyte orientation coherence is not larger is that averaging cannot correct systematic errors

(Haacke et al., 1999). One such source of systematic error is the fact that, in some voxels, the single tensor model provides a poor modeling for the diffusion process. To get a better insight into the averaging impact on primary cardiomyocyte orientation coherence discussed above, we also provide images (see Fig. 6) of the primary eigenvector field. These pinpoint in great detail the realignment of post-diagonalization diffusion tensor primary eigenvectors with atypical directions following tensor averaging.

Table 5: Quantification of the averaging impact on the primary cardiomyocyte orientation coherence as assessed by the IVDC index for one heart of this study. Results are shown when the Log-Euclidean distance function was used for tensor averaging. Similar results were obtained for the other two studied distance functions.

	IVDC
1st acquisition	0.9585
2nd acquisition	0.9196
3rd acquisition	0.9288
Mean of 1st and 2nd acquisitions	0.9602
Mean of 1st and 3rd acquisitions	0.9617
Mean of 2nd and 3rd acquisitions	0.9463
Mean of all three acquisitions	0.9649

Finally, IVDC could prove a useful quantitative tool for other DT-MRI applications, where structure-sensitive functions (derived from local orientation similarity) are needed. Such applications are, for example, segmentation (Lenglet et al., 2006), and interpolation comparison (Yang et al., 2011).

5. Conclusion

The comparison performed in this paper showed that the Euclidean, AIR and log-Euclidean tensor distance functions tend to provide equivalent results in terms of primary cardiomyocyte orientation coherence restoration performance. The tensor manifold for cardiac DT-MRI studies is a curved space of almost zero curvature. The results of our study are different from the findings of several works in the related literature that suggest the superiority of one of the three metrics. Moreover, it is worth noting that despite the solid mathematics that underlie the Riemannian framework, its superiority over the uncomplicated Euclidean and log-Euclidean frameworks could not be demonstrated here. For equivalent performances, the Euclidean and log-Euclidean frameworks offer the advantage of being much more (~ 13 times here) computationally inexpensive.

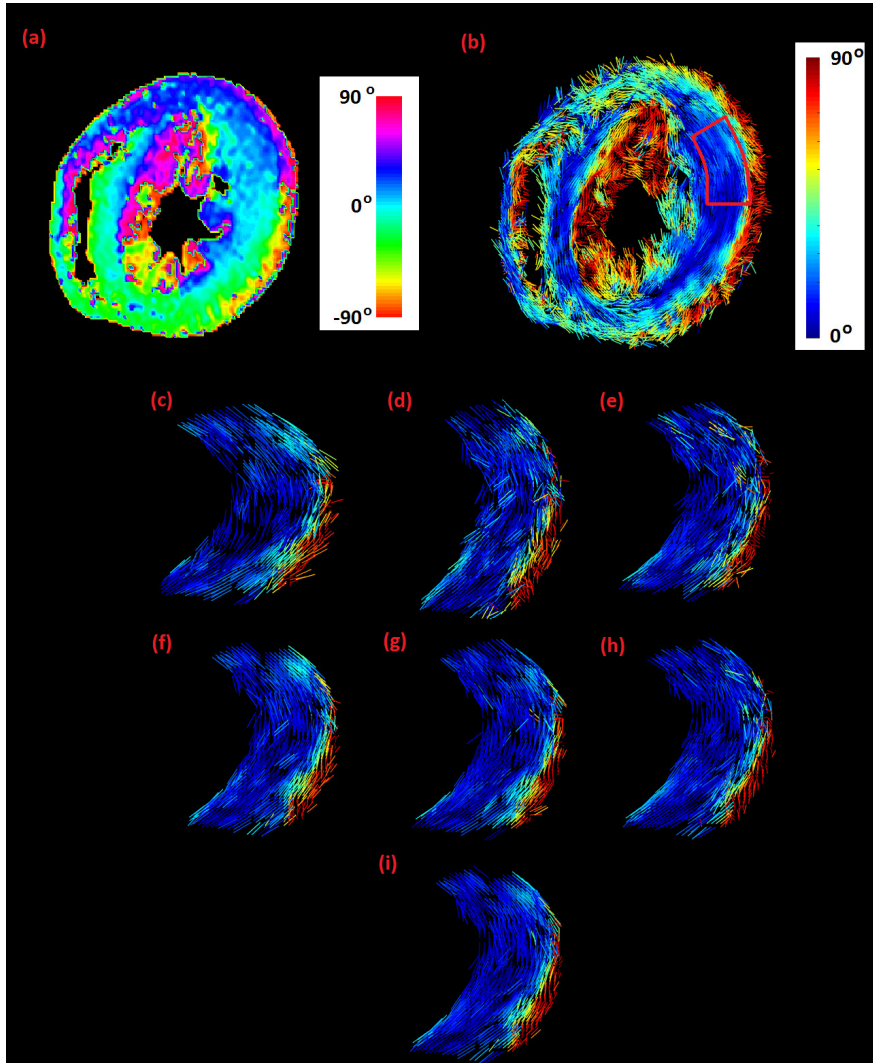


Figure 6: (a) Helix angle map, where helix angle is defined (Streeter et al., 1969) as the angle between the cardiac short-axis plane and the projection of the primary eigenvector onto the epicardial tangent plane. (b)-(i) Images of the primary eigenvector field, color encoded to the absolute local helix angle. (b) The whole equatorial short-axis slice. The red sector indicates the zoom-in area (lateral wall). Remaining images are close-ups that refer to the 1st acquisition (c), 2nd acquisition (d), 3rd acquisition (e), mean of the 1st and 2nd acquisitions (f), mean of the 1st and 3rd acquisitions (g), mean of the 2nd and 3rd acquisitions (h), mean of all three acquisitions (i). The tensor averaging impact is the realignment of some primary eigenvectors that had atypical directions due to noise. Results are shown when the Log-Euclidean distance function was used for tensor averaging. Similar results were obtained for the other two studied distance functions.

Finally, the “swelling effect” occurrence following Euclidean averaging was found to be too unimportant to be worth consideration. Hence, this topic is maybe more an academic rather than a practical one.

Appendix A. Invariance Properties of the Various Distance Functions

Assume d is a function used to measure distance between two diffusion tensors, P and Q . Then:

- d is invariant under reflection (or, else, re-ordering) if $d(P, Q) = d(Q, P)$.
- d is invariant under rotation if $d(P, Q) = d(S^T P S, S^T Q S)$, where S any orthogonal matrix.
- d is scale invariant if $d(P, Q) = d(\beta P, \beta Q)$, for all $\beta > 0$.
- d is invariant under inversion if $d(P, Q) = d(P^{-1}, Q^{-1})$.
- d is affine invariant, if $d(P, Q) = d(G^T P G, G^T Q G)$, where G a general full rank matrix.

Appendix B. Definition of Matrix Logarithm, Exponential of a Matrix, and Square Root of a Matrix

Assume $W = U D U^T$ is the spectral decomposition of a symmetric positive-definite matrix W , where U is the orthonormal matrix the columns of which are the eigenvectors of W , and $D = \text{diag}(d_i)$ is the diagonal matrix of the strictly positive eigenvalues d_i of W . Then:

- The exponential of W is $\exp(W) \equiv \sum_{k=0}^{\infty} \frac{W^k}{k!} = U \text{diag}(\exp(d_i)) U^T$.
- The matrix logarithm is represented by: $\log(W) \equiv - \sum_{k=1}^{\infty} \frac{(I-A)^k}{k} = U \text{diag}(\log(d_i)) U^T$ where I denotes the identity matrix.
- The formula for the square root of W is: $W^{\frac{1}{2}} \equiv \exp\left(\frac{1}{2} \log(W)\right)$.

Acknowledgments

The authors would like to thank Prof. Thomas F. Budinger, Dr. Jean-Marc Peyrat, Dr. Ziad El Bitar, Dr Pedro F. A. D. C. Ferreira, Dr Andrew D. Scott, and Steve Collins for their help and Prof. David Firmin for the generous hospitality. We are grateful to Stephanie Murphy for preparation of the excised heart

samples. This work was supported by: (i) The National Institutes of Health under Grant R01 EB007219. (ii) The Director, Office of Science, Office of Biological and Environmental Research, Medical Sciences Division of the U.S. Department of Energy under Contract DE-AC02-05CH11231. (iii) The NIHR Cardiovascular Biomedical Research Unit of Royal Brompton & Harefield NHS Foundation Trust and Imperial College London.

References

- Alexander, D., Gee, J., Bajcsy, R., 1999. Similarity measures for matching diffusion tensor images. Proceedings of the 10th British Machine Vision Conference (BMVC '99) , 93–102, Nottingham, UK.
- Arsigny, V., Fillard, P., Pennec, X., Ayache, N., 2006. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magn. Reson. Med.* 56, 411–421.
- Bao, L.J., Zhu, Y.M., Liu, W.Y., Croisille, P., Pu, Z.B., Robini, M., Magnin, I.E., 2009. Denoising human cardiac diffusion tensor magnetic resonance images using sparse representation combined with segmentation. *Phys. Med. Biol.* 54, 1435–1456.
- Basser, P.J., Mattiello, J., LeBihan, D., 1994. MR diffusion tensor spectroscopy and imaging. *Biophys. J.* 66, 259–267.
- Basser, P.J., Pajevic, S., 2000. Statistical artifacts in diffusion tensor MRI (DT-MRI) caused by background noise. *Magn. Reson. Med.* 44, 41–50.
- Basser, P.J., Pajevic, S., 2003. A normal distribution for tensor-valued random variables: Applications to diffusion tensor MRI. *IEEE Trans. Med. Imaging* 22, 785–794.
- Batchelor, P.G., Moakher, M., Atkinson, D., Calamante, F., Connelly, A., 2005. A rigorous framework for diffusion tensor calculus. *Magn. Reson. Med.* 53, 221–225.
- Beg, M.F., Dickie, R., Golds, G., Younes, L., 2007. Consistent realignment of 3D diffusion tensor MRI eigenvectors. In the Proceedings of SPIE Medical Imaging 2007 - Signal Processing 6512, 1–9. San Diego, CA, USA.
- Behrens, T.E.J., Woolrich, M.W., Jenkinson, M., Johansen-Berg, H., Nunes, R.G., Clare, S., Matthews, P.M., Brady, J.M., Smith, S.M., 2003. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn. Reson. Med.* 50, 1077–1088.

- Bernus, O., Radjenovic, A., Trew, M.L., LeGrice, I.J., Sands, G.B., Magee, D.R., Smaill, B.H., Gilbert, S.H., 2015. Comparison of diffusion tensor imaging by cardiovascular magnetic resonance and gadolinium enhanced 3D image intensity approaches to investigation of structural anisotropy in explanted rat hearts. *J. Cardiovasc. Magn. Reson.* 17, 31.
- Bingham, C., 1974. An antipodally symmetric distribution on the sphere. *Ann. Stat.* 2, 1201–1225.
- de Boor, C., 1978. A practical guide to splines. Applied Mathematical Sciences. Vol. 27.. Springer-Verlag, New York, Berlin.
- Bouchon, A., Noblet, V., Heitz, F., Lamy, J., Blanc, F., Armspach, J.P., 2015. Which manifold should be used for group comparison in diffusion tensor imaging? In the Proceedings of the 18th Medical Image Computing and Computer-Assisted Intervention International Conference (MICCAI 2015), Lecture Notes in Computer Science 9349 , 158–165, Munich.
- Breslow, N., 1970. A generalized Kruskal-Wallis test for comparing K samples subject to unequal patterns of censorship. *Biometrika* 57, 579–594.
- Chen, B., Hsu, E.W., 2005. Noise removal in magnetic resonance diffusion tensor imaging. *Magn. Reson. Med.* 54, 393–401.
- Collard, A., Bonnabel, S., Phillips, C., Sepulchre, R., 2014. Anisotropy preserving DTI processing. *Int. J. Comput. Vis.* 107, 58–74.
- Corouge, I., Fletcher, P.T., Joshi, S., Gouttard, S., Gerig, G., 2006. Fiber tract-oriented statistics for quantitative diffusion tensor MRI analysis. *Medical Image Analysis* 10, 786–798.
- Coulon, O., Alexander, D.C., Arridge, S., 2004. Diffusion tensor magnetic resonance image regularization. *Medical Image Analysis* 8, 47–67.
- Dryden, I.L., Koloydenko, A., Zhou, D., 2009. Non-Euclidean statistics for covariance matrices, with applications to diffusion tensor imaging. *Ann. Appl. Stat.* 3, 1102–1123.
- Fillard, P., Pennec, X., Arsigny, V., Ayache, N., 2007. Clinical DT-MRI estimation, smoothing, and fiber tracking with log-Euclidean metrics. *IEEE Trans. Med. Imaging* 26, 1472–1482.
- Fletcher, P.T., Joshi, S., 2007. Riemannian geometry for the statistical analysis of diffusion tensor data. *Sign. Process.* 87, 250–262.

- Frindel, C., Robini, M., Croisille, P., Zhu, Y.M., 2009. Comparison of regularization methods for human cardiac diffusion tensor MRI. *Medical Image Analysis* 13, 405–418.
- Gahm, J., Wisniewski, N., Klug, W.S., Garfinkel, A., Ennis, D.B., 2011. Statistical comparison of DT-MRI interpolation methods using cardiac DT-MRI data. In the Proceedings of the 19th Scientific Meeting of the Intl. Soc. Magn. Reson. Med. (ISMRM 2011) , 3894, Montreal, Quebec, Canada.
- Gahm, J.K., Kung, G.L., Ennis, D.B., 2013. Weighted component-based tensor distance applied to graph-based segmentation of cardiac DT-MRI. In the Proceedings of the 2013 IEEE 10th International Symposium of Biomedical Imaging (ISBI 2013) , 504–507, San Francisco, CA, USA.
- Giannakidis, A., Rohmer, D., Veress, A.I., Gullberg, G.T., 2012. Cardiac Mapping. John Wiley and Sons Ltd, Chichester, UK. chapter 53: Diffusion tensor MRI-derived myocardial fiber disarray in hypertensive left ventricular hypertrophy: Visualization, quantification and the effect on mechanical function. 4th edition. pp. 574–588.
- Haacke, E.M., Brown, R.W., Thompson, M.R., Venkatesan, R., 1999. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. Wiley-Blackwell, New York. 1st edition.
- Henkelman, R.M., 1985. Measurement of signal intensities in the presence of noise in MR images. *Med. Phys.* 12, 232–233.
- Hsu, E.W., Healy, L.J., Einstein, D.R., Kuprat, A.P., 2010. *Computational Cardiovascular Mechanics: Modeling and Applications in Heart Failure*. Springer, New York, NY, USA. chapter 2: Imaging-based assessment and modeling of the structures of the myocardium. pp. 23–39.
- Hsu, E.W., Muzikant, A.L., Matulevicius, S.A., Penland, R.C., Henriquez, C.S., 1998. Magnetic resonance myocardial fiber-orientation mapping with direct histological correlation. *Am. J. Physiol. Heart Circ. Physiol.* 274, H1627–H1634.
- Jones, D.K., Griffin, L.D., Alexander, D.C., Catani, M., Horsfield, M.A., Howard, R., Williams, S.C.R., 2002. Spatial normalization and averaging of diffusion tensor MRI data sets. *NeuroImage* 17, 592–617.
- Kay, D.C., 1988. *Schaum’s outline of theory and problems of tensor calculus*. McGraw-Hill, New York, NY, USA.

- Kindlmann, G., Estepar, R.S.J., Niethammer, M., Westin, C.F., 2007. Geodesic-
loxodromes for diffusion tensor interpolation and difference measurement. In
the Proceedings of the 10th Medical Image Computing and Computer-Assisted
Intervention International Conference (MICCAI 2007), Lecture Notes in Com-
puter Science 4791 , 1–9, Brisbane, Australia.
- Klingberg, T., Vaidya, C.J., Gabrieli, J.D.E., Moseley, M.E., Hedehus, M., 1999.
Myelination and organization of the frontal white matter in children: A diffusion
tensor MRI study. *NeuroReport* 10, 2817–2821.
- Koay, C.G., Chang, L.C., Pierpaoli, C., Basser, P.J., 2007. Error propagation frame-
work for diffusion tensor imaging via diffusion tensor representations. *IEEE
Trans. Med. Imaging* 26, 1017–1034.
- Kuribayashi, T., 1987. Spontaneously occurring hypertrophic cardiomyopathy in
the rat. I. Pathologic features. *Jpn. Circ. J.* 51, 573–588.
- Lebel, C., Benner, T., Beaulieu, C., 2012. Six is enough? Comparison of diffusion
parameters measured using six or more diffusion-encoding gradient directions
with deterministic tractography. *Magn. Reson. Med.* 68, 474–483.
- Lenglet, C., Rousson, M., Deriche, R., Faugeras, O., 2006. Statistics on the man-
ifold of multivariate normal distributions: Theory and application to diffusion
tensor MRI processing. *J. Math. Imaging Vis.* 25, 423–444.
- Martin, K.M., Papadakis, N.G., Huang, C.L.H., Hall, L.D., Carpenter, T.A., 1999.
The reduction of the sorting bias in the eigenvalues of the diffusion tensor. *Magn.
Reson. Imaging* 17, 893–901.
- Moakher, M., 2005. A differential geometric approach to the geometric mean of
symmetric positive-definite matrices. *SIAM J. Matrix Anal. Appl.* 26, 735–747.
- Nocedal, J., Wright, S., 2006. *Numerical Optimization*. Springer, New York. 2nd
edition.
- Pajevic, S., Basser, P.J., 2003. Parametric and non-parametric statistical analysis
of DT-MRI data. *J. Magn. Reson.* 161, 1–14.
- Pajevic, S., Pierpaoli, C., 1999. Color schemes to represent the orientation of
anisotropic tissues from diffusion tensor data: Application to white matter fiber
tract mapping in the human brain. *Magn. Reson. Med.* 42, 526–540.
- Papadakis, N.G., Xing, D., Huang, C.L.H., Hall, L.D., Carpenter, T.A., 1999. A
comparative study of acquisition schemes for diffusion tensor imaging using
MRI. *J. Magn. Reson.* 137, 67–82.

- Pasternak, O., Sochen, N., Basser, P.J., 2010. The effect of metric selection on the analysis of diffusion tensor MRI data. *NeuroImage* 49, 2190–2204.
- Pasternak, O., Sochen, N., Basser, P.J., 2012. *New Developments in the Visualization and Processing of Tensor Fields*, Part of the series *Mathematics and Visualization*. Springer-Verlag, Berlin Heidelberg. chapter : Metric selection and diffusion tensor swelling. pp. 323–336.
- Pasternak, O., Verma, R., Sochen, N., Basser, P.J., 2008. On what manifold do diffusion tensors live? Workshop of the 11th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI '08): Manifold Learning in Medical Imaging New York, NY, USA.
- Peeters, T.H.J.M., Rodrigues, P.R., Vilanova, A., ter Haar Romeny, B.M., 2009. *Visualization and Processing of Tensor Fields*, Part of the series *Mathematics and Visualization*. Springer-Verlag, Berlin Heidelberg. chapter : Analysis of Distance/Similarity Measures for Diffusion Tensor Imaging. pp. 113–136.
- Pennec, X., Fillard, P., Ayache, N., 2006. A Riemannian framework for tensor computing. *Int. J. Comput. Vis.* 66, 41–66.
- Pierpaoli, C., Jezzard, P., Basser, P.J., Barnett, A., Di Chiro, G., 1996. Diffusion tensor MR imaging of the human brain. *Radiol.* 201, 637–648.
- Roberts, D.E., Hersh, L.T., Scher, A.M., 1979. Influence of cardiac fiber orientation on wavefront voltage, conduction velocity, and tissue resistivity in the dog. *Circ. Res.* 44, 701–712.
- Rohmer, D., Sitek, A., Gullberg, G.T., 2007. Reconstruction and visualization of fiber and laminar structure in the normal human heart from ex vivo diffusion tensor magnetic resonance imaging (DTMRI) data. *Invest. Radiol.* 42, 777–789.
- Skare, S., Li, T.Q., Nordell, B., Ingvar, M., 2000. Noise considerations in the determination of diffusion tensor anisotropy. *Magn. Reson. Imaging* 18, 659–669.
- Streeter, Jr., D.D., Bassett, D.L., 1966. An engineering analysis of myocardial fiber orientation in pig's left ventricle in systole. *Anat. Rec.* 155, 503–511.
- Streeter, Jr., D.D., Spotnitz, H.M., Patel, D.P., Ross, Jr., J., Sonnenblick, E.H., 1969. Fiber orientation in the canine left ventricle during diastole and systole. *Circ. Res.* 24, 339–347.

- Sun, S.W., Song, S.K., Hong, C.Y., Chu, W.C., Chang, C., 2001. Improving relative anisotropy measurement using directional correlation of diffusion tensors. *Magn. Reson. Med.* 46, 1088–1092.
- Verma, R., Khurd, P., Davatzikos, C., 2007. On analyzing diffusion tensor images by identifying manifold structure using isomaps. *IEEE Trans. Med. Imaging* 26, 772–778.
- Wang, J., Lin, Y., Wai, Y., Liu, H., Lin, C., Huang, Y., 2008. Visualization of the coherence of the principal diffusion orientation: An eigenvector-based approach. *Magn. Reson. Med.* 59, 764–770.
- Wang, Z., Vemuri, B.C., 2005. DTI segmentation using an information theoretic tensor dissimilarity measure. *IEEE Trans. Med. Imaging* 24, 1267–1277.
- Wang, Z., Vemuri, B.C., Chen, Y., Mareci, T.H., 2004. A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from complex DWI. *IEEE Trans. Med. Imaging* 23, 930–939.
- Westin, C.F., Maier, S.E., Mamata, H., Nabavi, A., Jolesz, F.A., Kikinis, R., 2002. Processing and visualization for diffusion tensor MRI. *Medical Image Analysis* 6, 93–108.
- Whitcher, B., Wisco, J.J., Hadjikhani, N., Tuch, D.S., 2007. Statistical group comparison of diffusion tensors via multivariate hypothesis testing. *Magn. Reson. Med.* 57, 1065–1074.
- Yang, F., Zhu, Y.M., E., M.I., Luo, J.H., Croisille, P., Kingsley, P.B., 2012. Feature-based interpolation of diffusion tensor fields and application to human cardiac DT-MRI. *Medical Image Analysis* 16, 459–481.
- Yang, F., Zhu, Y.M., Rapacchi, S., Luo, J.H., Robini, M., Croisille, P., 2011. Interpolation of vector fields from human cardiac DT-MRI. *Phys. Med. Biol.* 56, 1415–1430.
- Zhou, D., Dryden, I.L., Koloydenko, A.A., Audenaert, K.M.R., Bai, L., 2015. Regularisation, interpolation and visualisation of diffusion tensor images using non-Euclidean statistics. *J. Appl. Stat.* In press, doi:10.1080/02664763.2015.1080671.