

Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Extracting skeletal muscle fiber fields from noisy diffusion tensor data

David I.W. Levin^{a,*}, Benjamin Gilles^a, Burkhard Mädler^b, Dinesh K. Pai^a^a Dept. of Computer Science, The University of British Columbia, ICICS/CS Building, 201-2366 Main Mall, Vancouver, British Columbia, Canada V6T 1Z4^b Dept. of Physics and Astronomy, The University of British Columbia, 6224 Agricultural Road, Vancouver, British Columbia, Canada V6T 1Z1

ARTICLE INFO

Article history:

Received 9 March 2010

Received in revised form 10 January 2011

Accepted 29 January 2011

Available online 22 February 2011

Keywords:

Musculoskeletal
Diffusion tensor imaging
Fiber estimation
Image processing
Noise removal

ABSTRACT

Diffusion Tensor Imaging (DTI) allows the non-invasive study of muscle fiber architecture but musculoskeletal DTI suffers from low signal-to-noise ratio. Noise in the computed tensor fields can lead to poorly reconstructed muscle fiber fields. This paper describes an algorithm for producing denoised muscle fiber fields from noisy diffusion tensor data as well as its preliminary validation. The algorithm computes a denoised vector field by finding the components of its Helmholtz–Hodge decomposition that optimally match the diffusion tensor field. A key feature of the algorithm is that it performs denoising of the vector field simultaneously with its extraction from the noisy tensor field. This allows the vector field reconstruction to be constrained by the architectural properties of skeletal muscles. When compared to primary eigenvector fields extracted from noisy synthetic data, the denoised vector fields show greater similarity to the ground truth for signal-to-noise ratios ranging from 20 to 5. Similarity greater than 0.9 (in terms of fiber direction) is observed for all signal-to-noise ratios, for smoothing parameter values greater than or equal to 10 (larger values yield more smoothing). Fiber architectures were computed from human forearm diffusion tensor data using extracted primary eigenvectors and the denoised data. Qualitative comparison of the fiber fields showed that the denoised fields were anatomically more plausible than the noisy fields. From the results of experiments using both synthetic and real MR datasets we conclude that the denoising algorithm produces anatomically plausible fiber architectures from diffusion tensor images with a wide range of signal-to-noise ratios.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

1.1. Background and motivation

Fiber architecture plays a large role in understanding the function of muscles (Blemker et al., 2007). In fact, skeletal muscles made of similar fiber types but with different architectures can have a 10–20 times difference in contraction speed or force output (Lieber et al., 1992). Observing and modeling muscle architecture is important for computational studies of muscle performance (Blemker and Delp, 2005) and in medical procedures such as tendon transfer operations (Brand et al., 1981). Recent work has shown that Diffusion Tensor Imaging (DTI) can be used to accurately estimate bulk muscle architecture parameters such as pennation angle (Lansdown et al., 2007). However, musculoskeletal DTI is difficult due to the magnetic resonance (MR) properties of muscle which lead to lower signal-to-noise ratio (SNR) in the diffusion weighted images. Once DTI data has been acquired, fiber tracking can be applied in order to construct muscle fiber fields.

* Corresponding author.

E-mail addresses: dilevin@cs.ubc.ca (D.I.W. Levin), bgilles@cs.ubc.ca (B. Gilles), bmaedler@physics.ubc.ca (B. Mädler), pai@cs.ubc.ca (D.K. Pai).

These fiber fields can then be used for muscle architecture analysis (Lansdown et al., 2007).

Most fiber tracking algorithms construct fibers by advecting a particle or surface through the field of primary eigenvectors of the measured Diffusion Tensor field (Ding et al., 2001; Basser et al., 2000; Mori and van Zijl, 2002; Damon et al., 2002). Even methods such as Tensorlines utilize the primary eigenvector weighted by the linear anisotropy of the underlying tensor (Weinstein et al., 1999). Underlying these methods is the assumption that the extracted primary eigenvectors are correct; the presence of noise in the imaging data does not make this necessarily so. The question we wish to address, within the context of musculoskeletal DTI is: How does one choose an appropriate fiber direction at each voxel given a potentially noisy tensor field?

1.2. Comparison to other work and contribution

The problem of choosing a denoised fiber direction vector field directly from a diffusion tensor field is not often addressed in literature. There has been a large focus on creating denoised tensor fields from noisy tensor or diffusion weighted data (Basu et al., 2006; Fillard et al., 2007; McGraw et al., 2009; Tschumperlé and Deriche, 2003; Wang et al., 2004; Westin et al., 2002; Neji et al., 2007; Tristán-Vega and Aja-Fernández, 2010) and on directly using

the primary eigenvectors of a tensor field as fiber directions during tracking (Ding et al., 2001; Basser et al., 2000; Chedfd'hotel et al., 2002; Mori and van Zijl, 2002; Damon et al., 2002; Tschumperlé and Deriche, 2003; Westin et al., 2002). In this paper we present a method that generates a globally optimal vector field from a noisy tensor field using assumptions about the biological structure of skeletal muscle (that it is divergence free in most areas). This is in contrast to the standard methodology of performing simple eigenvector estimation after denoising tensor data. The presented algorithm performs primary eigenvector extraction simultaneously with denoising and offers four advantages beyond resistance to noise in the acquired data. First, it allows constraints to be placed on the properties of the extracted vector field. The organizational properties of muscle fiber fields are well understood and are more easily formulated as constraints on a vector field than as those on a tensor field. Second, the method resolves sign ambiguities inherent in per-voxel eigenvector extraction. Third, the method automatically resolves degenerate cases in which several principal directions have the same principal value. Finally, the method could be used in conjunction with current state-of-the-art tensor denoising methods to extract corrected primary eigenvector fields in high noise situations since denoised tensor fields can be used as input into the algorithm. The algorithm's design is motivated by the Helmholtz–Hodge decomposition of a vector field. Computing this decomposition has been previously explored in the computer graphics literature by Tong et al. (2003). In this paper we present a new method for computing a denoised muscle fiber field from a densely sampled diffusion tensor imaging volume. This allows us to reconstruct noise-reduced fiber direction fields.

2. Methods

2.1. Theoretical background

Muscle fibers, at some level, are parallel and, except at specific regions (the insertion and origin of the muscle) do not converge or diverge. The presented algorithm is motivated by the Helmholtz–Hodge decomposition of the muscle fiber field which is defined as

$$\begin{aligned} \mathbf{u}(\mathbf{x}) &= \mathbf{h}(\mathbf{x}) + \mathbf{p}(\mathbf{x}) + \mathbf{g}(\mathbf{x}) \\ \nabla \cdot \mathbf{h} &= 0, \quad \nabla \times \mathbf{h} = \mathbf{0} \\ \nabla \cdot \mathbf{p} &= 0 \\ \nabla \times \mathbf{g} &= \mathbf{0} \end{aligned} \quad (1)$$

where \mathbf{u} is the fiber vector field, ∇ is the gradient operator, $\nabla \cdot$ is the divergence operator and $\nabla \times$ is the curl operator. The Helmholtz–Hodge decomposition contains two divergence free components \mathbf{h} and \mathbf{p} (\mathbf{h} is known as the harmonic field) as well as one more curl free component \mathbf{g} . Because \mathbf{g} is the only component with non-zero curl it will contain the insertions and origins of the muscle fibers while the other two components describe the gross divergence free structure of the fiber field.

We can further simplify the decomposition by interpreting a diffusion tensor as a description of the probable direction of the diffusion gradient. Since diffusion is driven by an underlying concentration function we can view the problem of fiber field reconstruction as finding the concentration function, u , that best fits the described diffusion gradient. The reconstructed vector field ∇u will be necessarily curl free. Because muscle fibers are locally parallel it is reasonable to assume that the rotational component of the field is likely to be noise, thus assuming that the reconstructed vector field is the gradient of a scalar function u should lead to better denoising performance. It also causes \mathbf{p} in Eq. (1) to be zero because it is the only component of the decomposition that has non-zero curl. These simplifying assumptions lead us to estimate a subset of the full Helmholtz–Hodge decomposition (\mathbf{h}

and \mathbf{g}) and below we describe how to estimate \mathbf{g} and \mathbf{h} from noisy tensor data.

Let $D(\mathbf{x})$ be the diffusion tensor at point \mathbf{x} , let $\mathbf{u}(\mathbf{x})$ be the concentration gradient at point \mathbf{x} and $u(\mathbf{x})$ be the concentration at point \mathbf{x} . By definition we know that

$$\mathbf{u}(\mathbf{x}) = \nabla u(\mathbf{x}) \quad (2)$$

We define the following cost function for the match between $\mathbf{u}(\mathbf{x})$ and $D(\mathbf{x})$:

$$c(D, \mathbf{u}) = \sum_{i=1}^n \frac{\mathbf{u}(\mathbf{x}_i)^T D(\mathbf{x}_i) \mathbf{u}(\mathbf{x}_i)}{\mathbf{u}(\mathbf{x}_i)^T \mathbf{u}(\mathbf{x}_i)} \quad (3)$$

where n is the number of voxels in the imaging volume.

This cost function is a sum of Rayleigh quotients, and each term in the sum is maximized when $\mathbf{u}(\mathbf{x})$ is in the same direction as the primary eigenvector of $D(\mathbf{x})$. In order to construct a noise free vector field that is consistent with muscle physiology we must constrain this maximization.

Instead of finding the set of vectors, \mathbf{u} , that maximize the cost in Eq. (3) we search for an optimal concentration u (related by Eq. (2)). The divergence free constraint then becomes a constraint on the Laplacian of u and the optimization performed is

$$\begin{aligned} h &= \arg \max_u \sum_{i=1}^n \frac{\nabla u(\mathbf{x}_i)^T D(\mathbf{x}_i) \nabla u(\mathbf{x}_i)}{\nabla u(\mathbf{x}_i)^T \nabla u(\mathbf{x}_i)} \\ \text{s.t.} \quad &\nabla \cdot \nabla u = 0 \end{aligned} \quad (4)$$

where $\mathbf{h} = \nabla h$ is the harmonic field in Eq. (1).

Next we must solve for \mathbf{g} which we have chosen to represent as a linear combination of radial basis functions (RBFs) (Eq. (5)). Our assumption is that \mathbf{g} contains the sources and sinks of the muscle fiber field. These sources and sinks alter fiber directions locally, with the effect dissipating with distance. RBFs are an appropriate set of basis functions because they mimic this behavior. We solve for the coefficients, a_i , by substituting this \mathbf{u} into the cost function (3) and performing an unconstrained optimization in which h is held constant.

$$\mathbf{g}(\mathbf{x}) = \sum_{i=1}^n a_i \nabla g_i(\mathbf{x}) \quad (5)$$

where g_i is a radial basis function (RBF) and a_i are coefficients. The complete vector field can then be represented as

$$\mathbf{u}(\mathbf{x}) = \nabla \left(h(\mathbf{x}) + \sum_{i=1}^n a_i g_i(\mathbf{x}) \right). \quad (6)$$

We solve for the coefficients, a_i , by substituting this \mathbf{u} into the cost function (3) and performing an unconstrained optimization in which h is held constant.

Again, we are assuming that the bulk of the muscle is divergence free. Muscle fibers converge or diverge to or from tendons at the aponeurosis. The computed \mathbf{g} corrects \mathbf{h} at these points.

We have found that Gaussian RBFs (Eq. (7)) offer the best denoising in our experiments. Each RBF is positioned at a point \mathbf{y} in space. In the current implementation we center an RBF at each grid point in the imaging volume. The parameter σ can be used to adjust the support of the function. In practice a larger σ causes the denoised vector field to be smoother while a smaller σ preserves more of the field's details.

$$g_i(\mathbf{x}) = \exp \left(-\frac{1}{\sigma^2} (\mathbf{x} - \mathbf{y}_i)^T (\mathbf{x} - \mathbf{y}_i) \right) \quad (7)$$

2.2. Numerical implementation

We reconstructed DTI image volumes from diffusion weighted images (from both synthetic and human subject data) using

non-linear least squares fitting. The reconstruction implementation enforces the constraint that each diffusion tensor in the DTI image must be symmetric positive definite (SPD) (Bergmann et al., 2005). This is done by expressing each diffusion tensor as a Cholesky factorization which allows the SPD constraint to be incorporated into the cost function minimized by the least squares procedure.

Each DTI image volume contains tensors at the discrete points indexed by the triple (i,j,k) . Finite differences are a straightforward method with which to estimate the gradient and the Laplacian of u . However, one-sided finite difference schemes, on a regular grid, do not incorporate information from the entire surrounding neighborhood of the discrete tensor field D_{ijk} . To remedy this problem we discretize u and D using a Marker and Cell (MAC) grid (Harlow and Welch, 1965). In this way we can accurately compute the gradient of u at the exact location of D using centered differences. In two dimensions the Laplacian operator at a node indexed by $(i - 1/2, j)$ can be defined as

$$L_{i-1/2,j} = \frac{1}{\Delta y^2}(u_{i-1/2,j-1} - 2u_{i-1/2,j} + u_{i-1/2,j+1}) + \frac{1}{\Delta x^2}(u_{i-3/2,j} - 2u_{i-1/2,j} + u_{i+1/2,j}) \quad (8)$$

where Δx is the pixel spacing in x and Δy is the pixel spacing in the y direction (Fig. 1). Similar formulas can be defined for nodes at $u_{i,j+1/2}$. Linear interpolation is used to calculate values of u that lie off of the MAC grid.

We also define the gradient of u , at node (i,j) as

$$G_{i,j} = \left[\frac{1}{\Delta x}(u_{i+1/2,j} - u_{i-1/2,j}), \frac{1}{\Delta y}(u_{i,j+1/2} - u_{i,j-1/2}) \right]^T \quad (9)$$

so that it is centered over each diffusion tensor D_{ij} in the image. Eqs. (8) and (9) are linear combinations which can be expressed as

$$l_t = \sum_{o=1}^s L_{to} u_o \quad (10)$$

and

$$\mathbf{u}_r = \sum_{o=1}^s \mathbf{G}_{ro} u_o \quad (11)$$

where t indexes the cell walls and r indexes the cell centers of the grid. \mathbf{L} and \mathbf{G} are the linear operators for computing the Laplacian

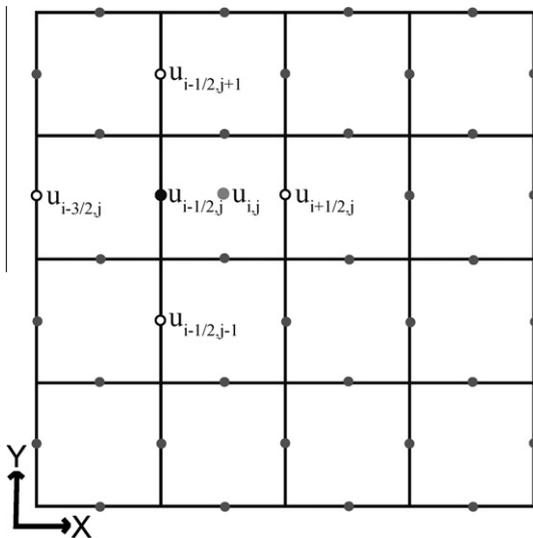


Fig. 1. A diagram of the nodes used to compute the Laplacian at node $u_{i-1/2,j}$ which is depicted by the solid black circle. Neighboring nodes are shown as hollow circles.

and the gradient and u_o is the value of u at each cell wall point o . Please note that the matrix \mathbf{G} has $2n$ (or $3n$) rows because the gradient at each point has two components (or three in three dimensions) (Eq. (9)) and so the indexing must increment by two (or three). Here we suppress this detail. As a further notational convenience we define \mathbf{G}_r which is a matrix consisting of the rows of \mathbf{G} required to compute \mathbf{u}_r in the following manner:

$$\mathbf{u}_r = \mathbf{G}_r \mathbf{u} \quad (12)$$

where \mathbf{u} is now the vector of all the discrete values of u .

Using the above operators we can finally define the discrete form of Eq. (4), used in the first step of the algorithm, as

$$h = \arg \max_u \sum_{i=1}^n \frac{u^T \mathbf{G}_i^T D_i \mathbf{G}_i u}{u^T \mathbf{G}_i^T \mathbf{G}_i u} \quad (13)$$

s.t. $\mathbf{L} \mathbf{u} = 0$.

When solving for g we begin by constructing the Jacobian matrix \mathbf{R} which is defined as

$$\mathbf{R} = [\nabla g_1 \nabla g_2 \cdots \nabla g_n] \quad (14)$$

where each column of the matrix, ∇g_i , is the gradient of an RBF (Eq. (5)) centered at one of the $1 \dots n$ MAC grid cell centers and evaluated at all of the $1 \dots n$ cell centers.

In a similar manner as before we can define \mathbf{R}_r such that

$$\mathbf{u}_r = \mathbf{h}_r + \mathbf{R}_r \mathbf{a} \quad (15)$$

where \mathbf{a} is the vector of coefficients from Eq. (5) and \mathbf{h}_r is $\mathbf{G}_r \mathbf{h}$. By substituting Eq. (15) into Eq. (3) we can define the discrete cost as a function of the coefficients, \mathbf{a} . We can then define the optimization used in the second stage of the algorithm as

$$\arg \max_a \sum_{i=1}^n \frac{(\mathbf{G}_i \mathbf{h} + \mathbf{R}_i \mathbf{a})^T D_i (\mathbf{G}_i \mathbf{h} + \mathbf{R}_i \mathbf{a})}{(\mathbf{G}_i \mathbf{h} + \mathbf{R}_i \mathbf{a})^T (\mathbf{G}_i \mathbf{h} + \mathbf{R}_i \mathbf{a})} \quad (16)$$

Our algorithm proceeds in two phases, the initial h-phase solves the optimization described by Eq. (13). This can be done using standard techniques. We use constrained steepest descent to minimize a negated Eq. (13). We maintain the constraints by projecting onto the nullspace of \mathbf{L} at each step. This projection is defined as

$$\mathbf{P} = I - \mathbf{L}^T (\mathbf{L} \mathbf{L}^T)^{-1} \mathbf{L} \quad (17)$$

where I is the identity matrix. The second phase, g-phase, solves the optimization described by Eq. (16). We use the limited memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) method (Byrd et al., 1999) to minimize a negated Eq. (16). The solution methods used in both the h-phase and g-phase of the algorithm require the gradients of Eqs. (13) and (16) which are given by

$$\nabla_u c = \sum_{i=1}^n -2 \mathbf{G}_i^T \frac{(\mathbf{u}_i^T D_i \mathbf{u}_i) \mathbf{u}_i - (\mathbf{u}_i^T \mathbf{u}_i) D_i \mathbf{u}_i}{(\mathbf{u}_i^T \mathbf{u}_i)^2} \quad (18)$$

and

$$\nabla_a c = \sum_{i=1}^n -2 \mathbf{R}_i^T \frac{(\mathbf{u}_i^T D_i \mathbf{u}_i) \mathbf{u}_i - (\mathbf{u}_i^T \mathbf{u}_i) D_i \mathbf{u}_i}{(\mathbf{u}_i^T \mathbf{u}_i)^2} \quad (19)$$

respectively.

It should be noted that if a vector field \mathbf{h} is a solution to Eq. (13) then $a\mathbf{h}$ is also a solution for any scalar a . However in terms of optimality (Eq. (3)) \mathbf{h} and $a\mathbf{h}$ are equal. We normalize \mathbf{h} such that the maximum vector length in the field is 1. As noted above, the new vector field is still a solution to Eq. (13) and its cost is unchanged but our starting harmonic for the g-phase of the algorithm becomes unambiguous. For real DTI datasets the optimization scheme was modified. The initial, constrained optimization, Eq. (13), is difficult to solve for large datasets since the number of constraints is on the order of the number of voxels in the imaging

Algorithm 1 Pseudocode for the Full Resolution, Dense Sampling Algorithm (FRDS).

1: $h0 \leftarrow (\mathbf{G}^T \mathbf{G})^{-1} \mathbf{G}^T \mathbf{e}_1(D)$
2: $\mathbf{P} \leftarrow \mathbf{I} - \mathbf{L}^T (\mathbf{L} \mathbf{L}^T)^{-1} \mathbf{L}$
3: $h0 \leftarrow \mathbf{P} h0$
4: $tol \leftarrow 1e^{-8}$
5: $rel \leftarrow 1e^8$
6: $c \leftarrow \sum_{i=1}^n \frac{h0^T \mathbf{G}_i^T D_i \mathbf{G}_i h0}{h0^T \mathbf{G}_i^T \mathbf{G}_i h0}$
h-phase 7: **while** $rel > tol$ **do**
8: $\mathbf{u} \leftarrow \mathbf{G} h0$
9: $\nabla c \leftarrow \sum_{i=1}^n -2 \mathbf{G}_i^T \frac{(\mathbf{u}_i^T D_i \mathbf{u}_i) \mathbf{u}_i - (\mathbf{u}_i^T \mathbf{u}_i) D_i \mathbf{u}_i}{(\mathbf{u}_i^T \mathbf{u}_i)^2}$
10: $\nabla c \leftarrow \mathbf{P} \nabla c$
11: $h0 \leftarrow \text{BrentsMethod}(h0, \nabla c)$ (Press et al., 1992)
12: $c_{prev} \leftarrow c$
13: $c \leftarrow \sum_{i=1}^n \frac{h0^T \mathbf{G}_i^T D_i \mathbf{G}_i h0}{h0^T \mathbf{G}_i^T \mathbf{G}_i h0}$
14: $rel \leftarrow \frac{c - c_{prev}}{c}$
15: **end while**
16: $a0 \leftarrow 0$
g-phase 17: $a \leftarrow \text{Minimize} - \sum_{i=1}^n \frac{(\mathbf{G}_i h0 + \mathbf{R}_i a^k)^T D_i (\mathbf{G}_i h0 + \mathbf{R}_i a^k)}{(\mathbf{G}_i h0 + \mathbf{R}_i a^k)^T (\mathbf{G}_i h0 + \mathbf{R}_i a^k)}$ using L-BFGS using $a0$ as an initial guess. $\{a^k$ are the coefficients at the k^{th} iteration of L-BFGS.
18: $\mathbf{u} \leftarrow \mathbf{G} h0 + \mathbf{R} a$
19: **return** *Normalize* (\mathbf{u})

Fig. 2. Pseudocode for the Full Resolution, Dense Sampling Algorithm (FRDS).

Algorithm 2 Pseudocode for the Coarse Resolution, Stochastic Sampling Algorithm (CRSS).

1: $T \leftarrow$ Trilinear prolongation operator
2: $h0 \leftarrow \min_h \|\mathbf{G} T \cdot h - \mathbf{e}_1(D)\|^2$ using Non-linear Least Squares (MINPACK)
 $\{T \cdot h$ interpolates h from the coarse grid resolution to the resolution of the imaging data.
3: $\mathbf{P} \leftarrow \mathbf{I} - \mathbf{L}^T (\mathbf{L} \mathbf{L}^T)^{-1} \mathbf{L}$
4: $h0 \leftarrow \mathbf{P} h0$
5: $tol \leftarrow 1e^{-8}$
6: $rel \leftarrow 1e^8$
7: $c \leftarrow \sum_{i=1}^n \frac{h0^T T^T \mathbf{G}_i^T D_i \mathbf{G}_i T h0}{h0^T T^T \mathbf{G}_i^T \mathbf{G}_i T h0}$
h-phase 8: **while** $rel > tol$ **do**
9: $\mathbf{u} \leftarrow \mathbf{G} T \cdot h0$
10: $\nabla c \leftarrow \sum_{i=1}^n -2 T^T \mathbf{G}_i^T \frac{(\mathbf{u}_i^T D_i \mathbf{u}_i) \mathbf{u}_i - (\mathbf{u}_i^T \mathbf{u}_i) D_i \mathbf{u}_i}{(\mathbf{u}_i^T \mathbf{u}_i)^2}$
11: $\nabla c \leftarrow \mathbf{P} \nabla c$
12: $c_{prev} \leftarrow c$
13: $c \leftarrow \sum_{i=1}^n \frac{h0^T T^T \mathbf{G}_i^T D_i \mathbf{G}_i T h0}{h0^T T^T \mathbf{G}_i^T \mathbf{G}_i T h0}$
14: $rel \leftarrow \frac{c - c_{prev}}{c}$
15: **end while**
16: $\mathbf{u} \leftarrow \mathbf{G} T \cdot h0$
17: $m \leftarrow$ Number of RBFs to fit at each iteration
18: $S \leftarrow$ Set of all unsampled voxel locations
g-phase 19: **while** S is not empty **do**
20: $S_m \leftarrow m$ randomly selected voxel locations from S
21: $S \leftarrow S - S_m$
22: $\mathbf{R}^m \leftarrow [\nabla g_{S_1} \dots \nabla g_{S_m}]$
23: $a_m \leftarrow \text{Minimize} - \sum_{i=1}^n \frac{(\mathbf{u} + \mathbf{R}_i^m a)^T D_i (\mathbf{u} + \mathbf{R}_i^m a)}{(\mathbf{u} + \mathbf{R}_i^m a)^T (\mathbf{u} + \mathbf{R}_i^m a)}$ using L-BFGS using 0 as an initial guess.
24: $\mathbf{u} \leftarrow \mathbf{u} + \mathbf{R}^m \cdot a_m$
25: **end while**
26: **return** *Normalize* (\mathbf{u})

Fig. 3. Pseudocode for the Coarse Resolution, Stochastic Sampling Algorithm (CRSS).

volume. In these cases we solve for the divergence free field on a coarse grid and then use a trilinear prolongation operator to expand this solution to full resolution such that the second optimiza-

tion can be performed. The Laplacian (Eq. (8)) has a large magnitude in areas where h , the scalar potential of the harmonic, has a high spatial frequency. Since we are constraining h to have

a Laplacian equal to zero we are searching for a function with low spatial frequencies. Such a function can be represented on a lower resolution spatial grid. In practice we limit the dimensions of the coarse grid to a maximum of $16 \times 16 \times 16$ voxels. Numerical experiments also show that all constraints remain satisfied after applying the prolongation operator. Solving for g is done using a stochastic method in which we randomly choose m voxel locations in the DTI data. We then solve for a using Eq. (16) and compute an updated u using Eq. (15). This procedure is repeated until all voxel locations in the DTI dataset have been sampled. Because the initially described method solves for h using a full resolution grid and uses a dense sampling of the imaging volume to compute g we term it Full Resolution, Dense Sampling (FRDS) whereas the second method uses a Coarse Resolution grid and Stochastic Sampling (CRSS). Pseudocode for each algorithm is shown in Algorithm 1 (Fig. 2) and Algorithm 2 (Fig. 3). The initial step in both algorithms is to compute the primary eigenvectors of each tensor D_i . This is done using LAPACK (Anderson et al., 1999) and in the pseudo code we denote this operation as $e(D)$. We use MINPACK (Moré et al., 1980) for all non-linear least squares computations.

2.3. Synthetic data generation

In order to test the effect of noise on the denoising algorithm we generated four synthetic datasets. The first was a constant vector field with the vector $(0,0,1)$ at each point. The second was a two dimensional (2D) bipennate muscle dataset. This dataset was an $8 \times 8 \times 8$ volume of tensors computed from a vector field defined by

$$\mathbf{v}_s(\mathbf{x}) = \begin{cases} (1, -1, 0)^T & \text{if } \mathbf{x}_x < 4 \\ (-1, -1, 0)^T & \text{if } \mathbf{x}_x > 4. \end{cases} \quad (20)$$

The third was a semi-circular parallel line dataset described by

$$\mathbf{v} = (-r \sin(\theta), r \cos(\theta), 0)^T \quad (21)$$

where $r = \|\mathbf{x} - \mathbf{x}_c\|$, \mathbf{x}_c is the center of the curves and θ is $\arctan(\frac{y}{x})$. The fourth was a synthetic 3D bipennate dataset described by

$$\mathbf{v}_s(\mathbf{x}) = \begin{cases} (1, -1, 1)^T & \text{if } \mathbf{x}_x < 4 \\ (-1, -1, 1)^T & \text{if } \mathbf{x}_x > 4. \end{cases} \quad (22)$$

Simulated Diffusion Weighted Images were computed from these vector fields using the algorithm of Bergmann et al. (2005). Rician noise was added to the artificial DWIs at varying SNR. Noisy tensor fields were then computed from these DWIs. Both algorithm versions were tested on SNR values ranging from 20 to 5. The primary eigenvector field of each noisy tensor field was extracted and compared to the original vector field using the sum of the absolute value of the dot product as a similarity measure:

$$s = \frac{1}{n} \sum_{i=1}^n \|\mathbf{v1}(\mathbf{x}_i)^T \mathbf{v2}(\mathbf{x}_i)\| \quad (23)$$

where n is the number of voxels in the dataset.

Denoised vector fields were then computed and these were compared to the original vector field using the same metric.

2.4. MRI data acquisition and segmentation

Imaging for this study was performed on a 3T Philips Achieva MRI Scanner with dual nova gradients (80 mT/m maximum gradient strength, 200 T/m/s maximum slewrate) and scanner software release 2.1.3.

During image acquisition the subject lay prone in the scanner with the left arm raised straight overhead to be placed as optimally as possible in the magnet's center. The subject's forearm was secured in an 8-element phased array knee coil with 15 cm inner diameter.

The imaging protocol consisted of a fast gradient echo T1W localizer for positioning and planning followed by low and high resolution T2W fast spin-echo (FSE) scans for reconstruction of bone and muscle/fat surface boundaries. The session was concluded with a high resolution Diffusion Tensor Imaging (DTI) scan for muscle fiber orientation and segmentation.

Quick low-resolution anatomical FSE-scans were acquired with an in-plane resolution of $1.5 \times 1.5 \text{ mm}^2$ and a slice thickness of 4 mm covering the entire lower arm.

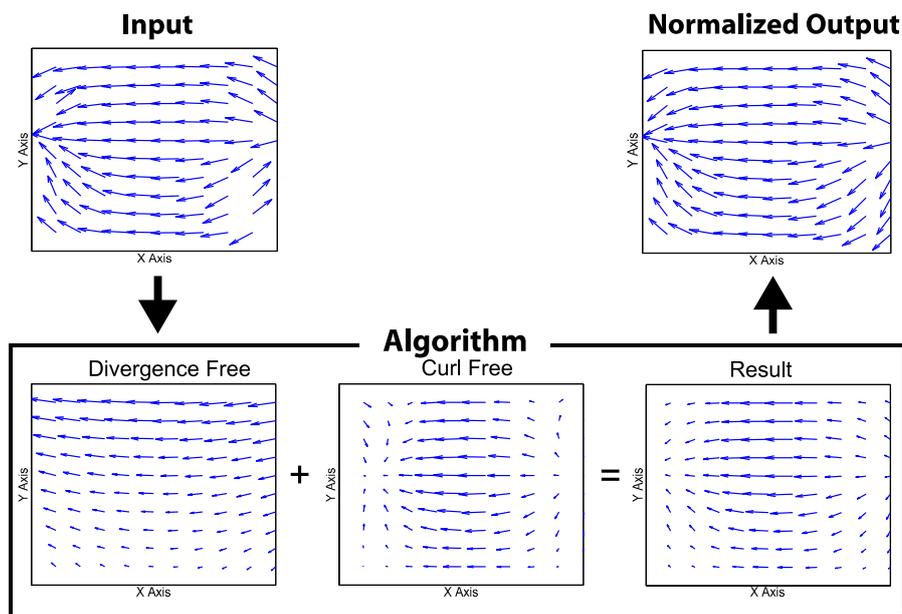


Fig. 4. The output of the various algorithm stages. This figure shows the principal eigenvector field of the input tensors (top left) as well as the final output of the algorithm (top right) and the output of the divergence free reconstruction, estimated curl free component and their sum.

The high resolution FSE was designed to match the Diffusion Tensor Imaging (DTI) scan in location, orientation and anatomy coverage with the following parameters: FSE-factor 12 with asymmetric profile order to give an effective echo time of $TE=50$ ms; field of view (FOV): $120 \times 120 \times 150$ mm³ with an in-plane resolution of 0.65×0.65 mm² and a slice thickness of 2 mm.

The lower resolution T2W-scan was used for segmenting bones and muscles that passed out of the field of view of the high resolution scan. Important parameters such as origin/insertion locations and bone coordinate systems could thus be obtained. Note that these two scans were run sequentially with the DTI scan and that the subject was immobilized. Therefore all volumes were closely aligned.

Diffusion Tensor Imaging (DTI) was performed with a single shot diffusion sensitized spin-echo Echo Planar Imaging (EPI) se-

quence involving 16 different gradient encoding directions at a maximum diffusion b -value of 500 s/mm². We used a reduced FOV of $120 \times 120 \times 150$ mm³, SENSE-factor of 2.0 and enhanced gradient performance to shorten the echo train length of the EPI-readout as much as possible for better compensation of susceptibility induced artifacts. Fat suppression was performed with a spectral spatial inversion prepared fat suppression technique. Further imaging parameters were as follow: $TE = 48$ ms, $TR = 6000$ ms, acquisition matrix 80×80 leading to an effective acquisition voxel size of $1.5 \times 1.5 \times 2.0$ mm³ and a scan time of 5 min. Additional image processing was done using the FSL software library. An affine registration of anatomical and diffusion scans was performed using FSL's linear registration application FLIRT (Jenkinson and Smith, 2001) and eddy correction of the diffusion images was performed using the utility FDT (FMRIB, 2006).

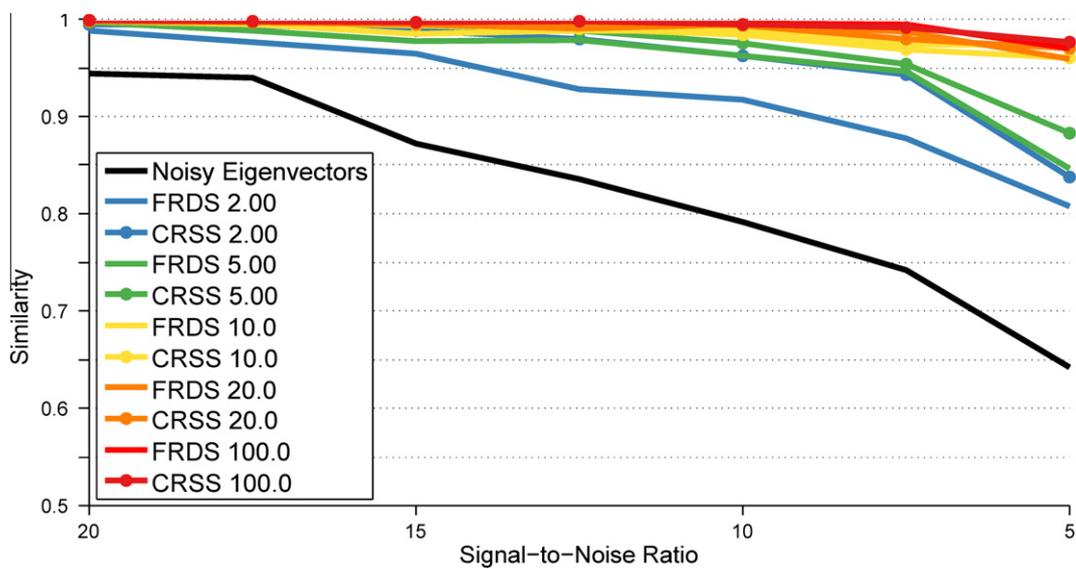


Fig. 5. Synthetic Z-Axis vector field: The similarity of both the denoised vector field and the noisy primary eigenvector field to the original, zero noise data. The result for each algorithm is denoted by the algorithm name followed by the value of the parameter σ .

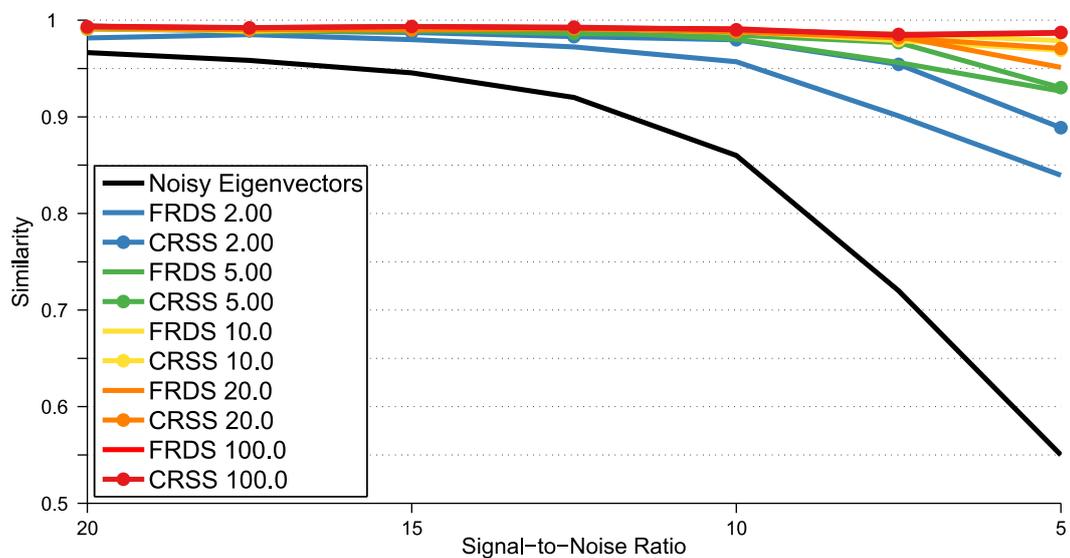


Fig. 6. Synthetic curved vector field: The similarity of both the denoised vector field and the noisy primary eigenvector field to the original, zero noise data. The result for each algorithm is denoted by the algorithm name followed by the value of the parameter σ .

Muscle surfaces were created by segmenting the anatomical MRI data using the fast shape matching algorithm of Gilles and Pai (2008). The three muscle surfaces generated were those of the brachioradialis (BR), extensor carpi radialis longus (ECRL) and the extensor carpi radialis brevis (ECRB). Volumetric masks were generated from these muscle surfaces and these were then used to mask the DTI data in preparation for denoising. Fiber tracking was performed on each muscle vector field (prior to and after denoising) using Paraview, an open source visualization toolkit developed by Kitware Inc. (Squillacote, 2008).

One limitation of our current diffusion acquisition protocol is that it fails to capture any forearm muscle in its entirety. Muscle segmentation was still possible due to the use of the aforementioned low-resolution anatomical scan which encompassed the entire arm of the subject. However we were unable to extract muscle fibers for a whole muscle. This prevented us from being able to accurately compute bulk muscle fiber metrics such as average fiber

length or physiological cross sectional area. We hope to address this issue in a future work and will discuss it further in later sections of this paper.

3. Results and discussion

Fig. 4 shows the output of each stage of the algorithm for a synthetic fusiform muscle tensor field. The top left sub figure shows the principal eigenvectors of the input tensor field. Notice that, as expected, the curl free component contains the sources and sinks at either ends of the synthetic fusiform muscle dataset. Because we are only interested in the direction of each vector we normalize the final summed vector field in order to produce a final output which is shown in the top right of Fig. 4 for comparison with the input vector field. It is also worth noting that the algorithm corrects the spurious directions present in the input vector

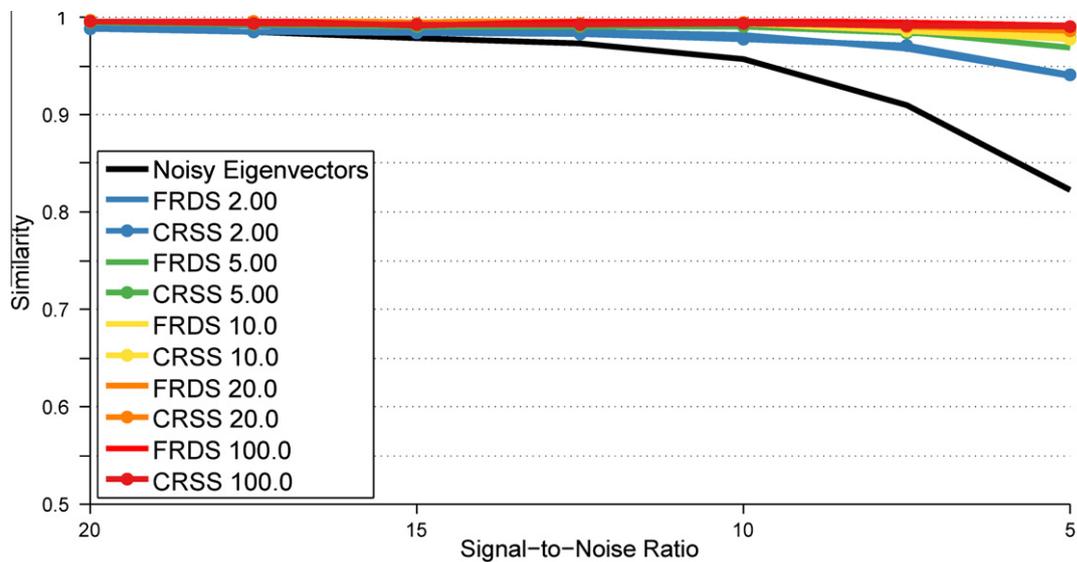


Fig. 7. Synthetic 2D bipennate vector field: The similarity of both the denoised vector field and the noisy primary eigenvector field to the original, zero noise data. The result for each algorithm is denoted by the algorithm name followed by the value of the parameter σ .

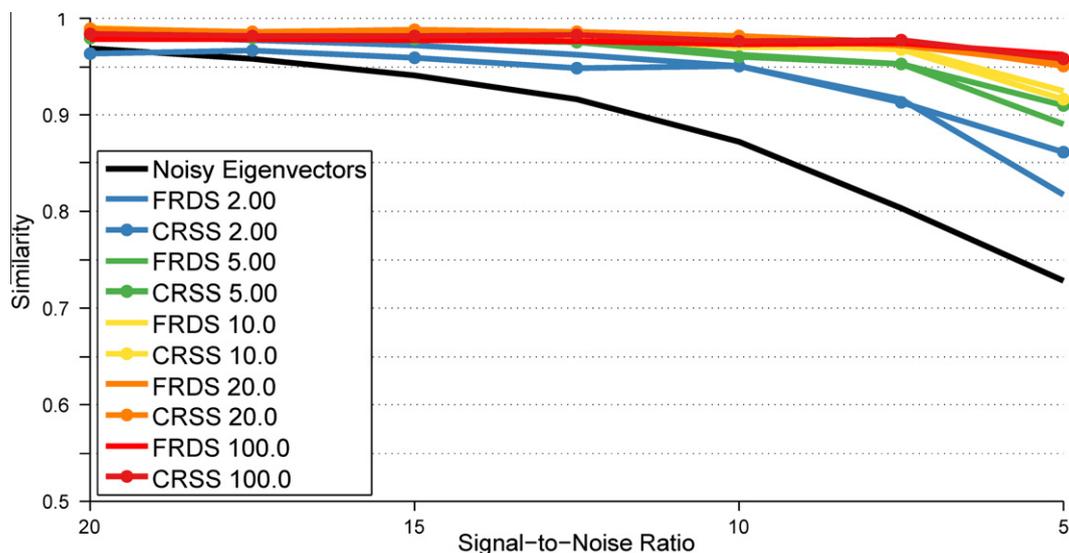


Fig. 8. Synthetic 3D bipennate vector field: The similarity of both the denoised vector field and the noisy primary eigenvector field to the original, zero noise data. The result for each algorithm is denoted by the algorithm name followed by the value of the parameter σ .

field. The consistent vector directions are a result of the initial, divergence free optimization. The gradient field of h neither converges nor diverges. This eliminates the flipped vectors in Fig. 4.

3.1. Synthetic data results

Figs. 5–8 summarize the results of applying the FRDS and the CRSS denoising algorithms to all four synthetic datasets. In all cases the following patterns can be observed. Firstly, for each combination of SNR and σ both denoising implementations produce vector fields that are more similar (Eq. (23)) to the ground truth than the primary eigenvector fields. Secondly, the larger two values of σ , 20, and 100, produce noticeably better results for low SNR. Both denoising implementations follow comparable patterns of decreasing similarity as SNR is reduced and σ is increased. However it is interesting that the CRSS algorithm produces more similar results than the FRDS implementation in the 3D bipennate case for $\sigma = 5$. The second stage of the FRDS algorithm attempts to fit RBFs at each voxel location simultaneously whereas the CRSS algorithm itera-

tively fits randomly selected RBF subsets. A sparse set of RBFs may provide a better approximation considering the small number of insertions and origins (i.e. sources and sinks) found in muscle whereas the dense sampling of the FRDS algorithm could be over-fitting the data.

Figs. 9–12 show the noisy tensor fields, the noisy eigenvectors fields and the denoised vector fields at zero noise (infinite SNR) and SNR values of 20, 10, and 5 for all synthetic datasets. Both denoising algorithms introduce some smoothing in the synthetic bipennate datasets. This is to be expected because we are attempting to fit a smooth field to the underlying tensor data. This smoothing is also a consequence of the RBF fitting stage of the algorithm which uses neighborhood information to estimate the underlying function u . Despite this, both algorithms always produce vector fields which are more similar to the original, synthetic field than the primary eigenvector fields (Figs. 5–8). As the SNR is decreased both the FRDS and the CRSS algorithms impart more smoothing but still produce results that capture the underlying structure of the synthetic datasets. These observations are supported quantita-

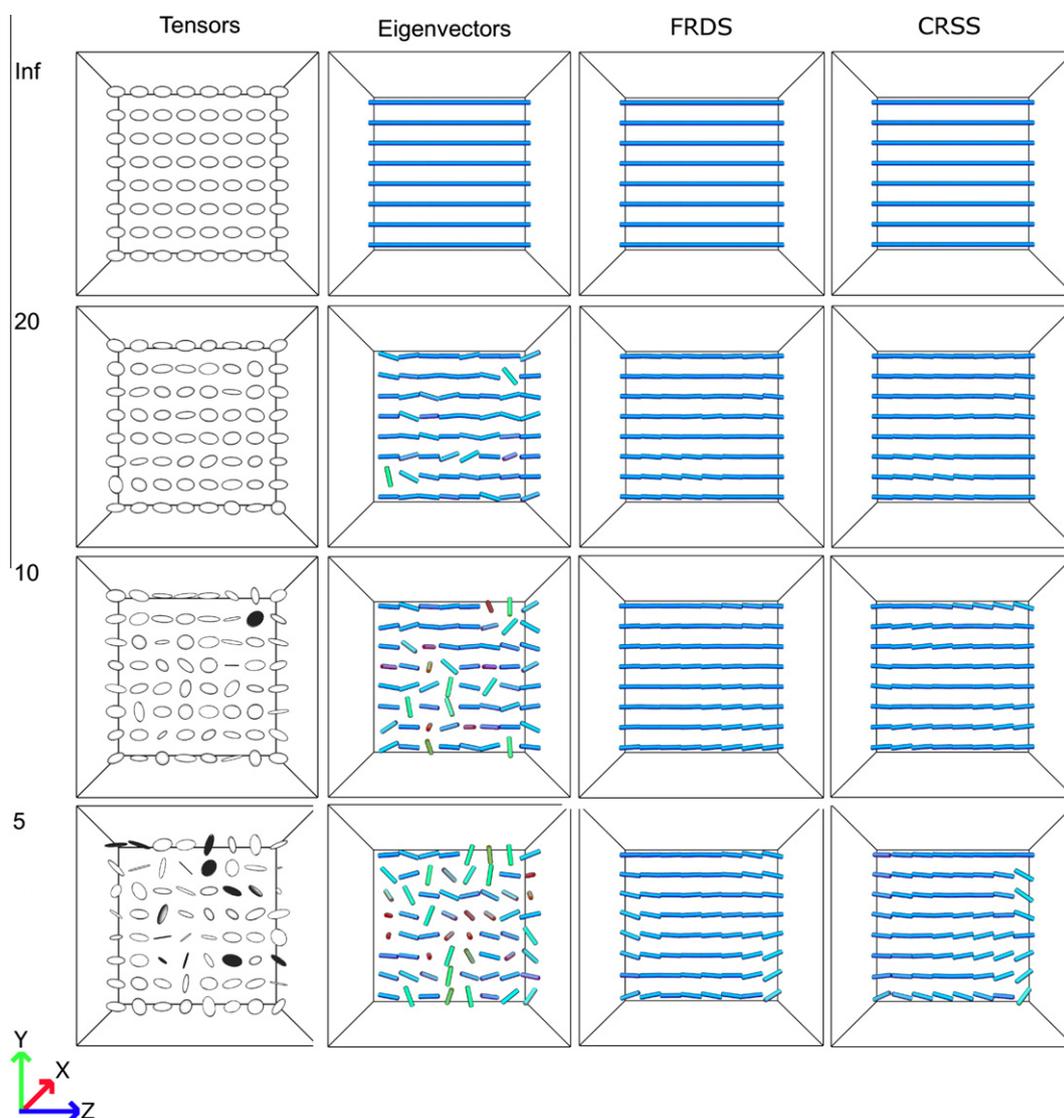


Fig. 9. A comparison of denoised results to noisy synthetic Z-Axis tensor and eigenvector fields for SNR values of 20 to 5. Each row shows images from a specific SNR with 0 noise indicated by Inf. The middle slice of each 3D dataset is shown. Displayed in order from left to right are the tensor volume, the primary eigenvectors of the tensor volume, the FRDS denoising result and the CRSS denoising result. Red, green and blue values of each vector are colored according to their magnitude in the coordinate directions. The red component is equal to the vector magnitude in x , green corresponds to the y magnitude and blue to the z magnitude. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

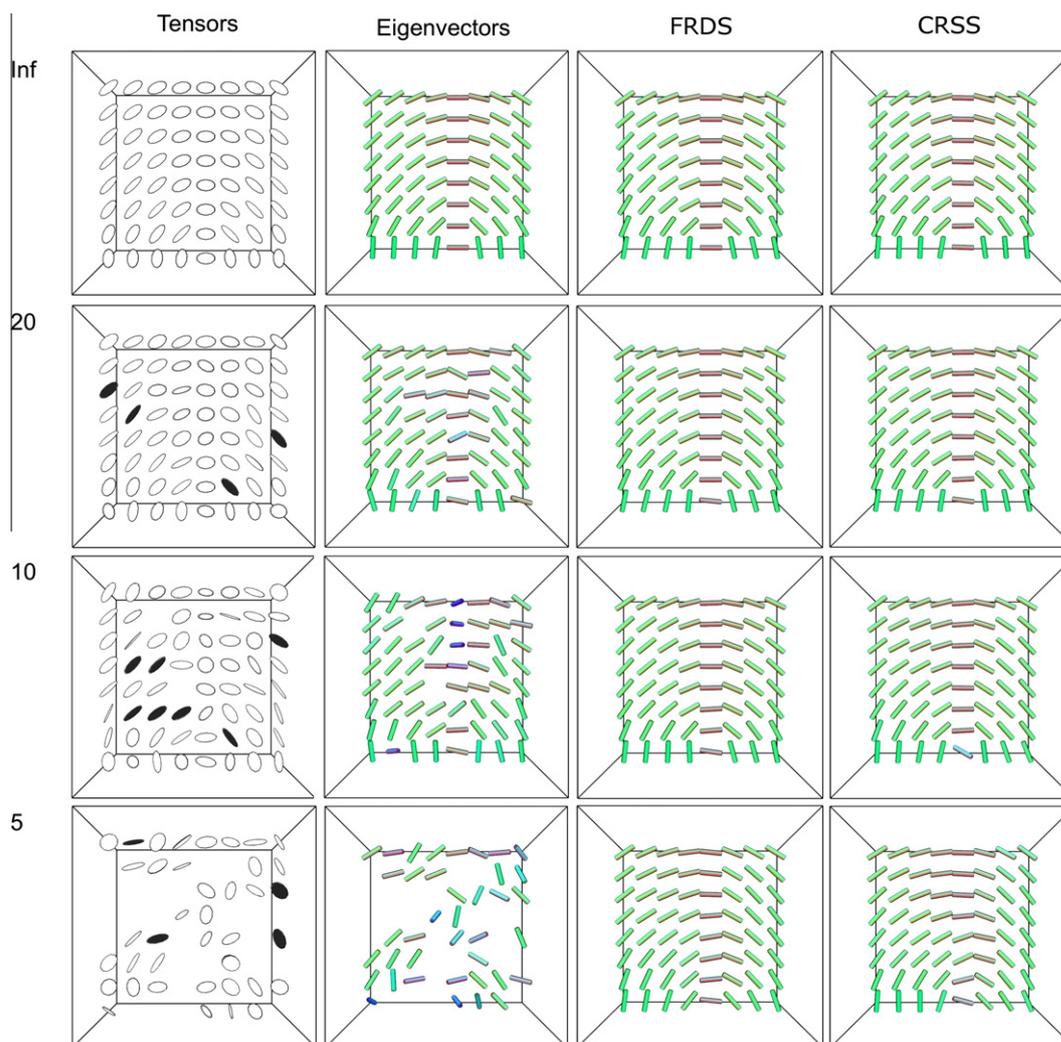


Fig. 10. A comparison of denoised results to noisy synthetic curved tensor and eigenvector fields for SNR values of 20 to 5. Each row shows images from a specific SNR with 0 noise indicated by Inf. The middle slice of each 3D dataset is shown. Displayed in order from left to right are the tensor volume, the primary eigenvectors of the tensor volume, the FRDS denoising result and the CRSS denoising result. Red, green and blue values of each vector are colored according to their magnitude in the coordinate directions. The red component is equal to the vector magnitude in x, green corresponds to the y magnitude and blue to the z magnitude. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

tively by Figs. 5–8 which show that even at low SNR values the denoised vector fields maintain a high (>0.9 for $\sigma \geq 10$) similarity with the synthetic vector fields.

Fig. 13 shows the effect that noise can have on fiber tracking algorithms. Streamlines were computed by integrating each synthetic dataset using Paraview (Squillacote, 2008). The paths were computed using a Runge–Kutta 4 integration scheme. Streamlines on the original, noise free fields follow the structure of the vector fields illustrated in Figs. 9–12. When the same procedure is repeated on the noisy (5 SNR) eigenvector fields, the results deviate greatly from the correct fiber architecture represented by the original vector field. Finally, the streamlines computed from the denoised vector field (FRDS) more closely represent the proper fiber architecture. This illustrates the dramatic effect that noise can have on fiber tracking.

Both the FRDS and CRSS algorithms were also compared with the recently published Joint Rician Linear Minimum Mean Squared Error (JRLMMSE) method of Tristán-Vega and Aja-Fernández (2010). JRLMMSE acts on Diffusion Weighted Images prior to tensor calculation in order to reduce the Rician noise. JRLMMSE was applied to the previously defined synthetic datasets (Eqs. 20, 22, 21) using the implementation of Tristán-Vega et al. available in

the 3D Slicer software platform (Pieper et al., 2006). The denoised DWIs were used to create a tensor field from which the eigenvector field was extracted. Since we are primarily interested in the performance of the presented algorithms under low SNR conditions we performed the comparison using an SNR of 5, the results of which are shown in Table 1. Parameters for JRLMMSE were manually tuned to achieve the best quantitative results. A σ value of 10 was used for the FRDS and CRSS algorithms. For this value of σ the results for the FRDS and CRSS are nearly identical (Figs. 5–8) thus we present similarity values for the FRDS only. Both algorithms produced denoised vectors that were more similar to the ground truth than those estimated by naive eigenvector computation. The results also show that the FRDS algorithm produced primary eigenvector fields that were quantitatively closer to the ground truth for three of the four datasets (Table 1). For the Z-Axis dataset JRLMMSE produces slightly better quantitative results, however for both the 2D and 3D Bipennate datasets the FRDS algorithm produces vector fields with considerably higher similarity (Table 1). These results illustrate that the FRDS and CRSS algorithms provide equivalent or better denoising performance than JRLMMSE in low SNR scenarios acting on skeletal muscle-like datasets.

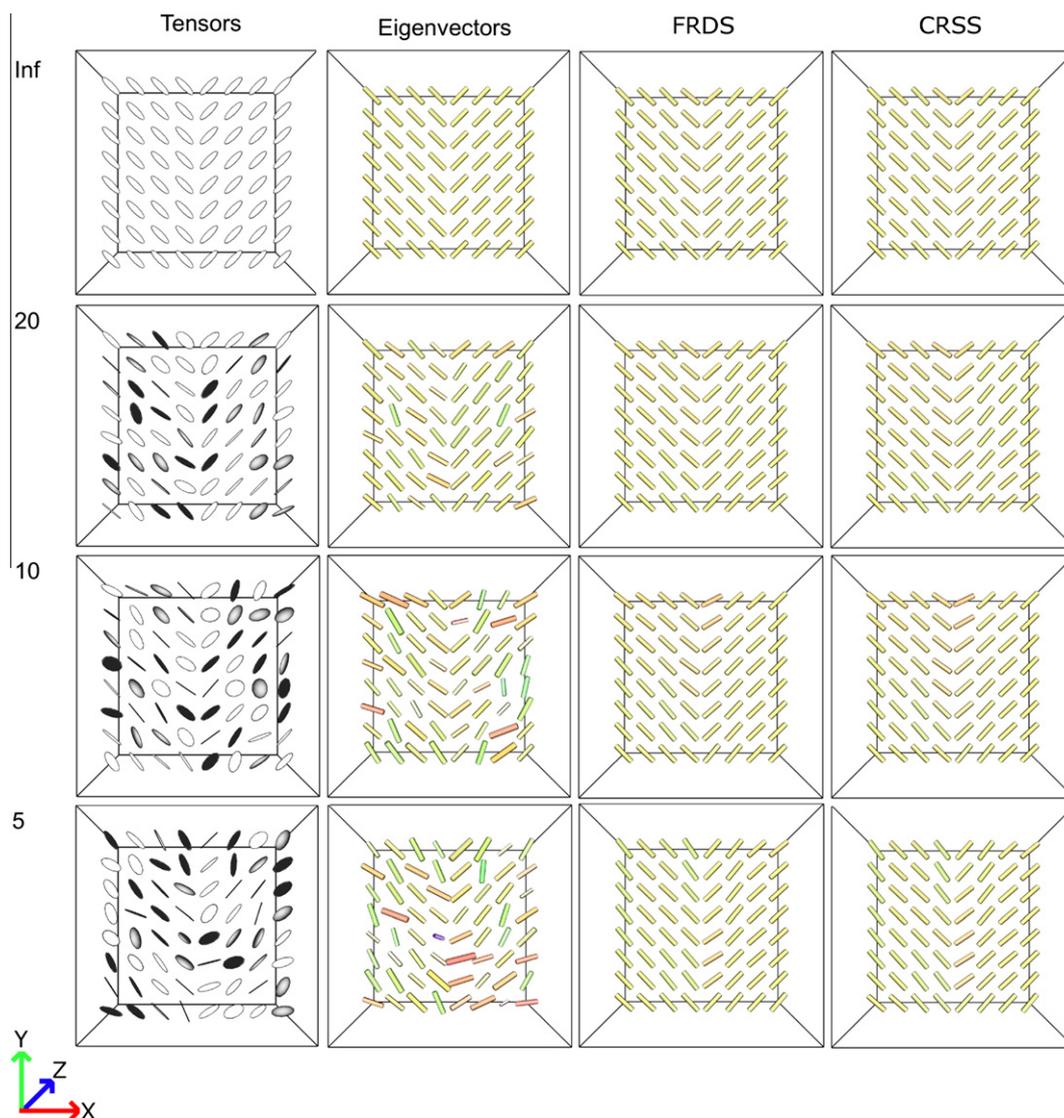


Fig. 11. A comparison of denoised results to noisy synthetic 2D bipennate tensor and eigenvector fields for SNR values of 20 to 5. Each row shows images from a specific SNR with 0 noise indicated by Inf. The middle slice of each 3D dataset is shown. Displayed in order from left to right are the tensor volume, the primary eigenvectors of the tensor volume, the FRDS denoising result and the CRSS denoising result. Red, green and blue values of each vector are colored according to their magnitude in the coordinate directions. The red component is equal to the vector magnitude in x, green corresponds to the y magnitude and blue to the z magnitude. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

3.2. Human subject data results

Fig. 14 shows the orientation of the DWI scan as well as three segmented muscle surfaces relative to the arm and hand of the subject. The three muscle surfaces shown are the brachioradialis (red), the extensor carpi radialis longus (green) and the extensor carpi radialis brevis (blue). Three slices of DWI data are also shown, the first and last slices illustrate the extent of the DWI acquisition.

The CRSS denoising algorithm was applied to DTI data computed from the forearm of a human subject. Specifically, we have produced denoised vector fields for portions of three muscles in the forearm, the ECRL, ECRB and BR. We only show subsets of each muscle due to the aforementioned limitations in the MRI acquisition protocol. Results of applying fiber tracking to both the noisy and denoised vector fields for each muscle are shown. Fiber tracking terminated at the tendon section of each muscle mesh because of the narrow structure of the geometry. The figure is oriented in the same manner as Fig. 14 such that its left edge is distal and the right edge is proximal. Most of the ECRL muscle body fell out-

side of our DWI field of view. However we do observe that the small fibers extracted from the denoised data do not exhibit the curled trajectories of those in the noisy data. The general behavior of the fibers is similar, moving from the proximal to distal side of the muscle where the ECRL inserts into the tendon that connects to the second metacarpal (Tortora and Grabowski, 2000). The removal of curl from the vector field is to be expected given the previously discussed formulation of the denoising algorithm. The ECRB data shows the effect of our algorithm more clearly, where the chaotic nature of the noisy streamlines is greatly smoothed out. Furthermore, the streamlines of the denoised ECRB can be observed converging at the distal end of the muscle, where the tendon that connects the muscle to the third metacarpal should arise (Tortora and Grabowski, 2000). In the noisy data this convergence is far less obvious. The bifurcated nature of the muscle fibers towards the distal end is similar to the fiber architecture observed by Ravichandiran et al. (2009). The denoised fibers extracted from the BR all curve towards a common insertion point at the distal end of the muscle (Tortora and Grabowski, 2000) whereas the noisy

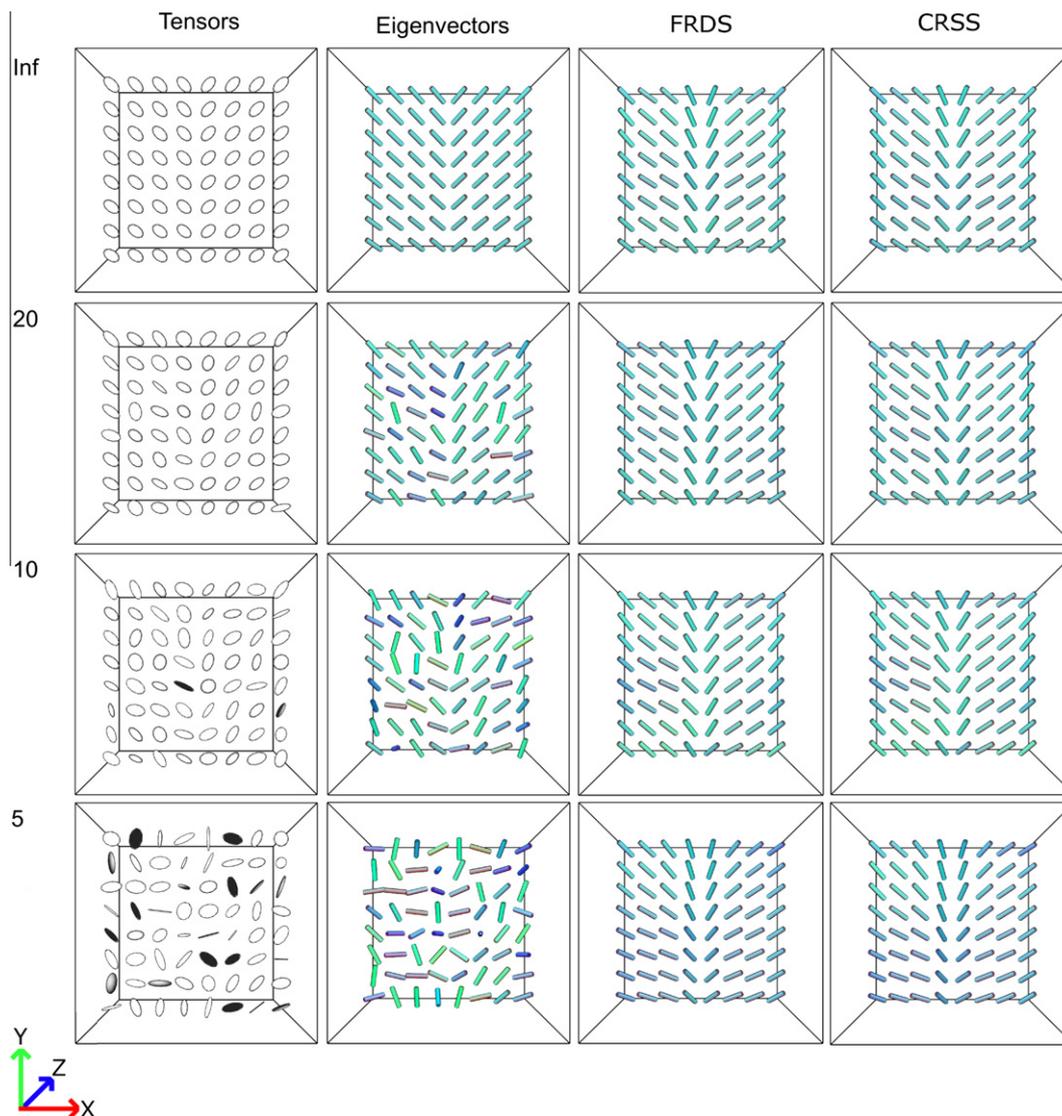


Fig. 12. A comparison of denoised results to noisy synthetic 3D bipennate tensor and eigenvector fields for SNR values of 20 to 5. Each row shows images from a specific SNR with 0 noise indicated by Inf. The middle slice of each 3D dataset is shown. Displayed in order from left to right are the tensor volume, the primary eigenvectors of the tensor volume, the FRDS denoising result and the CRSS denoising result. Red, green and blue values of each vector are colored according to their magnitude in the coordinate directions. The red component is equal to the vector magnitude in x, green corresponds to the y magnitude and blue to the z magnitude. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

data exhibits fibers that appear to diverge (particularly at the proximal edge of the image). Note that the denoised fibers are corrected so that they follow a path towards the distal end of the muscle. This fiber configuration is in agreement with the anatomical description of the BR in that the muscle inserts into the brachioradialis tendon at its distal end (Tortora and Grabowski, 2000).

Extraction of a physiologically plausible, noise-reduced vector field from noisy tensor data is not often explored. Both Basu et al. (2006) and Fillard et al. (2007) use probabilistic estimators to remove noise from the tensor image. They assume that the noise has a Rician distribution. This method has the advantage of not requiring an anatomical segmentation to produce good results. It also works well on DWI images from any anatomical region whereas the presented algorithm is designed specifically for musculoskeletal DTI. However, our method uses physical intuition to produce denoised results and thus should work well even when image noise is not Rician. Furthermore Basu et al. (2006) demonstrate results down to an SNR of 12.5 but only present tensor metrics such as fractional anisotropy as results, making direct

comparison impossible. They do provide images of the algorithm applied to synthetic images with an SNR of 15. We provide results for synthetic data with SNR as low as 5 and show the algorithm computes reasonable results. Finally the presented algorithm produced quantitatively better results on three of four synthetic datasets than the state-of-the-art JRLMMSE algorithm of Tristán-Vega and Aja-Fernández (2010) (Table 1). The similarity difference for the fourth dataset was 0.03. These results indicate the presented method competes well with tensor denoising algorithms in terms of being resilient to noise. However, it is important to note that the presented method does not replace existing tensor denoising algorithms in the imaging pipeline. The denoised vector fields offer both resistance to noise and, due to the divergence free constraint, are optimally “muscle like”. This algorithm could be applied to previously denoised tensors to yield a vector field that is more accurate than what could be achieved with either algorithm separately. This vector field could then be used for fiber tracking and tractography in conjunction with the denoised tensor field.

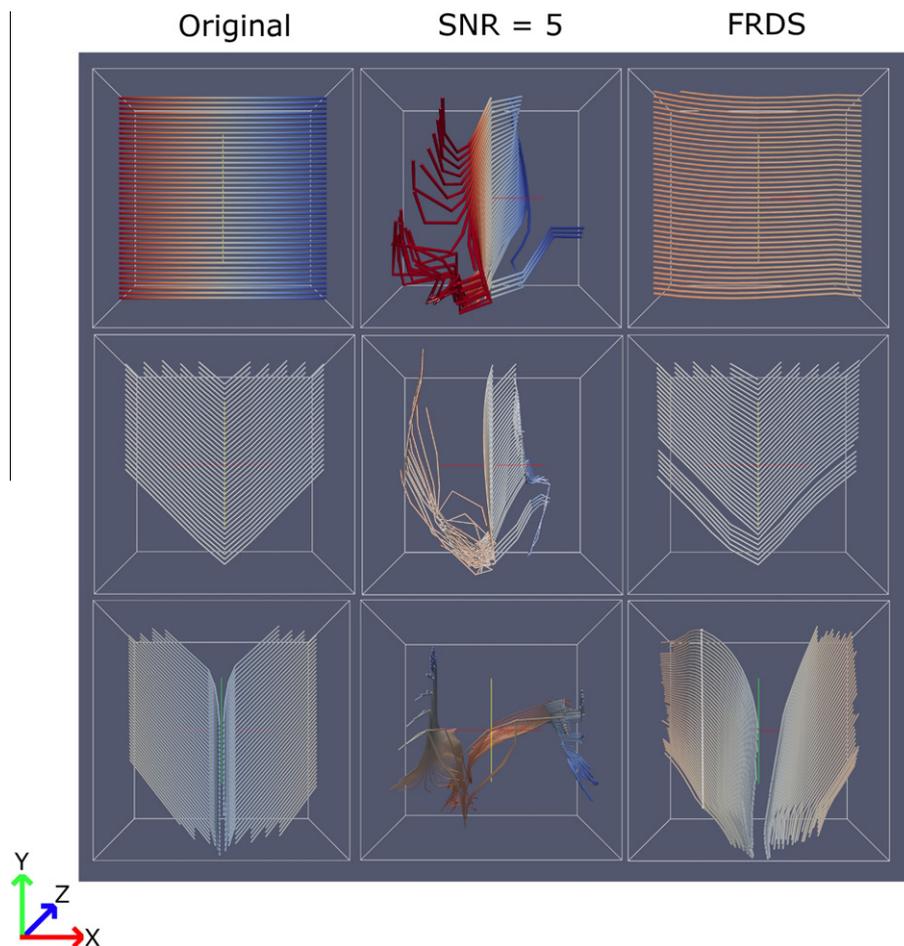


Fig. 13. The effect of vector field noise on fiber tracking. Each row of this image shows the result of applying fiber tracking to a synthetic dataset. Fiber tracking results are shown for clean data, noisy data with an SNR of 5 and finally for data denoised using the FRDS algorithm.

Table 1

A comparison of denoised vector fields produced by the Full Resolution, Dense Sampling algorithm and the Joint Rician Linear Minimum Mean Squared Error algorithm (Tristán-Vega and Aja-Fernández, 2010). Results are shown for four synthetic datasets with signal-to-noise ratio equal to 5.

	Z-Axis	2D BiPenn.	Curves	3D BiPenn.
FRDS	0.978	0.981	0.999	0.970
JRLMMSE	0.981	0.955	0.843	0.772

3.3. Limitations and future work

The presented algorithms have two main limitations. The first is that the underlying theory assumes that a significant part of the

vector field to be reconstructed is smooth. While we have shown good results on synthetic data with a discontinuity (Figs. 11 and 12) there is no way to entirely prevent smoothing. In practice this is less of a problem since we segment and process muscles separately and thus avoid the discontinuities caused by boundaries between muscles. The second limitation stems from the fact that we are reconstructing conservative vector fields. Conservative vector fields represent all irrotational flows within a simply connected domain, a domain without holes. The anatomical structure of most skeletal muscles observes this property however there are some exceptions such as sphincters, the ciliary muscle in the eye and the heart. A productive direction of future work would be to relax this restriction. The reliance on an anatomical segmentation is also a limitation. This algorithm's primary use is to provide more accu-

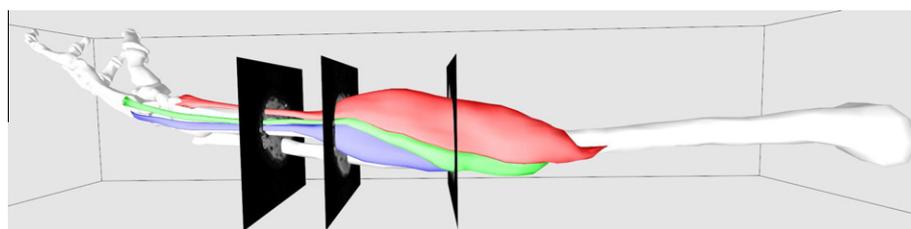


Fig. 14. The relative position of the acquired DWI scan as well as the muscle surfaces used for denoising and fiber tracking. The first, last and an arbitrarily chosen middle slice of the dataset are shown. The displayed muscles are the brachioradialis (red), extensor carpi radialis longus (green) and extensor carpi radialis brevis (blue). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

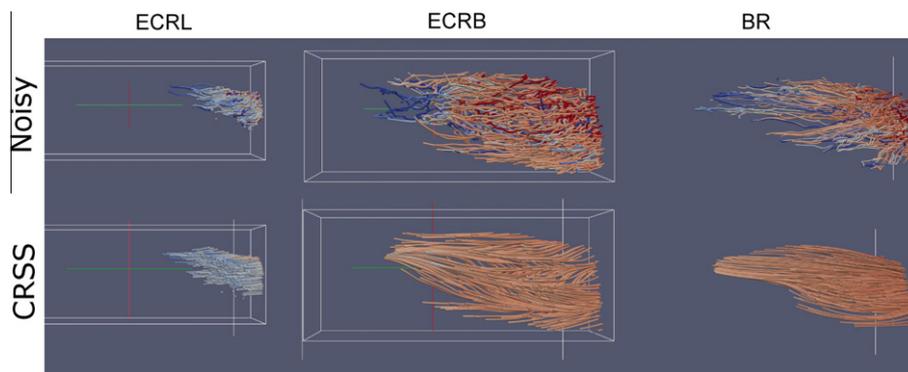


Fig. 15. Comparison of noisy and denoised (CRSS) vector field results using streamlines. Available portions of three muscles are shown, the distal end of the ECRL, and larger portions of the ECRB and the BR. the left side of the image is towards the distal end of the arm and the right hand side is towards the proximal end.

rate vector fields for fiber tracking. It is fair to assume that such a segmentation will be available because fiber tracking algorithms often require a segmented surface as a constraint.

While the results produced by the algorithm are promising there are several areas of future work we intend to explore. Firstly, the parameter σ is currently user defined. It would be advantageous to incorporate this into the optimization scheme using a prior model. Secondly, we would like to modify our MRI acquisition protocol in order to capture entire forearm muscles in the diffusion scan. This would allow the extraction of denoised fiber fields from entire muscles, facilitating the computation of physiological parameters and enabling further validation of the denoised fiber field results against the existing literature (Lieber et al., 1992; Ravichandiran et al., 2009). The effect of vector field noise on these physiological parameters could then be examined. Thirdly, the algorithm requires an SPD tensor field as input. Extending the algorithm to operate directly on diffusion weighted images could simplify vector field extraction further as well as lead to increased resistance to noise. Finally, we also intend to perform computational simulations of muscle behavior and examine how extracted fiber architectures affect the function of muscle.

Despite these limitations, the results for both synthetic and anatomical data experiments are promising. When run on synthetic data both the deterministic algorithm and its large scale variant show good similarity (>0.9 Figs. 5–8) for smoothing parameter values of 10 or greater. These results were visually verified as well (Figs. 9–12). Furthermore, fibers extracted from denoised vector fields computed from human forearm DTI data show greater smoothness than those extracted from the noisy data. The results for the ECRB and the BR are especially interesting because the denoised fields appear to either correct errors in the noisy vector fields (BR) or highlight salient features of the fiber architecture such as in the ECRB case (Fig. 15).

4. Conclusion

We have designed and implemented algorithms for constructing denoised skeletal muscle fiber vector fields from noisy diffusion tensor images. Using synthetic data we have shown that the denoised vector fields provide a much better approximation (>0.9 similarity with $\sigma \geq 10$) to ground truth noiseless data for SNR values greater than or equal to 5. Furthermore we have shown that the algorithms offer denoising performance exceeding that of a state-of-the-art DWI image filter at low SNR for muscle-like datasets. Finally, we show that the algorithms produce plausible results when applied to DTI data of a human forearm.

Acknowledgments

The authors would like to acknowledge Garrett Jones, Danny Kaufman and Shinjiro Sueda for helpful comments during the development of this manuscript. This work was supported in part by the Canada Research Chairs Program, NSERC, The Peter Wall Institute for Advanced Studies, Canada Foundation for Innovation and BC KDF.

References

- Anderson, E., Bai, Z., Bischof, C., Blackford, S., Demmel, J., Dongarra, J., Croz, J.D., Greenbaum, A., Hammarling, S., McKenney, A., Sorensen, D., 1999. LAPACK Users' Guide, third ed. Society for Industrial and Applied Mathematics, Philadelphia, PA.
- Basser, P.J., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A., 2000. In vivo fiber tractography using dt-mri data. *Magnetic Resonance in Medicine* 44 (4), 625–632.
- Basu, S., Fletcher, T., Whitaker, R., 2006. Rician noise removal in diffusion tensor mri. In: *Medical Image Computing and Computer-Assisted Intervention, MICCAI 2006*, pp. 117–125.
- Bergmann, O., Lundervold, A., Steihaug, T., 2005. Generating a synthetic diffusion tensor dataset. In: *Proceedings of the 18th IEEE Symposium on Computer-Based Medical Systems*, 2005, pp. 277–281.
- Blemker, S.S., Delp, S.L., 2005. Three-dimensional representation of complex muscle architectures and geometries. *Annals of Biomedical Engineering* 33, 661–673.
- Blemker, S.S., Asakawa, D.S., Gold, G.E., Delp, S.L., 2007. Image-based musculoskeletal modeling: applications, advances, and future opportunities. *Journal of Magnetic Resonance Imaging* 25 (2), 441–451.
- Brand, P.W., Beach, R.B., Thompson, D.E., 1981. Relative tension and potential excursion of muscles in the forearm and hand. *Journal of Hand Surgery* 3 (6), 209–219.
- Byrd, R.H., Lu, P., Nocedal, J., Zhu, C., 1999. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing* 16, 1190–1208.
- Chefd'hotel, C., Tschumperl, D., Deriche, R., Faugeras, O., 2002. Constrained flows of matrix-valued functions: application to diffusion tensor regularization. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (Eds.), *Computer Vision ECCV 2002, Lecture Notes in Computer Science*, vol. 2350. Springer, Berlin/Heidelberg, pp. 251–265.
- Damon, B.M., Ding, Z., Anderson, A.W., Freyer, A.S., Gore, J.C., 2002. Validation of diffusion tensor MRI-based muscle fiber tracking. *Magnetic Resonance in Medicine* 48 (1), 97–104.
- Ding, Z., Gore, J.C., Anderson, A.W., 2001. Case study: reconstruction, visualization and quantification of neuronal fiber pathways. In: *VIS '01: Proceedings of the Conference on Visualization '01*. IEEE Computer Society, Washington, DC, USA, pp. 453–456.
- Fillard, P., Pennec, X., Arsigny, V., Ayache, N., 2007. Clinical dt-mri estimation, smoothing, and fiber tracking with log-euclidean metrics. *IEEE Transactions on Medical Imaging* 26 (11), 1472–1482.
- FMRIB, 2006. Fmrib's diffusion toolbox – fdt v2.0. <<http://www.fmrib.ox.ac.uk/fsl/fdt/index.html>>.
- Gilles, B., Pai, D.K., 2008. Fast musculoskeletal registration based on shape matching. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'08)*, pp. 822–829.
- Harlow, F.H., Welch, J.E., 1965. Numerical calculation of time-dependent viscous incompressible flow of fluid with a free surface. *Physics of Fluid* 8 (12), 2182–2189.
- Jenkinson, M., Smith, S.M., 2001. A global optimisation method for robust affine registration of brain images. *Medical Image Analysis* 2 (5), 143–156.

- Lansdown, D.A., Ding, Z., Wadington, M., Hornberger, J.L., Damon, B.M., 2007. Quantitative diffusion tensor mri-based fiber tracking of human skeletal muscle. *Journal of Applied Physiology* 103 (2), 673–681.
- Lieber, R.L., Jacobson, M.D., Fazeli, B.M., Abrams, R.A., Botte, M.J., 1992. Architecture of selected muscles of the arm and forearm: anatomy and implications for tendon transfer. *The Journal of Hand Surgery* 17A, 787–798.
- McGraw, T., Vemuri, B., Ozarslan, E., Chen, Y., Mareci, T., 2009. Variational denoising of diffusion weighted MRI. *Inverse Problems and Imaging* 3 (4), 625–648.
- Moré, J.J., Garbow, B.S., Hillstom, K.E., 1980. User Guide for MINPACK-1, Report ANL-80-74. Argonne, Illinois, USA, <<http://www-fp.mcs.anl.gov/otc/Guide/softwareGuide/Blurbs/minpack.html>>.
- Mori, S., van Zijl, P., 2002. Fiber tracking: principles and strategies – a technical review. *NMR in Biomedicine* 15 (7–8), 468–480.
- Neji, R., Azzabou, N., Paragios, N., Fleury, G., 2007. A convex semi-definite positive framework for DTI estimation and regularization. *Advances in Visual Computing*, 220–229.
- Pieper, S., Lorensen, B., Schroeder, W., Kikinis, R., 2006. The na-mic kit: Itk, vtk, pipelines, grids and 3d slicer as an open platform for the medical image computing community. In: 3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro 2006, pp. 698–701.
- Ravichandiran, K., Ravichandiran, M., Oliver, M.L., Singh, K.S., McKee, N.H., Agur, A.M.R., 2009. Determining physiological cross-sectional area of extensor carpi radialis longus and brevis as a whole and by regions using 3d computer muscle models created from digitized fiber bundle data. *Computer Methods and Programs in Biomedicine* 95 (3), 203–212.
- Squillacote, A.H., 2008. *The Paraview Guide*. Kitware Inc.
- Tong, Y., Lombeyda, S., Hirani, A.N., Desbrun, M., 2003. Discrete multiscale vector field decomposition. *ACM Transactions on Graphics* 22 (3), 445–452.
- Tortora, G.J., Grabowski, S.R., 2000. *Principles of Anatomy and Physiology*, ninth ed. Biological Sciences Textbooks Inc.
- Tristán-Vega, A., Aja-Fernández, S., 2010. Dwi filtering using joint information for dti and hardi. *Medical Image Analysis* 14 (2), 205–218.
- Tschumperlé, D., Deriche, R., 2003. Variational frameworks for DT-MRI estimation, regularization and visualization, pp. 116–121.
- Wang, Z., Vemuri, B., Chen, Y., Mareci, T., 2004. A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from complex dwi. *IEEE Transactions on Medical Imaging* 23 (8), 930–939.
- Weinstein, D., Kindlmann, G., Lundberg, E., 1999. Tensorlines: advection-diffusion based propagation through diffusion tensor fields. In: *VIS '99: Proceedings of the Conference on Visualization '99*. IEEE Computer Society Press, Los Alamitos, CA, USA, pp. 249–253.
- Westin, C.-F., Maier, S.E., Mamata, H., Nabavi, A., Jolesz, F.A., Kikinis, R., 2002. Processing and visualization of diffusion tensor MRI. *Medical Image Analysis* 6 (2), 93–108.