# Perception-based Visualization of High-Dimensional Medical Images Using Distance Preserving Dimensionality Reduction

Ghassan Hamarneh, Chris McIntosh, and Mark S. Drew School of Computing Science, Simon Fraser University, Canada

#### Abstract

A method for visualizing high dimensional medical image data is proposed. The method operates on images in which each pixel contains a high dimensional vector, e.g. a time activity curve (TAC) in a dynamic positron emission tomography (dPET) image, or a tensor, as is the case in diffusion tensor magnetic resonance images (DTMRI). A nonlinear mapping reduces the dimensionality of the data to achieve two goals: Distance preservation and embedding into a perceptual color space. We use multidimensional scaling distance preserving mapping to render similar pixels (e.g. DT or TAC pixels) with perceptually similar colors. The 3D CIELAB perceptual color space is adopted as the range of the distance preserving mapping, with a final similarity transform mapping colors to a maximum gamut size. Similarity between pixels is determined analytically as geodesics on the manifold of pixels or approximated using manifold learning techniques. In particular, dissimilarity between DTMRI pixels is evaluated via a Log-Euclidean Riemannian metric respecting the manifold of the rank 3, 2nd order positive semi-definite DTs. Dissimilarity between TACs is approximated via ISOMAP. We demonstrate our approach via artificial highdimensional data, as well as clinical DTMRI and dPET images. Our results demonstrate the effectiveness of our approach in capturing, in a perceptually meaningful way, important structures in the data.

*Keywords:* High dimensional data, visualization, color, nonlinear dimensionality reduction, multidimensional scaling, diffusion tensor magnetic resonance imaging (DTMRI), dynamic positron emission tomography (dPET).

## I. INTRODUCTION AND MOTIVATION

High-dimensional data is becoming more prevalent in general. In medical imaging applications, in particular, the popularity of high-dimensional data is due to several reasons. Advances in acquisition hardware, such as faster electronics and more sensitive sensors, e.g. Gamma cameras, allow collection of more anatomical and functional data samples without increasing patient scan-time. Image acquisition manufacturers are now even designing hybrid imaging machines incorporating multiple modalities, e.g. hybrid single-photon emission and x-ray computed tomography (SPECT/CT) cameras. Further, multiple medical imaging modalities are also being spatially aligned, using medical image registration algorithms, and fused together providing complementary information about the underlying structures and biological processes within, e.g. T1-weighted, T2-weighted, and Positron Density magnetic resonance imaging (T1/T2/PD-MRI) is fused with functional MRI (f-MRI). Furthermore, hyper (or multi)-spectral or multi-energy imaging is becoming more common because it reveals more clinically useful information, as in, for example, the use of tracers or pharmaceuticals with dual or multiple isotopes in molecular imaging. The ability to capture and reconstruct dynamic behaviors/kinematics of tracers in tissues is resulting in

time activity curves (TAC) capturing photon counts recorded at each pixel, e.g. dynamic positron emission tomography (dPET) and d-SPECT. Additionally, new imaging protocols and modalities are designed from the outset to capture more exquisite data about the living body, such as brain white matter or cardiac muscle fiber organization obtained from diffusion tensor MRI (DTMRI) or high angular resolution diffusion imaging (HARDI). This explosion in the capabilities of medical imaging is, in turn, resulting in higher dimensional pixels and more complex spatial fields.

For patient-specific diagnosis or therapy, as well as for statistical population-based medical studies, there is a pressing need to efficiently interpret and analyze such data having increasing dimensionality and resolution. More specifically, there is a need for accurate, repeatable, and fast medical image analysis and interpretation algorithms. Medical image segmentation, which partitions an image into (two or more) different regions, is typically a necessary precursor to performing higher level representation and understanding of shapes and images. Segmentation methods typically rely on (i) identifying pixels with similar properties (e.g. CT pixel intensity in Hounsfield Units, DT, or TAC) and grouping them into homogenous regions; (ii) identifying local regions of pixel dissimilarities, or edges, and linking them to form separating boundaries between regions; and (iii) incorporating some form of prior knowledge of the different structural or functional regions to be segmented (e.g. prior knowledge of shape, appearance, spatial relationships, or temporal dynamics, as well as expert or domain-based knowledge). There are numerous techniques that attempt to automate this important segmentation step at the crux of the medical image interpretation task, including region-based approaches (e.g. seed-based region growing), graphtheoretic methods (e.g. graph cuts, random walker, intelligent scissors), pattern-recognition approaches (clustering and classification), atlas- or registration-based techniques, energy-minimizing methods (e.g. parametric or level-sets based deformable models) [24]. The criteria of a correct segmentation are not universal, although (i-iii) are typically incorporated in these algorithms in different forms.

The majority, if not all, of existing segmentation methods and subsequent interpretation methods rely on and are sensitive to user-initialized seeds, contours, or gestures, setting of low level parameters, and/or on an expert validated set of training data. Existing segmentation algorithms are yet to achieve full automation while producing completely correct results, and hence cannot be relied upon in clinical settings without user intervention. The fact remains that clinicians and radiologists continue to interpret medical images by relying largely on visual assessment, in which image partitioning and pixel grouping still relies on pixel dis/similarities and prior or expert domain knowledge, i.e. i-iii above. However, the segmentation is carried out largely in an implicit, subconscious way in "the expert's brain", in a way that is still not well understood, but known to benefit greatly from years of medical training and experience. Therefore, it is important to provide clinicians and radiologists with a medical image display or visualization system that is faithful to the underlying structural or functional data. The expert's image interpretation or diagnosis must not be affected by artificial manifestations of the visualization system, e.g. using inappropriate pseudo coloring or other rendering effects, that may either camouflage delicate information (false negatives, e.g. missing a tumor) or artificially introduce information (false positives) thereby misguiding their analysis.

In this work, we address these issues by proposing a high-dimensional medical image data visualization approach that (a) is faithful to the underlying medical image data; (b) respects models of human perception; and (c) relies on the essential information needed for segmentation (i-iii above: pixel dis/similarity and domain expertise). More specifically, in our method, we display high-dimensional medical image data to the domain expert as color images. The transformation of the high-dimensional data to human perception is facilitated via a nonlinear mapping that reduces the dimensionality of the data to three color channels in a way that preserves distances between pairs of pixels on high-dimensional manifolds and at the same time ensures data embedding into a perceptual color space.

# II. RELATED WORK

There exist numerous approaches for visualization of medical images and high-dimensional data and a complete survey is beyond the scope of this paper. We focus our review of the state of the art on approaches related to using color to visualize medical or general high-dimensional data. In [42], Wong and Bergeron surveyed multidimensional multivariate visualization techniques prior to 1997; their survey did not address medical image visualization in particular. In [18], Keim explored the formal basis and design decisions for the visualization of high-dimensional data using pixel-oriented visualization, which describes how pixel data is arranged and rendered on the screen.

Several methods have used color to visualize high-dimensional data by forming color channel responses through a linear projection of the original data onto basis functions. In [14], hyper-spectral data was rendered through projection on bases that reflect, through color-matching functions [43], what the human visual system would perceive had it been able to cover the range of the hyper-spectral data. Jacobson et al. proposed a method to design fixed, as in human vision, basis functions [15]. As early as 1973 [27], different methods were proposed that rely on representing the three color channels through the first three main data variation modes obtained through principal component analysis (PCA) [38]. Other methods relied on PCA in conjunction with wavelet methods, for example, by performing PCA on wavelet subbands to enhance edges at specific levels of detail [11] or to perform wavelet-based denoising followed by PCA [17]. As an alternative to PCA, independent component analysis (ICA) was performed as a means to reduce the dimensionality of the high-dimensional data into three color channels [40], [7]. Other approaches for high-dimensional visualization relied on providing iconic representation for each data point [9] or multi-colored texture elements [12]. In [37], the mapping, from high-dimensional pixel data to 3D color, which maximizes the mutual information between the original hyper-spectral bands and the color channels was used. In [10], Fang et al. used multi-dimensional scaling (MDS) as an alternative, nonlinear approach to dimensionality reduction [3]. In their approach, however, high-dimensional pixels, representing samples of temporal activity, are mapped to 2D in a way that preserves weighted distances between pixel locations and pixel dissimilarity. The 2D embedding is then used as a widget for specifying transfer functions for volume rendering.

To the best of our knowledge, works most related to our algorithmic approach are those of Rasche et al. and Brun et al. In [26], [25], color images (the higher-dimensional data in their case), are reduced for color deficient, mono- and di-chromats displays. In their paper they highlighted the importance of preserving contrast and maintaining luminance consistency. They based their method on the premise that perceived color difference between any pair of colors should be proportional to their perceived gray difference. An MDS-inspired objective function capturing this relationship was formulated and solved via constrained majorization [26]. In [4], a method for visualizing DTMRI fiber "traces" was proposed, in which a set of fiber traces is mapped to lower dimensional Euclidean space using Laplacian eigenmaps. The mapping was such that similar traces (defined as those with similar endpoints) were mapped to similar points in low-dimensional RGB color space. In contrast, our proposed method is not specific to color images or fibre traces, but rather to generic high-dimensional image fields, with a focus in this paper on medical image data (DTMRI and dPET, in particular). More importantly, we map the high-dimensional data, in a nonlinear, distance preserving way, into 3D perceptual color space and not to 1D or 2D and not to RGB. Further, we utilize any available knowledge of the manifold structure of the underlying pixels, e.g. diffusion tensors dissimilarity between pixels is evaluated via a Log-Euclidean Riemannian metric respecting the manifold of the rank 3, second order positive semi-definite DTs. When such knowledge about the manifold structure is absent, we resort to manifold learning techniques to capture dissimilarity between high-dimensional pixels.

## III. METHOD

#### A. Overview

Our objective is to present clinicians with images that are displayed and perceived in a way that best reflects the underlying medical image data. Given our focus on high-dimensional medical image data (e.g. DTMRI or dPET), where the pixel dimensionality is larger than the three color dimensions, we need to employ dimensionality reduction. Our goal is that after dimensionality reduction, pixels with similar DT or TAC pixels should be rendered with colors that are perceived similarly, and vice versa. More generally, we wish to display image pixels to the user such that pixels with similar high-dimensional data are rendered using perceptually similar colors (and different pixels using perceptually different colors). This raises two questions: (i) how to measure pixel dissimilarity and (ii) how to map pixels (with known dissimilarities) to perceptually meaningful colors. Our method addresses these two issues as follows. First, we assume that the high-dimensional pixel values are samples from an underlying manifold endowed with a distance metric. The manifolds are either learned (using manifold learning techniques, e.g. ISOMAP [36] and Locally Linear Embedding (LLE) [29]), or derived analytically, or approximated, based on the knowledge of the underlying data. Dissimilarity between any two high-dimensional pixels is measured as the geodesic distance between the two corresponding points on the manifold. We evaluate the similarity between DT pixels, in particular, via the Log-Euclidean Riemannian metric, which respects the rank 3 manifold of the DTs. For dPET, we use ISOMAP to learn the underlying manifold and approximate the distance between two TACs. Second, given known distance or dissimilarity between any pair of data points, we rely on a distance preserving mapping into perceptual color space. We use MDS for distance-preserving mapping in order to render similar DT or TAC pixels with perceptually similar colors. The 3D CIELAB perceptual color space is adopted as the range of the MDS mapping. A final rotation, scaling, and translation is still available without changing relative distance, and such a similarity transform is chosen so as to maximize the color gamut volume occupied.

# B. High-dimensional Medical Image Data

We focus on 2D or 3D fields of N pixels, i.e. each N-pixel image is represented as  $f(\mathbf{x}) : \mathbf{x} \in \mathbf{R}^d \to \mathbf{R}^n$ , where d is the spatial dimension 2 or 3, with  $\mathbf{x} = (x, y)$  or  $\mathbf{x} = (x, y, z)$ , respectively, and n is the dimensionality of the pixel data. At each location a vector  $f(\mathbf{x}) = [f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_n(\mathbf{x})]$  is sampled. Note that the intrinsic dimensionality  $\tilde{n}$  of the data can not be greater than n, i.e.  $\tilde{n} \leq n$ . For cases when n = 1, 2, or 3, then one, two or three color channels can be used without the need for nonlinear dimensionality reduction. For DTMRI data in particular, n = 6 and  $f(\mathbf{x}) : \mathbf{x} \in \mathbf{R}^d \to \mathbf{R}^6$ , since DTs are symmetric 3 matrices with 6 unique elements. Further, DTs must be positive semi-definite (with nonnegative eigenvalues) since they are interpreted as covariance matrices of 3D Gaussian probability density functions (PDF). The PDF models the probability of a water molecule diffusing to a particular location in 3D in a given time due to the underlying Brownian motion of molecules [33], [2]. This positive semi-definiteness, in turn, results in DTs being restricted to a convex half cone in 6D [23]. Therefore, dissimilarity between diffusion tensors must be calculated in a way that respects this underlying manifold.

#### C. Dissimilarity between high-dimensional Medical Image Data: Manifold Learning and Distance Metrics

An important issue in dealing with high-dimensional data is how to measure dissimilarity between observations. We distinguish between two primary cases: (a) The space of *n*-D observations forms a vector (linear) space and (b) when the space is nonlinear. In the vector space case  $L^p$  norms such as  $L^1$  (Manhattan),  $L^2$  (Euclidean) or related Chebyshev distance (Chessboard distance) can be used. In the nonlinear case, the observed variables do not form a vector space but rather their allowable values

are governed by nonlinear relationships forming an  $\tilde{n}$ -D subspace within the embedding n-D space, with intrinsic dimensionality  $\tilde{n} \leq n$ . For example, in molecular dynamic imaging applications, such as dPET, the TAC dimensionality n could be, say, 50 (i.e. 50 time samples). Nevertheless, in molecular imaging studies it is assumed that the underlying biological process can be modeled by a few kinetic parameters (e.g. 4 in a 2-compartment model) describing the partial differential equation of tracer transport, tissue perfusion, or tracer binding [5]. In the nonlinear case, we distinguish between two subclasses: (b1) The geodesic distance on the nonlinear manifold (or dissimilarity between data points) can be calculated analytically or approximated. In the case of DTMRI, for example, the distance on the manifold of PSD matrices is well defined and can be approximated numerically [1], [44], [41], [19]. In dPET, the dissimilarity may be formulated to reflect difference in functional behavior, through difference between kinetic parameters or system response or using other TAC dissimilarity metrics [10]. (b2) The underlying manifold and geodesic distance (or dissimilarity between points) are unknown, but many data samples are available. In this case, methods for learning the manifold structure are needed in order to allow for estimating geodesic distances and dissimilarity metrics. Given the locally Euclidean property of a manifold, the geodesic distances between two distant data points on the manifold can be approximated by the smallest possible aggregate of Euclidean hops between pairs of neighboring data points that connect the two distant points from start to finish (i.e. geodesic distance approximated by the shortest path made up of small Euclidean hops). This common approach requires the construction of a graph whose vertices represent the highdimensional sample points and whose edge weights are equal to the Euclidean distance between the two

Given the set of N high-dimensional, n-D, pixels, such as some or all of the pixels of a DTMRI or a dPET image, the treatment of any of the above cases (a, b1, or b2) results in an  $N \times N$  symmetric distance matrix D whose (i, j)th entry  $D_{ij} : \mathbf{R}^n \times \mathbf{R}^n \to \mathbf{R}^+$  stores the geodesic distance between the two values  $f(\mathbf{x}_i)$  at pixel i and  $f(\mathbf{x}_j)$  at pixel j, i.e.  $D_{ij} = d_{geodesic}(f(\mathbf{x}_i), f(\mathbf{x}_j))$ 

#### D. Distance Preserving Dimensionality Reduction into Color Spaces

high-dimensional points connected by the edge [36].

Given our goal of rendering pixels with similar high-dimensional data using similar colors (e.g. pixels capturing similar diffusion or metabolic processes, in DTMRI or dPET, respectively), and given the typical 3-channel representation of color spaces, the dimensionality of the data at each pixel must be reduced to 3D in such a way that the distances  $D_{ij}$  between all pairs of pixels is preserved (as much as possible) in such a dimensionality reduction. Clearly, linear (such as PCA or ICA [22], [6]), or even nonlinear, dimensionality reduction techniques that are not designed from the outset to preserve distance will be a poor choice towards achieving the aforementioned objective. The problem of performing a distance-preserving, nonlinear transformation of high-dimensional data points to lower dimension can be formulated as an optimization problem seeking transformation parameters and/or the new lower-dimensional representation of the data points such that the discrepancy between pairs of distances will be minimized. MDS is a well known approach that does exactly this [3]. We provide MDS with the N high-dimensional pixels, the geodesic distance matrix  $D_{ij}$  or an approximation thereof, and the target dimensionality: 3, to obtain N new pixels each of dimensionality 3. Each pixel can now be rendered in color, where the dissimilarity between colors (be it measured in RGB, HSV, or other color spaces) is equal (as much as possible) to the dissimilarity between the original pixel data. Our goal, however, is not only to render pixels with color, but rather to have equal differences between pixels be perceived as equally different. For this we resort to performing dimensionality reduction into a 3D *perceptual* color space.

## E. Distance-Preserving Dimensionality Reduction into a Perceptually-Uniform Color Space

To map high-dimensional pixels (such DTMRI or dPET pixels) into 3 dimensional color pixels, to preserve the pair-wise dissimilarity between pixels during such mapping, and to achieve a perceptual stimulus in observers (such as clinicians and radiologists), we must choose a perceptually uniform color space as the 3D range (target) of such a mapping. In a perceptually uniform color space, changes in color by a certain amount in that color space produce a change in the visual stimulus that is almost proportional to that amount. Hence, pairs of pixels with dissimilarities  $D_{ij}$  will be mapped to colors with perceptual difference proportional to  $D_{ij}$ . Therefore, pixels with similar high-dimensional data will be perceived (in color) similarly: our original objective. We choose the CIELAB perceptually uniform color space as the range of the mapping [43]. Given a DTMRI image, for example, with N pixels, with dissimilarity between pairs of pixels measured using an appropriate DT dissimilarity metric  $d_{geodesic}(T_i, T_j) = D_{ij}$ , where  $f(\mathbf{x}_i) = T_i$  is a diffusion tensor, the mapping from 6D (dimensionality of the extrinsic or ambient space of DTs) to CIELAB's 3D color space will result in corresponding pairs of pixels assigned colors separated by  $kD_{ij}$ , where k is a proportionality constant, i.e. perceived as similarly or as differently according to the value of  $D_{ij}$ . Clearly, the range of dissimilarity values between pairs of high-dimensional pixels in an image may be arbitrarily different than the range of possible Euclidean distances in the CIELAB color gamut. Therefore, a color normalization step must be performed.

Typically, CIELAB color difference thresholds are dependent on the desired application and thresholds for perceptibility judgments are significantly lower than thresholds for acceptability judgments. To correlate with human visual performance, differences in color are defined in terms of Euclidean distance in CIELAB (or L\*a\*b\*) units. A CIELAB residual, or  $\Delta E_{ab}^*$ , corresponds approximately to human judgements of perceptual difference, where CIELAB errors of 2 or 3 represent just noticeable color differences detectable by humans [35]. A difference of  $1 \Delta E_{ab}^*$  is sometimes the tolerance used for accepting or rejecting color tolerances in e.g. the dying of colored fabrics. Here we are interested in using this perception-based measure to delineate difference in medical data. Since there is a standard transform from CIELAB to color, we can indeed display colors according to their discriminability and perceptual distance.

## F. Perceptual Color Normalization

Given the difference between the range of  $D_{ij}$  values and the range of possible distances in the CIELAB gamut, a color normalization step is performed to *isotropically* scale the 3D points to new 3D points that, ideally, neither lie outside the CIELAB gamut nor leave parts of the gamut unutilized. The isotropy in the scaling is essential so as to preserve the relative distances between pairs of points. Isotropic scaling is not the only 3D-3D transformation that can be performed on 3D points that will preserve the relative pair-wise distances: translations and rotations in 3D can also be performed. Therefore, we formulate the normalization of the 3D points in the perceptual color space more generally as follows: We seek the 3D isotropic scaling, translation, and rotation transformation that best utilizes the CIELAB gamut. There can be several ways to formulate an objective function to capture this general criterion. The approach we adopt is to specify three key data points (e.g. three DTs) and specify which colors these three samples should be approximately transformed to. This, actually, is related to the Procrustes alignment or the absolute orientation problem for two sets of points, which can be solved analytically in closed form to find the rotation, translation, and isotropic scaling that, when applied to transform one set of points, will yield the smallest sum of squared distances between corresponding points [39].

#### G. Algorithm

To summarize the proposed method, algorithm 1 highlights the steps of our algorithms for rendering images with high-dimensional pixels such that pixels with similar physical characteristics (e.g. brownian

motion or diffusion or tracer dynamics) are perceived in color similarly.

Algorithm 1 Perceptual Visualization of High Dimensional Data

## Input:

- $\{f(\mathbf{x}_i)\}_{i=1}^N$ , where  $f(\mathbf{x}_i) : \mathbf{x}_i \in \mathbf{R}^d \to [f_1(\mathbf{x}_i), f_2(\mathbf{x}_i), \cdots, f_n(\mathbf{x}_i)] \in \mathbf{R}^n$ , i.e. N n-D pixels forming a d-dimensional image; d=2 for 2D images with  $\mathbf{x}_i = (x_i, y_i)$  or d=3 for 3D images with  $\mathbf{x}_i = (x_i, y_i, z_i)$ .
- $\mathbf{x}_{i} = (x_{i}, y_{i}, z_{i}).$   $\{(f(\mathbf{x}_{i}), C(\mathbf{x}_{i})\}_{i=1}^{P}, P \geq 3, \text{ i.e. at least 3 colors } C(\mathbf{x}_{i}) = [c_{1}(\mathbf{x}_{i}), c_{2}(\mathbf{x}_{i}), c_{3}(\mathbf{x}_{i})] \text{ (e.g. in RGB space } C(\mathbf{x}_{i}) = [R(\mathbf{x}_{i}), G(\mathbf{x}_{i}), B(\mathbf{x}_{i})]) \text{ associated with three different pixel values } f(\mathbf{x}_{i}), i = 1, \dots, P, P \geq 3. \text{ Without loss of generality, we assume these are the first } P \text{ elements in } \{f(\mathbf{x}_{i})\}_{i=1}^{N}.$
- Optionally, d<sub>geodesic</sub>(f(x<sub>i</sub>), f(x<sub>j</sub>)): R<sup>n</sup> × R<sup>n</sup> → R<sup>+</sup>; a helper function that calculates the geodesic distance (or a meaningful dissimilarity metric) between a pair of data points. For two DT pixels we use the LogEuclidean distance metric [1],

 $D_{ij} = d_{T_{LE}}(T_1, T_2) = \sqrt{trace((logm(T_1)) - (logm(T_2)))^2}$ , where  $T_1$  and  $T_2$  are two DTs and logm is the matrix logarithm, which is defined via the decomposition  $T = U\Lambda U^t$  as  $logm(T) = Udiag(\log(diag(\Lambda)))U^t$ .

# **Output:**

•  $\{C(\mathbf{x}_i)\}_{i=1}^N$ , where  $C(\mathbf{x}_i) : \mathbf{x}_i \in \mathbf{R}^d \to [c_1(\mathbf{x}_i), c_2(\mathbf{x}_i), c_3(\mathbf{x}_i)] \in \mathbf{R}^3$ , i.e. N 3-D pixels forming a d-dimensional color (e.g. RGB) image, such that perceptual color distance  $d_{perceptual}(C(\mathbf{x}_i), C(\mathbf{x}_j)) \propto d_{geodesic}(f(\mathbf{x}_i), f(\mathbf{x}_j)) \forall i, j \in \{1, 2, \dots, N\}$ , i.e. differences between pixel values are mapped to proportional differences in perception or visual stimulus.

## **Procedure:**

- Step 1. If  $D_{ij} = d_{geodesic}(f(\mathbf{x}_i), f(\mathbf{x}_j))$  is known (e.g.  $d_{T_{LE}}$  for DTs) go to Step 4.
- Step 2. Calculate d<sub>Euclidean</sub>(f(**x**<sub>i</sub>), f(**x**<sub>j</sub>)) = |f(**x**<sub>i</sub>) f(**x**<sub>j</sub>)|<sub>2</sub> ∀i, j ∈ 1, 2, ..., N s.t. f(**x**<sub>i</sub>) is connected to f(**x**<sub>j</sub>) in n-D. One of two connectedness criteria is applied: (i) |f(**x**<sub>i</sub>) f(**x**<sub>j</sub>)|<sub>2</sub> ≤ ε (ε-ISOMAP); (ii) f(**x**<sub>i</sub>) is connected to its K-closest (using |.|<sub>2</sub>) neighbors (K-ISOMAP) [36]. This generates a graph G(V, E) whose N vertices correspond to f(**x**<sub>i</sub>) and edges connect vertices that satisfy the connectedness criteria and are weighted by d<sub>Euclidean</sub>.
- Step 3. Approximate  $D_{ij} = d_{geodesic}(f(\mathbf{x}_i), f(\mathbf{x}_j))$  as the shortest path on the weighted graph (e.g. using Dijkstra's algorithm).
- Step 4. Calculate {g(**x**<sub>i</sub>)}<sup>N</sup><sub>i=1</sub>, where g(**x**<sub>i</sub>) : f(**x**<sub>i</sub>) ∈ **R**<sup>n</sup> → [g<sub>1</sub>(**x**<sub>i</sub>), g<sub>2</sub>(**x**<sub>i</sub>), g<sub>3</sub>(**x**<sub>i</sub>)] ∈ **R**<sup>3</sup>, such that d(g(**x**<sub>i</sub>), g(**x**<sub>j</sub>)) = d<sub>geodesic</sub>(f(**x**<sub>i</sub>), f(**x**<sub>j</sub>))∀i, j ∈ {1, 2, · · ·, N}, or make the difference as small as possible, i.e. perform D<sub>ij</sub> distance-preserving dimensionality reduction to 3D using MDS [3]. The {f(**x**<sub>i</sub>)}<sup>P</sup><sub>i=1</sub> samples (second input above) are now mapped to {g(**x**<sub>i</sub>)}<sup>P</sup><sub>i=1</sub> in 3D. All resulting g(**x**<sub>i</sub>) points are in perceptually uniform CIELAB 3D space, however they have arbitrary scale, rotation, and translation.

continued on the next page ...

## H. Polynomial Regression for Non-Linear Dimensionality Reduction

MDS (and ISOMAP which utilizes MDS) operates on a dissimilarity matrix D of dimensions equal to  $N \times N$  (where N is the number of pixels in the image). Given that N can be large when operating on 2D

or 3D images (e.g. a small  $100^3$  volume or a large  $1000^2$  2D image will result in a 1,000,000×1,000,000 dissimilarity matrix), it is important to address the issues of MDS complexity. The time complexity of MDS can be reduced to O(NlogN) [16], but at the price of increasing the space complexity to  $O(N^2)$  for the matrix of pre-computed distance values. Since this matrix is generally non-sparse, main memory size becomes a limiting issue. Therefore, we adopt a practical approach to performing MDS dimensional reduction as follows. Firstly, we run MDS on a random subsample of the high dimensional (e.g. diffusion tensor) data. Then, we calculate the mapping according to Algorithm 1 for this subsample only. Afterwards, we use the resulting mapping from high-dimensional *n*-D to 3-D and formulate a polynomial regression fit over this smaller data set. Finally, we apply the resulting regression from *n*-D to 3-D to all *n*-D pixel data points. In practice, we use a degree-2 polynomial (with no constant term). In Section IV, we demonstrate the effectiveness of this approach, by reducing the computation time while still achieving an accurate mapping on all pixels.

Step 5. Convert C(x<sub>i</sub>) (second input above) to CIELAB coordinates (e.g. RGB to CIELAB). Note that, generally, {g(x<sub>i</sub>)}<sup>P</sup><sub>i=1</sub> will not coincide with {C(x<sub>i</sub>)}<sup>P</sup><sub>i=1</sub> as desired.
Step 6. Transform {g(x<sub>i</sub>)}<sup>P</sup><sub>i=1</sub> using rotation R = USV<sup>t</sup>, isotropic scaling c = <sup>1</sup>/<sub>σ<sup>2</sup><sub>C</sub></sub>trace(D × S),

• Step 6. Transform  $\{g(\mathbf{x}_i)\}_{i=1}^{P}$  using rotation  $\mathbf{R} = USV^t$ , isotropic scaling  $c = \frac{1}{\sigma_C^2} trace(D \times S)$ , and translation  $\mathbf{t} = \mu_g - cR$ , which are calculated such that  $\{g(\mathbf{x}_i)\}_{i=1}^{P}$  are as close as close as possible to the corresponding  $\{C(\mathbf{x}_i)\}_{i=1}^{P}$ , where [39]:  $\mu_C = \frac{1}{P} \sum_{i=1}^{P} C(\mathbf{x})_i$ ,  $\mu_g = \frac{1}{P} \sum_{i=1}^{P} g(\mathbf{x}_i)$ ,  $\sigma_C^2 = \frac{1}{P} \sum_{i=1}^{P} ||C(\mathbf{x}_i) - \mu_C||^2$ ,  $\Sigma = \frac{1}{P} \sum_{i=1}^{P} (g(\mathbf{x}_i) - \mu_g) (C(\mathbf{x}_i) - \mu_C)^t$ ,  $\Sigma = UDV^t$ ,  $S = \begin{cases} I & , \ det(\Sigma) \ge 0 \\ diag(1, 1, ..., 1, -1) & , \ det(\Sigma) < 0 \end{cases}$ 

# **IV. RESULTS**

We begin our results by evaluating the accuracy of the polynomial regression approximation. Then, we test our method on hand-crafted synthetic images with 3 dimensional pixel data sampled from a variety of underlying distributions. We then apply our method to DTMRI brain and heart data with 6 dimensional pixels (positive semi-definite  $3 \times 3$  matrix), and to dynamic PET brain image data with a 27-dimensional time activity curve at each pixel.

#### A. Evaluation of Polynomial Regression

Figure 1 shows the disparity between the low-D representation obtained from running MDS on the full set of diffusion tensors versus the low-D representation obtained by running MDS on a random subsample of the tensors to learn a polynomial regression mapping from 6-D to 3-D, with interpolation generating 3D co-ordinates for the full set of tensors (Section III-H). We measure disparity as the  $L^2$  norm of the pairwise tensor distances in 3D, after running MDS, minus the pairwise tensor distances in 3D after interpolation. We used 1173 tensors from a mid-sagittal image of the corpus callosum (CC) (see Figure 7), and measure the disparity between all possible pairs of tensors. We see from Figure 1(a) that utilizing about 25% of the data provides a reasonable balance between accuracy and data size, in that the largest gain in accuracy is achieved going from sampling at 15% of the data to just above 20%. Figure 1(b) demonstrates that the regression approach does adequately well in representing a full MDS analysis.



Fig. 1. Evaluation of polynomial regression. (a) Error obtained by approximating MDS on the whole data set by polynomial regression on a subset of the data. (b) Scatter plot of distances approximated via polynomial regression vs. the original tensor distances. The 1:1 line is shown along the diagonal. A sample size of 25% of the total number of pixels is employed in the approximation.

## B. Synthetic data with pixel dimensionality 3

We begin with simple synthetic image data. We created a  $60 \times 30$ -pixel 2D image divided into two regions: upper and lower halves, U and L, each of size  $30 \times 30$ . Each pixel is a 3-dimensional vector. Four types of images were synthesized to produce different distributions of the 3D pixels (Figure 2): (i) U and L 3D pixel vectors are sampled from two different 3D Gaussian distributions with different means and anisotropic covariance matrices, with the axes of maximal variation extending parallel to the first extrinsic dimension (Figure 2(a)); (ii) Same as (i) but with an oblique axis of maximal variation (Figure 2(b)); (iii) the U and L pixels are sampled from two Gaussian distributions that together form a nonlinear space of samples in 3D (Figure 2(c)); and (iv) the 3D vectors of the image pixels form a "Swiss roll" in 3D, where one half of the swiss roll corresponds to L pixels and the other half to U pixels (Figure 2(d)). We show in Figure 3 the results of different approaches to coloring this synthetic data. There are two issues to be examined here: (i) How the images change as the method of mapping from data space to color space is changed; (ii) How the images change as the color space being mapped to is changed. As expected, when the data is linearly separable PCA methods are able to display the main variability of the data. We find that CIELAB colors do not provide any convincing advantage in these cases. However, for the more complicated synthetic example of the swiss roll data, only manifold learning provides an acceptable degree of visual separation between the two classes. Moreover, in our method difference is represented as color keyed to perceptual difference so the bottom-right image, displaying the highest visual separation of complexly interwoven data, provides convincing evidence justifying the suitability of the proposed approach.

#### C. Simulated DTMRI

Medical doctors typically resort to viewing scalar images derived from the DTMRI field. One common scalar field is a Fractional Anisotropy (FA) image. At each pixel in the DT image, the FA is proportional to the amount of anisotropy in DT at that pixel. FA is a function of the eigenvalues  $\lambda_i$ , i = 1..3, of the DT and is defined as  $FA = \sqrt{3/2} \sqrt{\sum_{i=1}^{3} (\lambda_i - \overline{\lambda})^2} / \sqrt{\sum_{i=1}^{3} \lambda_i^2}$ , where  $\overline{\lambda}$  is the Mean Diffusivity (MD),



Fig. 2. Different types of synthetic data with pixel dimensionality 3 used to populate the upper (U) and lower (L) halves of the image. Blue x's correspond to U pixels and green squares to L pixels. (a) Data sampled from two Gaussians whose principal directions of variability align with the first extrinsic dimension. (b) Data sampled from two Gaussians whose principal direction of variability does not align with any of the extrinsic dimensions. (c) Data sampled from two Gaussians forming a nonlinear subspace. (d) Data sampled from a nonlinear Swiss roll.

 $1/3 \sum_{i=1}^{3} \lambda_i$ . MD is another scalar field derived from DTMRI data that is typically used to explore DT data in clinical practice and constitutes an average measure of diffusion at a particular DT pixel. Clearly, scalar images are not able to capture the variability of tensors in 6D. The simulated DT image in Figure 4 is designed to highlight this situation. In Figure 4(right), our method clearly shows perceptual changes in color along different directions in the image, and indeed the simulated tensors do change in those directions. However, the same DT field visualized via the commonly used FA and MD maps in Figure 4(left and middle) is insensitive to the change in DTs that occurs as we move vertically in the image (constant value in any column in the FA or MD maps).

A second synthetic DTMRI data set is presented in Figure 5. Here, our coloring method clearly shows a gradual transition in the DTs as we move vertically in the image, with the top half different from the lower half. The FA map shows changes in one half of the image but fails to discriminate between tensors in the top half vs. the bottom half. The MD, on the other hand, is completely oblivious to the change in tensors.



(a) Gaussians, variability aligns with extrinsic



(b) Gaussians, variability does not align with extrinsic





(d) Swiss roll

Fig. 3. Coloring the 3D pixels using different approaches. The four rows from top to bottom correspond to the four cases in Figure 2. The four columns correspond to coloring using (from left to right): (i) Extrinsic dimensions as RGB; (ii) Principal components as RGB; (iii) Principal components as CIELAB; and (iv) Intrinsic dimensions as CIELAB (i.e. our approach).



Fig. 4. First synthetic DTMRI example (see text in section IV-C for details). The FA and MD images capture changes along the horizontal direction but fail to capture changes along the vertical direction. Our approach captures changes in DT along both directions.



Fig. 5. Second synthetic DTMRI example (see text in section IV-C for details). Although the FA image captures changes in DT along the vertical direction, it fails to distinguish between the top and bottom halves of the image. MD fails to capture any change. Our method is able to distinguish between the top and bottom halves.

## D. Real Brain and Heart DTMRI

Here we present results of experiments on real DTMRI data. 2D slices from 3D DTMRI volumes are used. In the first example, we focus on the CC in the brain (Figure 6). In Figure 7, we show a close up of the visualization of of the CC region in a mid-sagittal 2D slice. Note that using our coloring method, the CC body appears whitish and the fornix appears with a brownish hue indicating a difference in the underlying diffusion tensors. In contrast, it is difficult to see any difference between the colors of the fornix and the CC body in the structural MRI (Figure 6). Similarly, in the MD and FA images, the coloring of two regions is almost indistinguishable. Figure 8 shows additional examples showing how our coloring approach reveals differences and renders them in a perceptually accurate way compared to traditional MD and FA visualization. In Figure 9, we compare MD, FA and our method of coloring high dimensional images applied to 2D short axis cardiac DTMRI slices. Note how the colors change reflecting the myocardial sheet organization (the transmural rotation from the endocardial to the epicardial surface

# [21], [31]).

No false-coloring applied to either MD or FA will rectify this disadvantage since equal grayscales will simply map to equal colors, thus failing to disambiguate regions which our method clearly shows.



Fig. 6. Corpus callosum (CC) anatomy referred to in elsewhere in the paper.



Fig. 7. Close up on the CC. Not how the fornix appears with the same color as the CC body when either MD or FA are used. Using our method, the fornix appears with a different (brown-ish) color, where as the CC body appears whitish.

# E. Real DTMRI image with simulated pathology

Several clinical works have demonstrated that different pathologies, such as tumor progression and growth, multiple sclerosis lesions, and high grade gliomas are manifested as changes in diffusion tensor properties [30], [34], [8], [20], [32]. We performed the next experiment to mimic the existence of pathology in brain DTMRI. The DTs in a particular region of interest (ROI) in the CC between the rostrum and genu (Figure 6) were manipulated to simulate a pathological condition (Figure 10). More specifically, the tensors in that ROI are rotated and then the eigenvalues of the DT are modified such that both MD and FA remain the same. Figure 10(a) shows the CC rendered with our method compared to displays using FA and MD maps, before the simulation of pathology. In Figure 10(b), the simulated pathology is clearly shown in the rendering with our method whereas, as expected, the FA and MD maps remain completely unchanged.



Fig. 8. Other examples of brain DTMRI visualization (sagittal and axial views). Each row represents one image. The columns represent different methods for coloring, from left to right: MD, FA, and our approach.



Fig. 9. Visualization of cardiac DTMRI slices. Short axis slices from two data sets are shown (the two rows). From left to right, we show MD, FA and our approach to exploring the DTMRI data.



(b) After

Fig. 10. Simulating pathology within a real DTMRI image. (a) before pathology simulation, (b) after. Left to right: MD, FA, and our method. The pathology is clearly visible (in red) using our method.

## F. Dynamic PET

In dPET, a time activity curve is collected at each pixel. The time activity is typically sampled nonuniformly. In the first frames, the sampling interval is shorter to account for more rapidly changing tracer dynamics. This results in a lower photon count and, therefore, a lower signal to noise ratio (SNR). In later frames, the tracer dynamics stabilize and longer sampling intervals are used, so that more photon counts are collected yielding higher SNR [13]. Therefore, clinicians often examine the last PET frame of a dynamic study. This can be misleading because the activity in the last frame can be the same for different tissues with very different dynamic behavior (i.e. they just happen to have similar value for activation in the last frame). The same is argued for the integral under the time activity curve; clinicians sometimes look at a static scalar field with these integral values at each pixel. However, vastly different curves can have similar integrals, and hence displaying these scalar fields can also be misleading. To better illustrate the point, in 11(d), we examine an axial brain dPET slice with a time activity curve of dimensionality 27 at each pixel. We quantify the number of pairs of neighboring pixels that are erroneously visualized with a similar color when the scalar TAC integral is used for visualization (i.e. the change in the TAC integral is less than a threshold  $T_{scalar}$ ), whereas visualizing them using our method shows perceptual difference between pixels (larger than  $T_{perceptual}$ ) as a result of the difference in their underlying 27-D TACs. To further illustrate how our method can improve over the standard visualization techniques for dPET, we provide additional examples.

Figure 12 shows an ROI in another dPET slice (Image 1), comparing the results of four visualization methods: the last time frame; a sum-over-frames; a mapping of Principal Component (PC) weights to RGB; and finally, a mapping of PC weights to CIELAB. As before, visualizing neither the last frame, nor a sum-over-frames conveys the full variability of the data. Using PCA allows us to map the 27-D data to a 3-D color spaces (RGB and CIELAB), but as the data is non-linear in origin, PCA is not ideally suited for this task. Though both color images technically represent the same degree of variability, as they are both 3D color spaces, only in the CIELAB image does the visual difference in colors correspond to the underlying differences. In Figure 13, manifold learning is used instead of a linear mapping (PCA). For comparison, we examined learning both 2D and 3D manifolds, with 3D manifolds providing the better results. As expected, utilizing a 2D-manifold produces less information than does using a 3D-manifold, and more details are discernible. Figures 14 and 15 show similar results on another dPET slice (Image 2). Note how our approach not only highlights the putamen in the brain but also shows different tracer uptake properties within the putamen itself (reddish in the medial part of the putamen changing laterally to white in Figure 15(d)).

## V. CONCLUSIONS

CIELAB is a 3-dimensional color metric based on Weber's law, that provides an approximately uniform color-distance measure. Here we perform dimensionality reduction with a view to having high-dimensional distances properly mapped to color, while at the same time utilizing the CIELAB space to best advantage. The key idea of our method is to visualize similar data points using perceptually similar colors. The dissimilarity between data points is measured as the geodesic distance on the underlying manifold, with the manifold either known analytically (as in DTMRI) or learned (as in dPET data). To map the high dimensional data into color, a nonlinear mapping is calculated in such a way that distance between data points is preserved as much as possible when transforming the data into a perceptually uniform 3D color space. We tested our method on synthetic data and simulated and real medical image (DTMRI and dPET) data.

We foresee no obstacles in using our method for other (non-medical) high dimensional data field visualization, e.g. geospatial data. We also anticipate added value if, in addition to 3 dimensional color,



(c) TACs

(d) pixel count

Fig. 11. Misleading dPET visualization. (a) ROI of a dPET image with the sum of time activity visualized at each pixel. (b) The same ROI visualized using our method. (c) The distribution of TACs in this ROI (each TAC contains n=27 time samples). (d) The number of neighboring pixel pairs (p, q) of which sub-figure (a) generates misleading visualization. The number of pixels are those that satisfy the following criteria: The difference in the summation of activity (values in (a)) is smaller than threshold  $T_{scalar}$  and at the same time the CIELAB distance between p and q is larger than  $T_{perceptual}$ .



(a) Last time frame of a 27-frame dPET image.



(b) Last frame

(c) Sum-over-frames

(d) PC as RGB

Fig. 12. Real dPET brain image. Linear results of dPET Image 1. The last frame of a dPET image (top row), and a zoomed in region of the image (bottom row) displayed using four alternative methods.



Fig. 13. Intrinsic results of PET Image 1, where the n-D manifold is learned via ISOMAP, learning either a 2D- or 3D-manifold. The initial distance metric given to ISOMAP is simply the Euclidean distance between PET samples treated as vectors. (a) Intrinsic learned 2D-manifold as RGB. (b) Intrinsic 2D-manifold as CIELAB. (c) Intrinsic learned 3D-manifold as RGB. (d) Intrinsic 3Dmanifold as CIELAB.



(a) Last time frame of a 27-frame dPET image.



(b) Last frame

(c) Sum-over-frames

(d) PC as RGB

Fig. 14. Linear results of dPET Image 2. The last frame of a dPET image (top row), and a zoomed in region of the image (bottom row) displayed using four alternative methods.



Fig. 15. Intrinsic results of PET Image 2, where the n-D manifold is learned via ISOMAP, learning either a 2D- or 3D-manifold. The initial distance metric given to ISOMAP is simply the Euclidean distance between PET samples treated as vectors. (a) Intrinsic learned 2D-manifold as RGB. (b) Intrinsic 2D-manifold as CIELAB. (c) Intrinsic learned 3D-manifold as RGB. (d) Intrinsic 3Dmanifold as CIELAB.

opacity is also used in the visualization. More powerful visualization may also be obtained if the colors extracted at each data using our method are used to color different types of glyphs.

One of the important next steps of our work is to collaborate closely with doctors or radiologists to assess the objective clinical value of our approach.

#### VI. ACKNOWLEDGEMENTS

GH and MD are funded through the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery grants. CM is funded from by NSERC Graduate (Doctoral) Scholarship and Michael Smith Foundation for Health Research (MSFHR) Senior Graduate Studentship. We acknowledge the PET-SORTEO project [28], the source of the dPET data. We thank Drs. Patrick A. Helm and Raimond L. Winslow at the Center for Cardiovascular Bioinformatics and Modeling and Dr. Elliot McVeigh at the National Institute of Health for provision of the heart DTMRI data. We thank the Johns Hopkins University for providing the DTMRI data. We used the manifold learning software for MATLAB provided by Todd Wittman, Department of Mathematics, University of Minnesota<sup>1</sup>. We thank Eli Gibson for assistance in preparing the synthetic tumor example.

#### REFERENCES

- [1] Vincent Arsigny, Pierre Fillard, Xavier Pennec, and Nicholas Ayache. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. Magnetic Resonance in Medicine, 56(2):411-421, August 2006.
- [2] P J Basser, J Mattiello, and D LeBihan. MR diffusion tensor spectroscopy and imaging. Biophys. J., 66(1):259-267, 1994.
- [3] Ingwer Borg and Patrick J. F. Groenen. Modern Multidimensional Scaling: Theory and Applications (Springer Series in Statistics). Springer, Berlin, 2nd ed. edition, September 2005.
- [4] Anders Brun, Hae-Jeong Park, Hans Knutsson, and Carl-Fredrik Westin. Coloring of DT-MRI fiber traces using Laplacian eigenmaps. In Roberto Moreno Diaz and Alexis Quesada Arencibia, editors, Computer Aided Systems Theory (EUROCAST'03), Lecture Notes in Computer Science 2809, pages 564-572, Las Palmas de Gran Canaria, Spain, February 24-28 2003. Springer Verlag.
- [5] Richard E Carson. Tracer kinetic modeling in pet. Positron Emission Tomography, pages 127–159, 2005.
- [6] Pierre Comon. Independent component analysis, a new concept? Signal Process., 36(3):287-314, 1994.
- [7] Mark S. Drew and Ghassan Hamarneh. Visualizing diffusion tensor dissimilarity using an ICA based perceptual color metric. In Color Imaging, pages 42-47, 2007.
- [8] I. Elshafiey. Diffusion tensor magnetic resonance imaging of lesions in multiple sclerosis patients. Radio Science Conference, 2002. (NRSC 2002). Proceedings of the Nineteenth National, pages 626-633, 2002.
- [9] Robert F. Erbacher, David L. Gonthier, and Haim Levkowitz. Color icon: a new design and a parallel implementation. volume 2410, pages 302-312. SPIE, 1995.
- [10] Zhe Fang, Torsten Moeller, Ghassan Hamarneh, and Anna Celler. Visualization and exploration of time-varying medical image data sets. In Graphics Interface (GI), pages 281-288, 2007.
- [11] M.R. Gupta and N.P. Jacobson. Wavelet principal component analysis and its application to hyperspectral images. Image Processing, 2006 IEEE International Conference on, pages 1585–1588, Oct. 2006.
- [12] C.G. Healey and J.T. Enns. Large datasets at a glance: combining textures and colors in scientific visualization. Visualization and Computer Graphics, IEEE Transactions on, 5(2):145-167, Apr-Jun 1999.
- [13] Kewei Chen Hongbin Guo, Rosemary Renaut and Eric Reiman. Clustering huge data sets for parametric pet imaging. Biosystems, 71(1-2):81-92, 2003.
- [14] N.P. Jacobson and M.R. Gupta. Design goals and solutions for display of hyperspectral images. Geoscience and Remote Sensing, IEEE Transactions on, 43(11):2684-2692, Nov. 2005.
- [15] N.P. Jacobson, M.R. Gupta, and J.B. Cole. Linear fusion of image sets for display. Geoscience and Remote Sensing, IEEE Transactions on, 45(10):3277-3288, Oct. 2007.
- [16] F. Jourdan and G. Melangon. Multiscale hybrid mds. Information Visualisation, 2004. IV 2004. Proceedings. Eighth International Conference on, pages 388-393, July 2004.
- [17] S. Kaewpijit, J. Le Moigne, and T. El-Ghazawi. Automatic reduction of hyperspectral imagery using wavelet spectral analysis. Geoscience and Remote Sensing, IEEE Transactions on, 41(4):863-871, April 2003.
- [18] D.A. Keim. Designing pixel-oriented visualization techniques: theory and applications. Visualization and Computer Graphics, IEEE Transactions on, 6(1):59-78, Jan-Mar 2000.

20

<sup>1</sup>http://www.math.umn.edu/~wittman/mani/

- [19] Gordon Kindlmann, Raul San Jose Estepar, Marc Niethammer, Steven Haker, and Carl-Fredrik Westin. Geodesic-loxodromes for diffusion tensor interpolation and difference measurement. In *Tenth International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'07)*, Lecture Notes in Computer Science 4791, pages 1–9, Brisbane, Australia, October 2007.
- [20] H. Kitzler, W. Benger, A. Werner, H. Bartsch, A. Shumilina, H.-C. Hege, G. Schackert, and R. von Kummer. Diffusion tensor imaging: visualization of brain tumours using the method of tensor patterns. 57th Annual Meeting of the German Society of Neurosurgery Joint Meeting with the Japanese Neurosurgical Society, 2006.
- [21] I. J. Legrice, P. J. Hunter, and B. H. Smaill. Laminar structure of the heart: a mathematical model. Am J Physiol Heart Circ Physiol, 272(5):H2466–2476, 1997.
- [22] K. Pearson. On lines and planes of closest fit to systems of points in space. Philosophical Magazine, 2(6):559-572, 1901.
- [23] Xavier Pennec, Pierre Fillard, and Nicholas Ayache. A riemannian framework for tensor computing. Int. J. Comput. Vision, 66(1):41–66, 2006.
- [24] D. L. Pham, C. Xu, and J. L. Prince. A survey of current methods in medical image segmentation. In Annual Review of Biomedical Engineering, volume 2, pages 315–338. 2000.
- [25] K. Rasche, R. Geist, and J. Westall. Detail preserving reproduction of color images for monochromats and dichromats. *Computer Graphics and Applications, IEEE*, 25(3):22–30, May-June 2005.
- [26] Karl Rasche, Robert Geist, and James Westall. Re-coloring images for gamuts of lower dimension. Computer Graphics Forum, 24:423–432(10), September 2005.
- [27] P. Ready and P. Wintz. Information extraction, snr improvement, and data compression in multispectral imagery. Communications, IEEE Transactions on [legacy, pre - 1988], 21(10):1123–1131, Oct 1973.
- [28] A. Reilhac, G. Batan, C. Michel, N. Costes, and A.C. Evans. Pet-sorteo: a platform for simulating realistic pet studies. In Nuclear Science Symposium Conference Record, 2004 IEEE, volume 7, pages 4053–4057, Oct. 2004.
- [29] Sam T. Roweis and Lawrence K. Saul. Nonlinear Dimensionality Reduction by Locally Linear Embedding. Science, 290(5500):2323–2326, 2000.
- [30] M Schluter, B Stieltjes, H K Hahn, J Rexilius, O Konrad-verse, and H o Peitgen. Detection of tumour infiltration in axonal fibre bundles using diffusion tensor imaging. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 1(3):80–86, 2006.
- [31] D. F. Scollan, Alex Holmes, Raimond Winslow, and John Forder. Histological validation of myocardial microstructure obtained from diffusion tensor magnetic resonance imaging. Am J Physiol Heart Circ Physiol, 275(6):H2308–2318, 1998.
- [32] Price S.J., Burnet N.G., Donovan T., Green H.A.L., Pena A., Antoun N.M., Pickard J.D., Carpenter T.A., and Gillard J.H. Diffusion tensor imaging of brain tumours at 3t: A potential tool for assessing white matter tract invasion? *Clinical Radiology*, 58:455–462(8), June 2003.
- [33] E Stejskal and J Tanner. Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient. Journal of Chemical Physics, 42:288–292, January 1965.
- [34] Bram Stieltjes, Mathias Schlter, Bernd Didinger, Marc-Andr Weber, Horst K. Hahn, Peter Parzer, Jan Rexilius, Olaf Konrad-Verse, Heinz-Otto Peitgen, and Marco Essig. Diffusion tensor imaging in primary brain tumors: Reproducible quantitative analysis of corpus callosum infiltration and contralateral involvement using a probabilistic mixture model. *NeuroImage*, 31(2):531 – 542, 2006.
- [35] M. Stokes, M.D. Fairchild, and R.S. Berns. Precision requirements for digital color reproduction. ACM Transactions on Graphics, 11(4):406–422, October 1992.
- [36] Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. Science, 290(5500):2319–2323, 2000.
- [37] V. Tsagaris and V. Anastassopoulos. Multispectral image fusion for improved rgb representation based on perceptual attributes. *International Journal of Remote Sensing*, 26:3241–3254(14), 10Aug2005.
- [38] J.S. Tyo, A. Konsolakis, D.I. Diersen, and R.C. Olsen. Principal-components-based display strategy for spectral imagery. Geoscience and Remote Sensing, IEEE Transactions on, 41(3):708–718, March 2003.
- [39] S. Umeyama. Least-squares estimation of transformation parameters between two point patterns. Pattern Analysis and Machine Intelligence, IEEE Transactions on, 13(4):376–380, Apr 1991.
- [40] Jing Wang and Chein-I Chang. Independent component analysis-based dimensionality reduction with applications in hyperspectral image analysis. *Geoscience and Remote Sensing, IEEE Transactions on,* 44(6):1586–1600, June 2006.
- [41] Zhizhou Wang and B.C. Vemuri. An affine invariant tensor dissimilarity measure and its applications to tensor-valued image segmentation. Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on, 1:I–228–I–233 Vol.1, June-2 July 2004.
- [42] Pak Chung Wong and R. Daniel Bergeron. 30 years of multidimensional multivariate visualization. In Scientific Visualization, Overviews, Methodologies, and Techniques, pages 3–33, Washington, DC, USA, 1997. IEEE Computer Society.
- [43] Günther Wyszecki and W. S. Stiles. Color Science: Concepts and Methods, Quantitative Data and Formulae (Wiley Series in Pure and Applied Optics). Wiley-Interscience, 2 edition, August 2000.
- [44] Erdem Yorukk, Burak Acar, and Roland Bammer. A physical model for dt-mri based connectivity map computation. Medical Image Computing and Computer-Assisted Intervention (MICCAI), pages 213–220, 2005.