

Perception-based Visualization of Manifold-Valued Medical Images using Distance-Preserving Dimensionality Reduction

Ghassan Hamarneh, *Senior Member, IEEE*, Chris McIntosh, and Mark S. Drew, *Member, IEEE*

Abstract—A method for visualizing manifold-valued medical image data is proposed. The method operates on images in which each pixel is assumed to be sampled from an underlying manifold. For example, each pixel may contain a high dimensional vector, such as the time activity curve (TAC) in a dynamic positron emission tomography (dPET) or a dynamic single photon emission computed tomography (dSPECT) image, or the positive semi-definite tensor in a diffusion tensor magnetic resonance image (DTMRI). A nonlinear mapping reduces the dimensionality of the pixel data to achieve two goals: distance preservation and embedding into a perceptual color space. We use multi-dimensional 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 either determined analytically as geodesics on the manifold of pixels or is 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, 2^{nd} order positive semi-definite DTs, whereas the dissimilarity between TACs is approximated via ISOMAP. We demonstrate our approach via artificial high-dimensional, manifold-valued data, as well as case studies of normal and pathological clinical brain and heart DTMRI, dPET, and dSPECT images. Our results demonstrate the effectiveness of our approach in capturing, in a perceptually meaningful way, important features in the data.

Keywords: Manifold-valued data; high dimensional data; visualization; color; nonlinear dimensionality reduction; multi-dimensional scaling; distance-preserving mapping; diffusion tensor magnetic resonance imaging (DTMRI); dynamic positron emission tomography (dPET); dynamic single photon emission computed tomography (dSPECT).

I. INTRODUCTION

HIGH dimensional, manifold-valued data is becoming more prevalent in medical imaging applications [1], [2], [3]. This increase of data complexity is due to several reasons, including advances in acquisition hardware that allows for collecting more anatomical, functional, and hybrid data; fusion of medical images from multiple sources to provide complementary information; faster acquisitions that enable reconstructing tracer dynamics in tissues, resulting in time activity curves (TAC) captured at each pixel, as in dynamic single-photon emission computed tomography (dSPECT) and

positron emission tomography (dPET); and new imaging modalities 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).

For patient-specific diagnosis and therapy, as well as for population studies, there is a need for accurate, repeatable, and fast algorithms to analyze large numbers of complex data with increasing dimensionality and resolution. Despite commendable efforts towards automating the medical image interpretation task, human intervention remains unavoidable and intuitive image display methods persist in the clinical workplace. In contrast to the ubiquitous display of 2D scalar (grayscale) or true-color RGB images, developing methods for the visualization of higher dimensional, manifold-valued pixels and complex spatial fields remains an area of active research [4]. The objective of this work is to develop a visualization method for manifold-valued medical images that is intuitive yet faithful to the underlying variability in the high-dimensional image pixel data.

Medical image segmentation, a key image analysis task that partitions an image into different regions, is often a necessary precursor to the higher level understanding of 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 [5], and although a universal criteria for a correct segmentation is likely impossible, almost all algorithms attempt to incorporate (i-iii) in different forms. However, the majority 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 may not be relied upon in clinical settings without user intervention. The fact remains that expert human operators in the clinic (e.g. radiologists)

The authors GH and CM are with the Medical Image Analysis Lab and MSD is with the Vision and Media Lab, School of Computing Science, Simon Fraser University, Canada. Corresponding author: hamarneh@cs.sfu.ca

continue to interpret medical images by depending largely on visual assessment, in which image partitioning and pixel grouping still rely on pixel dis/similarities and prior or expert domain knowledge, i.e. i-iii above. It is therefore imperative to provide experts with a medical image display or visualization system that is faithful to the underlying structural or functional data.

We argue that 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 tumour) or artificially introduce information (false positives) thereby misleading their analysis. In this work we address these issues by proposing a visualization approach for high-dimensional, manifold-valued medical image data 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). We display high-dimensional medical image data to the domain expert as color images. The transformation of the high-dimensional data to a form suitable for 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. Our display method is geared toward research clinicians and medical image analysis researchers faced with the challenging task of distilling high dimensional data to clinically relevant information.

II. RELATED WORK

In [6], Wong and Bergeron surveyed multidimensional multivariate visualization techniques prior to 1997. In [7], 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. More recently, Peng and Laramee surveyed a variety of approaches for visualization of high-dimensional data [4]. We focus our review on approaches related to using color to visualize medical or general high-dimensional data.

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 [8], hyperspectral image data¹ was rendered through projection on bases that reflect, through color-matching functions [9], [10], what the human visual system would perceive had it been able to cover the range of the hyperspectral data. Jacobson et al. proposed a method to design fixed basis functions mimicking human vision [11]. As early as 1973 [12], 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)

¹Whereas color images typically capture red, green and blue color channels per pixel (i.e. three-channel multispectral image), hyperspectral images capture many more channels that may include electromagnetic radiation outside the range of colors visible by humans (e.g. infrared or ultraviolet).

[13]. Other methods relied on PCA in conjunction with wavelet methods, for example, by performing PCA on wavelet sub-bands to enhance edges at specific levels of detail [14] or to perform wavelet-based denoising followed by PCA [15]. 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 [16], [17]. Other approaches for high-dimensional visualization relied on providing iconic representation for each data point [18] or multi-colored texture elements [19]. In [20], the mapping, from high-dimensional pixel data to 3D color, which maximizes the mutual information between the original hyperspectral bands and the color channels was used. In [21], Fang et al. used multi-dimensional scaling (MDS) as an alternative, nonlinear approach to dimensionality reduction [22], where 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 approach are those of Rasche et al. and Brun et al. In [23], [24], color images (the higher-dimensional data in their case), are reduced for color deficient, mono- and di-chromats displays. They highlighted the importance of preserving contrast and maintaining luminance consistency by basing 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 [24]. In [25], 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 eigen-maps. 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 fiber tracts, but rather to generic high-dimensional image fields, with a focus in this paper on medical image data (DTMRI, dPET, and dSPECT 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 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 DTs: rank 3, second order positive semi-definite (PSD) with nonnegative eigenvalues. When such knowledge about the manifold structure is absent, we resort to manifold learning techniques to capture dissimilarity between high-dimensional pixels.

Finally, we note that there exist approaches that produce pseudo-color images, like our method does. However, these approaches differ in two main aspects: the input to the algorithm, and the criteria governing the creation of the output image. Two examples of such methods stand out. Neophytou and Mueller proposed a tool for manipulating regions of color images by artists and designers who operate directly on the color gamut of an image region by applying geometric

transformations in the CIELAB perceptually uniform color space [26]. Similar to our approach, Neophytou and Mueller generate a pseudo-color image based on processing performed in the CIELAB space. However, the input to their method is a color image whereas the input to our method is an image with high-dimensional (manifold-valued) pixels. An input RGB color image is a special case in our technique, as we accommodate (and focus on) higher dimensional pixels, e.g. the 6 dimensions of DTMRI and the tens dimensions of time activity curves in dPET, and others. Further, key to our approach is the distance preserving dimensionality reduction (from high-dimensional pixels to 3D color space) and the use of geodesic distances on the manifold of pixel data (either analytically or using manifold learning). In [27], Wang et al. propose a different semi-interactive colorization technique, in which the user chooses, based on personal preferences, the hue of features in the image. Then, saturation and luminance parameters are calculated based on aesthetic rules, such as color harmony, user preferences, or scene parameters, e.g. feature size. The calculations are done in CIELAB perceptual color space. Although Wang et al. generate a pseudo color image like we do, their method expects as input a segmented image (i.e. pixels assigned to classes or labels), whereas our method works with the original data and is geared specifically towards high dimensional (manifold-valued) pixels and not an image of scalar labels. Another important difference is that, in [27], pixels with the same image values might receive different colors if they belong to different classes. In our approach this is unacceptable; pixels with the same high-dimensional data values receive the same color. This is done to ensure that the pseudo-color image faithfully reflects the underlying data, and to avoid data misinterpretation. Furthermore, Wang et al.'s method resorts to aesthetics and other cues to control the resulting colors. Our strategy is different; our critical rule is to be faithful to the underlying data as much as possible, rendering similar high-dimensional pixels with perceptually similar colors, regardless of the aesthetic appeal that is of negligible importance in medical applications.

III. METHODS

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 [28] or locally linear embedding [29]), derived analytically, or approximated. 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 $f(\mathbf{x}) : \mathbf{x} \in \mathbf{R}^d \rightarrow \mathbf{R}^n$, 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.

As we focus on DTMRI, dPET and dSPECT images in this paper, we provide a brief review of the basics of these image modalities. DTMRI is a 3D medical imaging modality that captures the probability of water diffusing in different spatial directions. The diffusion properties are captured by assigning at every voxel a DT, which is estimated from multiple diffusion measurements along different directions, dictated by diffusion-sensitizing gradients that are activated during the MR imaging protocol. The resulting DTMRI image of tensors is useful for understanding the underlying microstructure of fibrous tissues, such as the white matter of the brain. DTs are symmetric 3×3 matrices, or second-order rank 3 diffusion tensors, with 6 unique elements, i.e. $n = 6$ and $f(\mathbf{x}) : \mathbf{x} \in \mathbf{R}^3 \rightarrow \mathbf{R}^6$. Further, DTs must be PSD as 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 [30], [31]. This positive semi-definiteness, in turn, results in DTs being restricted to a convex half cone in 6D [32]. For more information on DTMRI we refer the reader to [33], [34].

In dSPECT medical imaging, radiopharmaceuticals carrying gamma-emitting isotopes are injected into the body. A rotating gamma camera collects the gamma photons from different orientations and computed tomography image reconstruction methods use this data to reconstruct a dynamic 3D image.

In dPET positron-carrying radiopharmaceutical tracers are injected in the body. The electron-positron annihilation causes anti-parallel gamma photons to be emitted, which are captured by the surrounding cylindrical gamma coincidence detector. The data is collected in ‘list mode’ and reconstructed into time activity curves assigned to different spatial positions. Both dPET and dSPECT are functional medical imaging modalities that capture the spatio-temporal patterns of the injected tracers, which is important for revealing underlying physiology, e.g. metabolic activity. dSPECT or dPET are represented as $f(\mathbf{x}) : \mathbf{x} \in \mathbf{R}^3 \rightarrow \mathbf{R}^n$ where, n is the number of the TAC samples at each voxel. The reader may consult with [35], [36], [37] for further details on dSPECT and dPET.

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 [38]. 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 [39], [40], [41], [42]. 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 [21]. (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 space 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 high-dimensional sample points and whose edge weights are equal to the Euclidean distance between the two high-dimensional points connected by the edge [28].

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 $D_{ij} : \mathbf{R}^n \times \mathbf{R}^n \rightarrow \mathbf{R}^+$ 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

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 maximally preserved. Clearly, linear (such as PCA [43] or ICA [44]), 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 [22]. 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 3D 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 [9], [10]. CIELAB is based on a one-third power law that closely

resembles a logarithmic encoding of human-visual-system tri-stimulus values. Consequently, CIELAB forms an expression in terms of color vision of the well-known Weber’s law from psychophysics, which attempts to describe the relationship between the physical magnitudes of stimuli and the perceived intensity of the stimuli, in order to formulate an approximately perceptually uniform metric.

HSV is one of many color spaces that assign an axis to brightness and two others axes effectively to hue and saturation, which may be intuitively seen as a geometrically-derived color space transformation that tilts the RGB cube such that the “black-white” intensity axis aligns with the “Value” axis in HSV. According to the encyclopedia article Color Perception [45], HSV does correspond to the perceptual correlates of hue, chroma, and lightness (human vision descriptors) but are not meant to be understood as even approximately forming a uniform measure. Instead, CIELAB color space is meant to be a much better approximation of an isotropic perceptually uniform metric color space.

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; thresholds for perceptibility judgments (just-noticeable difference in a laboratory setting) are significantly lower than thresholds for acceptability judgments (usability of colors for printing or viewing, compared to target colors). 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 ΔE_{ab}^* , corresponds approximately to human judgements of perceptual difference, where CIELAB errors of 2 or 3 represent just noticeable color differences detectable by humans [46]. A difference of a few ΔE_{ab}^* is sometimes the tolerance used for accepting or rejecting color tolerances in e.g. the dyeing of colored fabrics. Here we are interested in using this perception based measure to delineate difference in medical image data. Since there is a standard transform from CIELAB to display-device RGB 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 [47]. Finally, we note that not all CIELAB colors map uniquely to a reproducible RGB color. Since gamut mapping is therefore involved [48], we can expect some clipping. We first address any outliers in CIELAB space by applying a robust outlier detection algorithm, in particular the least median of squares (LMS) method in [49], which operates entirely automatically. We then go on to clamp out-of-gamut colors to the closest RGB color, as done in [26].

G. Algorithm

To summarize the proposed method, Algorithm 1 highlights the steps of our algorithm 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.

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 -voxel volume or a large 1000^2 -pixel 2D image will result in a $1,000,000 \times 1,000,000$ dissimilarity matrix), it is important to address the issues of MDS complexity. The time complexity of MDS can be reduced to $O(N \log N)$ [50], 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 dimensionality 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 3D and formulate a polynomial regression fit over this smaller data set. Finally, we apply the resulting regression from n -D to 3D to all n -D pixel data points. In practice, we use

Algorithm 1 Perceptual Visualization of High-Dimensional, Manifold-Valued Medical Image Data

Input:

- 1) $\{f(\mathbf{x}_i)\}_{i=1}^N$, where $f(\mathbf{x}_i) : \mathbf{x}_i \in \mathbf{R}^d \rightarrow [f_1(\mathbf{x}_i), f_2(\mathbf{x}_i), \dots, f_n(\mathbf{x}_i)] \in \mathbf{R}^n$, 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 $\mathbf{x}_i = (x_i, y_i, z_i)$.
- 2) $\{(f(\mathbf{x}_i), C(\mathbf{x}_i))\}_{i=1}^P$, $P \geq 3$, i.e. at least 3 colors $C(\mathbf{x}_i) = [c_1(\mathbf{x}_i), c_2(\mathbf{x}_i), c_3(\mathbf{x}_i)]$ (e.g. in RGB space $C(\mathbf{x}_i) = [R(\mathbf{x}_i), G(\mathbf{x}_i), B(\mathbf{x}_i)]$) associated with at least three different pixel values $f(\mathbf{x}_i)$, $i = 1, \dots, P$, $P \geq 3$. Without loss of generality, we assume these are the first P elements in $\{f(\mathbf{x}_i)\}_{i=1}^N$.
- 3) Optionally, $d_{geodesic}(f(\mathbf{x}_i), f(\mathbf{x}_j)) : \mathbf{R}^n \times \mathbf{R}^n \rightarrow \mathbf{R}^+$; 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 Log-Euclidean distance metric [39], $D_{ij} = d_{T_{LE}}(T_1, T_2) = \sqrt{\text{trace}((\log m(T_1)) - (\log m(T_2)))^2}$, where T_1 and T_2 are two DTs and $\log m$ is the matrix logarithm, which is defined via the decomposition $T = U\Lambda U^t$ as $\log m(T) = U \text{diag}(\log(\text{diag}(\Lambda))) U^t$.

Output:

- 1) $\{g(\mathbf{x}_i)\}_{i=1}^N$, where $g(\mathbf{x}_i) : \mathbf{x}_i \in \mathbf{R}^d \rightarrow [g_1(\mathbf{x}_i), g_2(\mathbf{x}_i), g_3(\mathbf{x}_i)] \in \mathbf{R}^3$, i.e. N 3-D pixels forming a d -dimensional color (e.g. RGB) image, such that perceptual color distance $d_{perceptual}(g(\mathbf{x}_i), g(\mathbf{x}_j)) \propto d_{geodesic}(f(\mathbf{x}_i), f(\mathbf{x}_j)) \forall i, j \in \{1, 2, \dots, N\}$, 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 $d_{Euclidean}(f(\mathbf{x}_i), f(\mathbf{x}_j)) = \|f(\mathbf{x}_i) - f(\mathbf{x}_j)\|_2 \forall i, j \in \{1, 2, \dots, N\}$ s.t. $f(\mathbf{x}_i)$ is connected to $f(\mathbf{x}_j)$ in n -D. One of two connectedness criteria is applied: (i) $\|f(\mathbf{x}_i) - f(\mathbf{x}_j)\|_2 \leq \epsilon$ (ϵ -ISOMAP); (ii) $f(\mathbf{x}_i)$ is connected to its K -closest (using $\|\cdot\|_2$) neighbors (K -ISOMAP) [28]. This generates a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ whose N vertices (\mathcal{V}) correspond to $f(\mathbf{x}_i)$ and edges (\mathcal{E}) connect vertices that satisfy the connectedness criteria and are weighted by $d_{Euclidean}$.
 - **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 $\{g(\mathbf{x}_i)\}_{i=1}^N$, where $g(\mathbf{x}_i) : f(\mathbf{x}_i) \in \mathbf{R}^n \rightarrow [g_1(\mathbf{x}_i), g_2(\mathbf{x}_i), g_3(\mathbf{x}_i)] \in \mathbf{R}^3$, such that $d(g(\mathbf{x}_i), g(\mathbf{x}_j)) = d_{geodesic}(f(\mathbf{x}_i), f(\mathbf{x}_j)) \forall i, j \in \{1, 2, \dots, N\}$, or make the difference as small as possible, i.e. perform D_{ij} distance-preserving dimensionality reduction to 3D using MDS [22]. The $\{f(\mathbf{x}_i)\}_{i=1}^P$ samples (second input above) are now mapped to $\{g(\mathbf{x}_i)\}_{i=1}^P$ in 3D. All resulting $g(\mathbf{x}_i)$ points are in perceptually uniform CIELAB 3D space, however they have arbitrary scale, rotation, and translation.
 - **Step 5.** Convert $C(\mathbf{x}_i)$ (second input above) to CIELAB coordinates (e.g. RGB to CIELAB). Note that, generally, $\{g(\mathbf{x}_i)\}_{i=1}^P$ will not coincide with $\{C(\mathbf{x}_i)\}_{i=1}^P$ as desired.
 - **Step 6.** Transform $\{g(\mathbf{x}_i)\}_{i=1}^N$ using rotation $\mathbf{R} = USV^t$, isotropic scaling $s = \frac{1}{\sigma_g} \text{Tr}(DS)$, and translation $\mathbf{t} = \mu_C - sR\mu_g$, which are calculated such that $\{g(\mathbf{x}_i)\}_{i=1}^P$ are as close as possible to the corresponding $\{C(\mathbf{x}_i)\}_{i=1}^P$, where [47]: $\mu_C = \frac{1}{P} \sum_{i=1}^P C(\mathbf{x}_i)$, $\mu_g = \frac{1}{P} \sum_{i=1}^P g(\mathbf{x}_i)$, $\sigma_g^2 = \frac{1}{P} \sum_{i=1}^P \|g(\mathbf{x}_i) - \mu_g\|^2$, $\Sigma = \frac{1}{P} \sum_{i=1}^P (C(\mathbf{x}_i) - \mu_C)(g(\mathbf{x}_i) - \mu_g)^t$, $\Sigma = UDV^t$, $S = \begin{cases} I & , |\Sigma| \geq 0 \\ \text{diag}(1, 1, \dots, 1, -1) & , |\Sigma| < 0 \end{cases}$.
 - **Step 7.** Perform LMS robust outlier detection in CIELAB space [49], clamping outliers to the inlier boundary. Then, convert the resulting $\{g(\mathbf{x}_i)\}_{i=1}^N$ from CIELAB to the desired color space (e.g. RGB) for display and clamp out-of-gamut colors to the closest allowable color [26].
-

IV. RESULTS

In this section, we evaluate the accuracy of the polynomial regression approximation. Then, we test our method on hand-crafted synthetic images with three-dimensional pixel data sampled from a variety of underlying distributions. Next, we provide results on simulated DT data. We then test our method on case studies of normal subjects demonstrating our method's ability to highlight several anatomical and functional features in real medical images: the connectivity patterns through the corpus callosum relative to the nearby fornix assessed from brain DTMRI; the normal twisting pattern in the laminar sheet in cardiac DTMRI; and the heterogeneous putamen tracer distribution in a brain dynamic PET. We then provide a hybrid real-simulated example, in which we implant a simulated tumour into a real DTMRI image. Finally, through three other case studies, we examine our methods ability to highlight anatomical and functional pathologies: multiple sclerosis lesions in brain DTMRI; brain tumour (glioblastoma) in DTMRI; and renal functional abnormality in dynamic SPECT.

A. Evaluation of Polynomial Regression and Distance Preservation

A DTMRI display method should ideally preserve the original DT distances between pixels. The scatter plots of Figure 1 capture the disparity between original DT distances and the distances obtained following different dimensionality reduction operations. Pair-wise distances between all possible pairs of 1173 tensors from a brain region of a mid-sagittal DTMRI slice are used in this experiment. The original DT distances are measured in the Log-Euclidean space. Reducing the dimensionality of the DTs via MDS to 3D introduces some disparity compared to the original DT distances, which is captured by the spread of the scatter plot data points away from the diagonal 1:1 line (Fig. 1(a)). The approximation via polynomial regression also introduces some disparity compared to the original DT distances (Fig. 1(b)). Figure 1(c) isolates the disparity due to the polynomial regression approximation alone. To further quantify the latter disparity, Fig. 1(d) shows the error as a function of the percentage of data points (out of the whole image data) used to learn the regression parameters (Section III-H). We see from Fig. 1(d) 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%. Therefore, the regression approach does adequately well in representing a full MDS analysis. One commonly visualized scalar image derived from DTMRI is the 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 DT eigenvalues λ_i , $i = 1, 2, 3$, and is defined as $FA = \sqrt{\frac{3}{2} \sqrt{\sum_{i=1}^3 (\lambda_i - \bar{\lambda})^2} / \sqrt{\sum_{i=1}^3 \lambda_i^2}}$, where

$\bar{\lambda} = \frac{1}{3} \sum_{i=1}^3 \lambda_i$ is the mean diffusivity (MD). 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. Figures 1(e) and (f) show the result of using the common MD and FA measures to reduce the dimensionality of each DT to 1D. The large spread of the scatter plot data points away from the 1:1 line is a clear demonstration of the amount of information loss, i.e. DT distances not preserved, when using MD or FA.

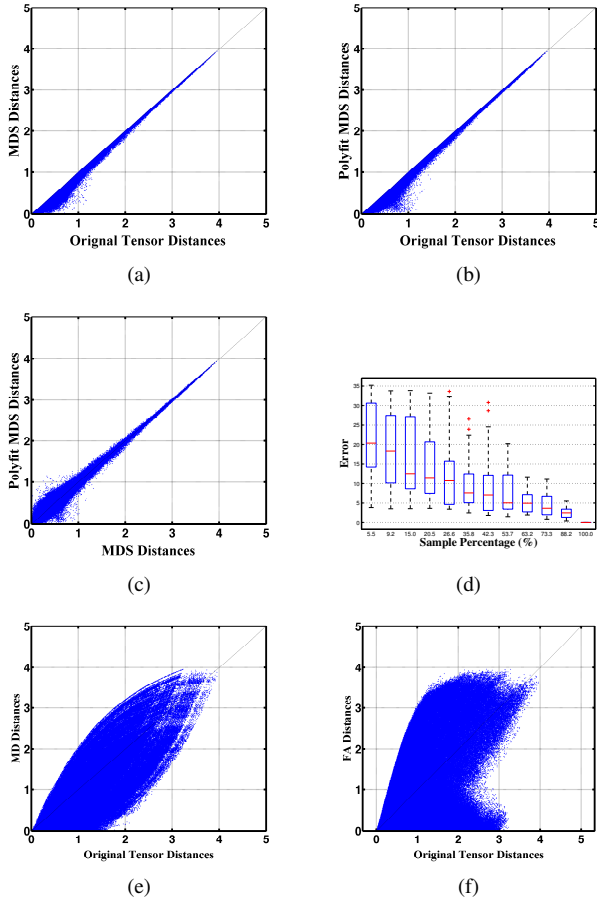


Fig. 1. Evaluation of polynomial regression. (a) Scatter plot of distances (between pairs of DT-pixels) after performing MDS (i.e. without the polynomial regression) vs. the original tensor distances. The 1:1 line is shown along the diagonal. (b) Scatter plot of the distances approximated via polynomial regression vs. the original tensor distances. (c) Scatter plot of distances when polynomial regression is used vs. when only MDS is used. (d) Error introduced by the polynomial regression (i.e. approximating MDS on the whole data set by polynomial regression on a subset of the data). The plot shows L^2 error versus percentage of pixels used in the approximation. A sample size of about 25% of the total number of pixels is employed in the polynomial regression approximation. (f,e) Scatter plot of MD and FA differences vs. the original DT distances. The large spread of the data points in (e) and (f) clearly demonstrate inferior distance preservation. Note that the MD and FA distances have been scaled to match the range of the original DT distances.

B. Synthetic Data with Pixel Dimensionality 3

We created 30×60 -pixel 2D images, each divided into two regions: right and left halves, R and L, each of size 30×30 . Each pixel is a three-dimensional vector. Four types of images were synthesized to produce different distributions of the pixels (Fig. 2): (i) R 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, i.e. the red channel (Fig. 2(a)); (ii) Same as (i) but with an oblique axis of maximal variation (Fig. 2(b)); (iii) the R and L pixels are sampled from two Gaussian distributions that together form a nonlinear space of samples in 3D (Fig. 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 R pixels (Fig. 2(d)). We show in Fig. 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 example of the more complicated 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.

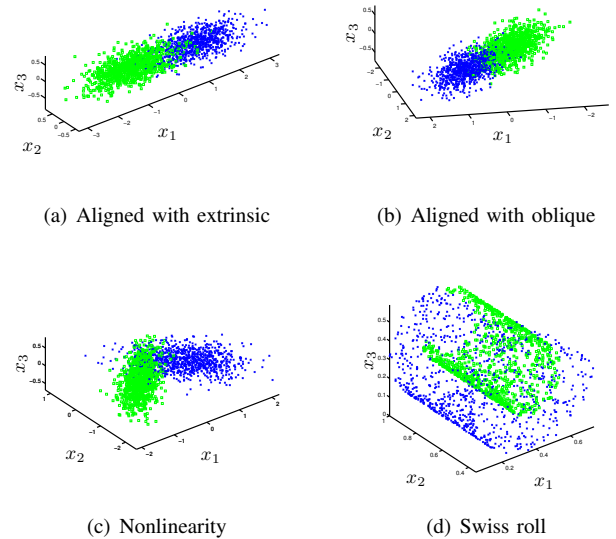


Fig. 2. Different types of synthetic data with pixel dimensionality 3 used to populate the right (R) and left (L) halves of images. Blue x’s correspond to R pixels and green squares to L pixels. (a) Data sampled from two Gaussians whose principal directions of variability align with the first extrinsic 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.

C. Simulated DTMRI

Neuroradiologists typically resort to viewing scalar images derived from the DTMRI field such as FA and MD. Clearly,




Fig. 3. Coloring the 3D pixels using different approaches. The four rows from top to bottom correspond to the four cases in Fig. 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).

scalar images are not able to capture the variability of tensors in 6D as seen in the two examples (I and II) in Fig. 4. In Example I, 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 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). In Example II of Fig. 4, 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. Although the FA is sensitive to DT changes as we move vertically in the image, the top half of the FA map appears symmetric to the bottom half. The MD, on the other hand, is completely oblivious to the change in tensors. Note that the colors used in our method to represent the DTs in Example I and Example II of Fig. 4 follow the legend in Fig. 5. In fact, all of our DTMRI results, from here forward, use the same color map to maintain a consistent representation of the diffusion tensors.

D. Case Studies of Normal Anatomy and Function

1) *Distinguishing the corpus callosum tissue from the fornix in brain DTMRI*: We start by showing in Fig. 6 qualitative DTMRI results of axial brain images of 4 normal subjects. Compared to MD and FA maps, note how our display method reveals differences in the underlying DTs via vivid color variations.

We turn our attention now to specific features in the brain; the CC and the fornix (Fig. 7). In Fig. 8, 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 (Fig. 8(c)), the CC body appears yellowish, whereas the fornix appears with reddish hues. The perceptual difference in color capture the difference in the underlying diffusion tensors, which in turn reflect the difference in the underlying connectivity patterns between these two regions: the corticocortical connections through the CC in contrast to the fibres through the fornix [51], where the latter connect the hippocampus to the septum,

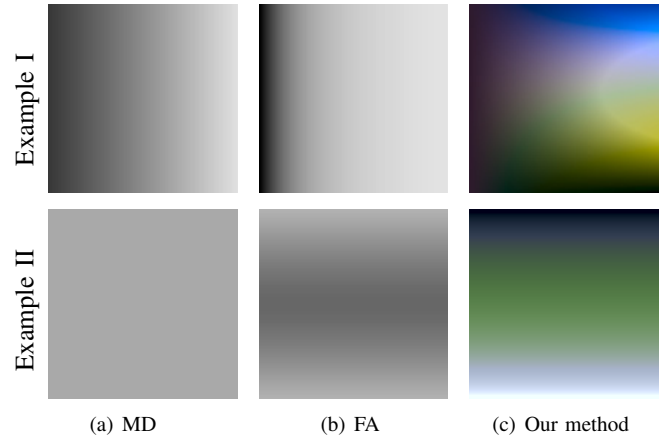


Fig. 4. Display of two example synthetic DTMRI. (Example I): The FA and MD images capture changes along the horizontal direction but fail to capture changes along the vertical direction. Our approach captures changes in DT along both directions. (Example II): Although the FA image captures changes in 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. The coloring is according to the legend in Fig. 5.




Fig. 5. Legend for the DTMRI visualization experiments. (a) The 3 key colors (c.f. Section III-F and Algorithm 1) with the following associated DTs: Pink: $DT_1 = \text{diag}(\epsilon, \epsilon, 1)$; Green: $DT_2 = \text{diag}(1, \epsilon, \epsilon)$; Blue: $DT_3 = \text{diag}(\epsilon, 1, \epsilon)$. (b) Three unit-square color palettes that act as a legend for visualizing DTMRI using our method. The palettes (from left to right) depict DT_2 vs. DT_1 , DT_3 vs. DT_1 , and DT_3 vs. DT_2 . The color at coordinate $(0 \leq \alpha \leq 1, 0 \leq \beta \leq 1)$ of the unit-square, corresponds to a DT interpolated in the Log-Euclidean space, e.g. $\text{expm}(\alpha \log m(DT_1) + \beta \log m(DT_2))$ for the first palette.

anterior nucleus of the thalamus and the mamillary bodies [52]. In contrast, the colors of the two regions are almost indistinguishable in the MD and FA maps (Fig. 8(a)-(b)), as well as in the structural MRI (Fig. 7).

2) *Highlighting the twisting pattern in the laminar sheet in cardiac DTMRI*: In the following case study, we explore the myocardial fibre microstructure. We compare MD, FA and our method for displaying 2D short axis cardiac DTMRI slices². Using our approach, the visualization of Fig. 9(c) shows the pixel colors changing as the myocardium is traversed from the epicardial surface radially inward towards the endocardial surface of the heart. The coloring is based on the underlying diffusion tensor image of the myocardium. Different tensors are colored with different colors. The key criterion of our method is that the displayed colors are chosen such that perceptual differences between the colors match the differences in the underlying diffusion tensors (as measured using the Log-Euclidean metric). The gradual epi- to endo-cardial change in

²The data was preprocessed using DTMRI bilateral filtering [53].




Fig. 6. Qualitative visualization of axial brain DTMRI slices of four subjects. Note how, compared to MD and FA, our method reveals differences in the underlying DTs via vivid color variations. The coloring is according to the legend in Fig. 5.




Fig. 7. Corpus callosum (CC). (left) Mid-sagittal plane in brain MR with a white box drawn around the CC. (right) The different anatomical parts of the CC referred to elsewhere in the paper (anterior to posterior): genu, rostrum, body, fornix, and splenium.




Fig. 8. Close up on the CC. Note how the fornix appears with the same color as the CC body when either MD or FA are used (as well as in the structural MR image in Fig. 7). Using our method, the fornix appears with a different (reddish) color, whereas the CC body appears yellowish reflecting the differences in the underlying DTs and connectivity patterns. The coloring is according to the legend in Fig. 5.

color indicates that diffusion tensors, and hence the diffusion properties, of the myocardium are changing in a unique way. Examining the results more carefully, we see that the colors are changing from greenish to reddish. Consulting the legend and key colors in Fig. 5, we note that the color green encodes tensors that have a primary diffusion direction along the 3D vector $(1, 0, 0)$, whereas the reddish color reflects diffusion close to the $(0, 0, 1)$ axis (pink in the legend). Therefore, the primary diffusion direction and, accordingly, the underlying fibre direction are gradually rotating as the myocardium is traversed. In fact, Tseng et al. [54] have observed a similar rotation in bovine myocardium using conventional histology with optical imaging (following fixing, sectioning and inking). This transmural rotation that our method highlights using the perceptual color mapping also agrees with other works that demonstrated the laminar sheet twisting in the human heart, including [55], [56].




Fig. 9. Visualization of short axis cardiac DTMRI slices. MD, FD, and our display method are compared in (a-c). The coloring is according to the legend in Fig. 5.

3) *Demonstrating heterogeneous putamen tracer distribution in a brain dynamic PET*: The photon activity in dPET is typically sampled non-uniformly with time. 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 to yield higher SNR [57]. 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 activity 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 reflected through the brightness of each pixel. However, vastly different curves can have similar integrals, and hence displaying these scalar fields can also be misleading (Fig. 10(a)-(c)).

To better illustrate the point above, in Fig. 10(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.

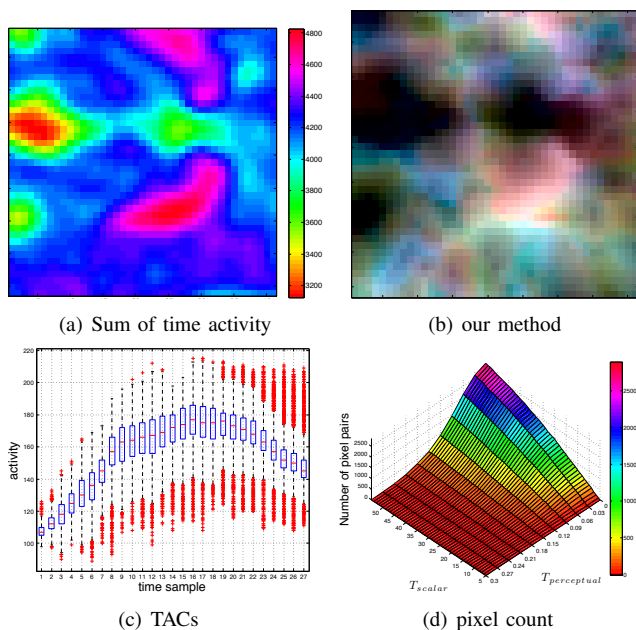


Fig. 10. Misleading brain dPET visualization. (a) Region of interest (ROI) of a dPET image with the sum of time activity visualized at each pixel. (b) The same ROI visualized using our method. (c) The distribution of TACs in this ROI (each TAC contains $n=27$ time samples). (d) The number (along the vertical axis) of neighboring pixel pairs (p, q) of which sub-figure (a) generates misleading visualization. This number is equal to the number of pixels pairs 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 $T_{perceptual}$.

To further illustrate how our method can improve over the standard visualization techniques for dPET, we provide results on two additional dPET data sets in Examples I and II of Fig. 11. The dPET in Example I is a Raclopride dPET, with spatial dimensions $128 \times 128 \times 63$, voxel size $2.11 \times 2.11 \times 2.43$ mm³, a 26 time-step TAC sampled at each voxel using the following non-uniform sampling: 6×30 s (i.e. 6 samples at 30 seconds sampling interval), 7×60 s, 5×120 s and 8×300 s. Example II has similar spatial dimensions and voxel size but uses FDopa dPET with 27 time-step TACs sampled at 6×30 s, 7×60 s, 5×120 s, 4×300 s, and 5×600 s. The similarity between two TACs, $f(x_i, t)$ and $f(x_j, t)$ at positions x_i and x_j , is measured using the weighted L^2 distance $d(f(x_i, t), f(x_j, t))^2 = \sqrt{\sum_{t=1}^T (f(x_i, t) - f(x_j, t))^2 z_t}$ where z_t is a weight that accounts for the non-uniform time sampling encountered in dPET [58]: shorter sampling intervals with lower SNR are weighted less.

We compare the results of five visualization methods: the last time frame of the TAC; the summation of the activity over all dPET frames; a mapping of principal component (PC) weights to the RGB color space; a mapping of PC weights to perceptual CIELAB; and finally our method. Neither visualizing the last frame nor the sum-over-frames conveys the full variability of the data. Using PCA allows us to map the 27-D data to a 3D color spaces (RGB or CIELAB), but as the data is non-linear in origin, PCA is not ideally suited for this task. In our method, on the other hand, visual difference in

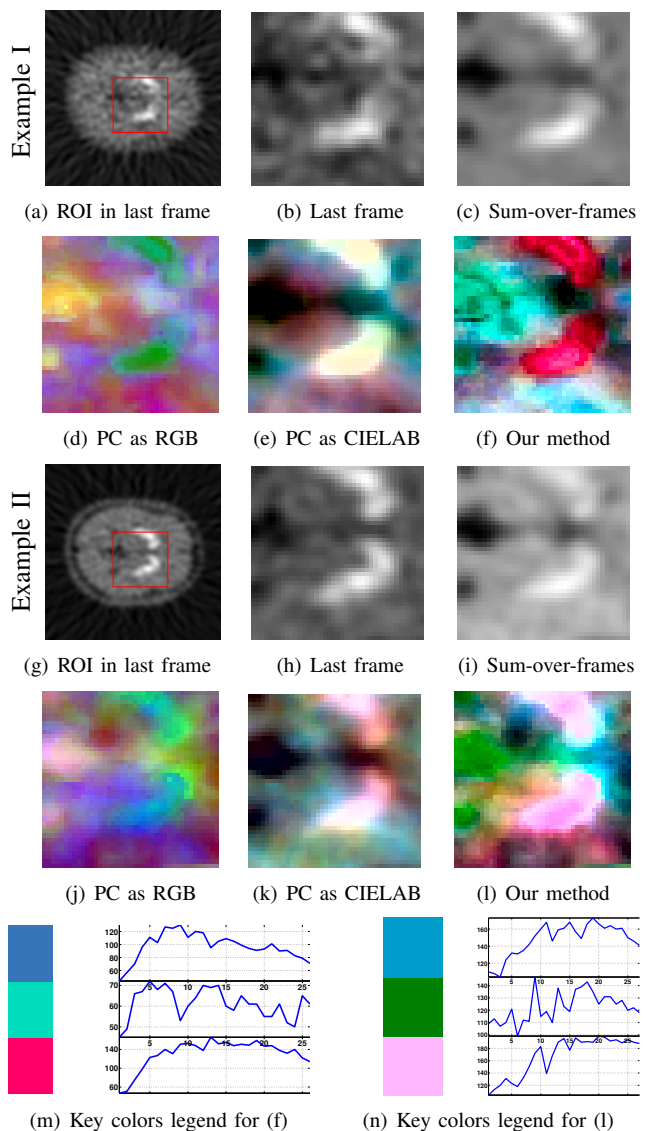


Fig. 11. Visualizing dPET brain images. Two data sets are used: Example I (a-f) and Example II (g-l). (a,g) An axial slice with the explored ROI outlined in red. The different methods of visualization used are: (b,h) the time activity in the last frame; (c,i) the summation of activity over all the TAC; (d,j) the weights of the first 3 PC modes used as the R, G, and B color channels; (e,k) the PC weights as CIELAB channels; and (f,l) our method. (m) and (n) show the key TACs and corresponding colors used as a legend for visualizing (f) and (l), respectively (as explained in Section III-F and the 2^{nd} input of Algorithm 1).

colors are optimally chosen to correspond to the underlying differences in TACs. We note in Fig. 11(f) and (l) how our approach not only highlights the putamen in the brain but also shows different tracer uptake properties within the putamen itself (dark to bright red in (f) and pink to white in (l)).

It is important to emphasize that since the image data visualized contains, at each pixel, a 27-dimensional vector sampling the TAC, the “ground truth” changes from pixel to pixel can only be asserted by examining the changes in these 27-dimensional vectors. However, there is a tradeoff between providing a display that is perfectly faithful to the underlying changes in 27-dimensional space, and providing a display that is simple and intuitive to examine. Our method displays color

images that are as faithful to the high-dimensional changes as much as possible given the 3 color channel limit, i.e. our method can, at best, capture as much changes as the 3-color limit allows, not more. Nevertheless, our method is the one most faithful to the underlying ground truth changes when compared to other commonly used display methods (displaying the last frame, the integration of activity), or even linear or nonlinear dimensionality reduction based color display methods (e.g. PCA or MDS) that do not utilize a perceptual color space.

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 [59], [60], [61], [62], [63]. We performed the next experiment to mimic the existence of a region of abnormality in a brain DTMRI. The DTs in a particular region of interest (ROI) in the CC between the rostrum and genu (Fig. 7) were manipulated to simulate a pathological condition (Fig. 12). This implanted simulated pathology is designed to be ‘transparent’ to classical display techniques and, at the same time, be clearly picked up by our coloring approach. The idea from this experiment is to show a concrete (albeit hand-crafted) example, with valid diffusion tensors, in which striking differences in diffusion tensors could be completely unseen using traditional display methods. 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. Fig. 12(c) shows the CC rendered with our method (before and after simulating the pathology) compared to displays using FA and MD maps (Fig. 12(a)-(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. 12. Simulating pathology within a real brain DTMRI image. Before (top) and after (bottom) pathology simulation. Left to right: MD, FA, and our method. The pathology is clearly visible (in red) using our method. This hand-crafted example of a simulated pathology, with valid diffusion tensors, show how striking differences in diffusion tensors could be completely unseen using traditional (MD and FA) display methods.

F. Case Studies of Pathology

1) *Distinct appearance of multiple sclerosis lesions in brain DTMRI:* In this case study, we evaluate our method on a DTMRI data of a multiple sclerosis patient. The original 3D data is of dimensions $60 \times 112 \times 112$, with 2.2 mm isotropic

resolution. A single acquisition with 32 gradient directions and a B-value 750 s/mm^2 was used. Fig. 13, shows an axial slice in the DTMRI volume displayed using MD (Fig. 13(a)); FA (Fig. 13(b)); and using our method (Fig. 13(c)). Note how the lesion appears in blueish color in the close up shown in Fig. 13(d). We also note that our method preserved 6 times more variations of the original DTMRI data compared to either MD or FA.




Fig. 13. Brain DTMRI with multiple sclerosis lesions. Three display methods are used: MD map in (a); FA map in (b); and using our proposed method in (c). A close up on the ROI surrounding the lesion is shown in (d). The key colors and the legend palettes of Fig. 5 are used for this figure as well. Note the appearance of the lesion as a dark blue spot in (d).

2) *Highlighting areas of edema and axonal disorganization surrounding brain tumours:* We now examine the appearance of a frontal glioblastoma (a highly malignant brain tumor) in DTMRI. The original 3D data is of dimensions $256 \times 256 \times 20$, with $0.86 \times 0.86 \times 5.0 \text{ mm}^3$ voxel size, B-value 1000 s/mm^2 and 6 gradient directions in a single acquisition [64]. In Fig. 14, we display an axial slice in the DTMRI volume using an MD map (Fig. 14(a)); an FA map (Fig. 14(b)); and using our method (Fig. 14(c)). Close ups on the regions of the glioblastoma are shown in Fig. 14(d-f). Note how using our method a blueish hue appears around the area of the tumour (Fig. 14(f)) which may reflect the axonal disorganization and edema surrounding the tumour that affects the underlying fibre connections and the measured DTs [65], [66]. We also note that our method preserved 6 times more variations of the original DTMRI data compared to MD, and 7.5 times more than FA.

3) *Demonstrating areas of renal functional abnormality in dynamic SPECT:* We explore in this case study a kidney dSPECT data set used to assess renal function. Every 2 mm isotropic pixel stores a 48-sample TAC collected during the acquisition of a dynamic $64 \times 64 \times 64$ volume. Fig. 15 depicts coronal slices through the field of view showing both the left and the right kidneys using different display methods: integral of TAC (Fig. 15(a)); the last (48th) frame of the dynamic data (Fig. 15(b)); and the result of applying our visualization method (Fig. 15(c)). Our display method highlights an abnormality in the lower third of the left kidney (rendered in a yellowish color). This abnormality was clinically verified in




Fig. 14. Brain DTMRI of a frontal glioblastoma. Three display methods are used: (a) MD; (b) FA; and (c) our proposed method. (d-f) Close ups on the ROI surrounding the glioblastoma. The key colors and the legend palettes of Fig. 5 are used for this figure as well. Note the textured blue region around the tumour in (f), which may be attributed to the edema and axonal disorganization which affect the fibre patterns.

[67].

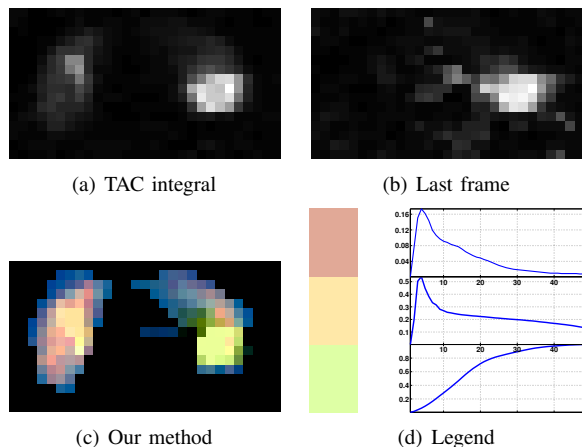


Fig. 15. Left and right kidneys shown in a coronal slice from a dSPECT study of renal function. Three display methods are compared: (a) The integral of the dSPECT TAC at every pixel; (b) the activity of the last frame in the TAC; and (c) our method. Note how the pathological lower third of the left kidney (in the right side of the image in (c)) stands out in bright yellow, relative to the other regions of the kidneys. (d) The three key colors (brown, orange, yellow) and corresponding TACs (Section III-F and Algorithm 1) used to produce (c).

V. CONCLUSIONS

There is an increasing urgency for developing methods to interpret high-dimensional, manifold-valued medical image (e.g. DTMRI, dPET, or dSPECT). Popular methods for visualizing these modalities rely on reducing the dimensionality without minimizing the loss of the underlying information and/or without taking into account perceptual factors. We propose a display method whose key criterion is to color pixels such that perceptual differences between the colors remain as faithful as possible to the dissimilarities between the underlying high-dimensional pixels. Our method relies on three key ideas: (i) estimating geodesic distances on the manifolds of pixel data (using analytical derivations or manifold learning methods); (ii) reducing the dimensionality of the data to a 3-channel

color image via a distance-preserving transformation (using MDS or ISOMAP); and (iii) mapping the high-dimensional dissimilarities into perceptual color differences while maximizing the utilization of the color gamut (by operating in the CIELAB color space). We presented results on several synthetic, hand-crafted, and real (normal and pathological) case studies (cardiac, neurological, and renal in DTMRI, dPET and dSPECT), which demonstrated how standard display methods fall short of reflecting the true underlying variability in the data, while our coloring method highlights anatomical and functional features of interest while being faithful (as much as possible given the 3-color channel limit) to the underlying data.

We anticipate improved display results from our method if, in addition to 3D color, opacity is also used in the visualization as a 4th degree of freedom. Further, more powerful visualization may be obtained if the colors extracted at each datum using our method are used to color different types of glyphs. We also foresee no evident obstacles in using our method to color other (possibly non-medical) high-dimensional data fields or scattered data, e.g. geospatial data.

Visualization of multi-modal data at the same time (with a single color image) is also an interesting application to explore in future work. There are, however, two primary prerequisites for this exploration: the need for proper spatial alignment (registration) and the need for a proper distance dissimilarity metric that combines both modalities. Registration of multi-modal data is an actively pursued research area and we can benefit from published results and available software to perform the alignment. However, the latter problem requires further exploration (e.g. is a weighted sum of multiple dissimilarity metrics, one for each modality, justifiable?)

At the moment our approach is not used to directly support any specific clinical task. Nevertheless, our technique's ability to discern information that is otherwise lost using the traditional display approaches will be used for in-depth studies related to specific clinical application, e.g. the display of multiple sclerosis lesions from DTMRI, appearance of white matter injury in DTMRI of pre-term infants, and dPET imaging for the study of Parkinson's patients. An important next step is therefore to collaborate closely with doctors or radiologists to assess the objective clinical value of our approach.

To perform gamut mapping [48], our algorithm included the last step of LMS outlier detection [49] followed by clipping out-of-gamut colors to the nearest reproducible color [26]. A potential alternative approach is to modify the criteria we optimize such that it not only encodes high-dimensional dissimilarity as perceptual differences, but also encourages the resulting colors to be reproducible. A related potential extension could involve the adoption of automatic color scheme selection algorithms, for example using the ColorBrewer [68], [69].

Finally, we point out that we have made the implementation of our algorithm available to the scientific community³, so that other researchers can use our method to possibly gain further

³The software is available for download from <http://perceptvis.cs.sfu.ca>

insight into their own data sets.

VI. ACKNOWLEDGEMENTS

We thank the reviewers and editors for their valuable feedback that resulted in an improved paper. GH and MD are funded through the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery grants. CM is funded by an NSERC Graduate (Doctoral) Scholarship and a Michael Smith Foundation for Health Research (MSFHR) Senior Graduate Studentship. We thank Drs. Patrick A. Helm and Raimond L. Winslow at the Center for Cardiovascular Bioinformatics and Modeling and Dr. Elliot McVeigh at the Dept. of Biomedical Engineering, Johns Hopkins University School of Medicine, for provision of the heart DTMRI data; Drs. Alex MacKay, Roger Tam, and Corree Laule for the DTMRI multiple sclerosis lesions data sets; Drs. Elizabeth Bullitt and Marcel Prastawa for providing the brain tumour DTMRI data; Prof. Anna Celler for dSPECT data; and the PET-SORTEO for the dPET images [70]. We used the manifold learning software for MATLAB provided by Todd Wittman, Department of Mathematics, University of Minnesota⁴. We thank Brian Booth and Eli Gibson from the Medical Image Analysis Lab at Simon Fraser University for assistance in, respectively, reconstructing the diffusion tensors and preparing the synthetic tumor example.

REFERENCES

- [1] P. Fletcher, S. Venkatasubramanian, and S. Joshi, "The geometric median on riemannian manifolds with application to robust atlas estimation," *Neuroimage*, vol. 45, no. 1, pp. S143–S152, 2009.
- [2] X. Shi, M. Styner, J. Lieberman, J. Ibrahim, W. Lin, and H. Zhu, "Intrinsic regression models for manifold-valued data," in *Medical Image Computing and Computer-Assisted Intervention*, ser. Lecture Notes in Computer Science, G.-Z. Yang, D. Hawkes, D. Rueckert, A. Noble, and C. Taylor, Eds. Springer Berlin / Heidelberg, 2009, vol. 5762, pp. 192–199.
- [3] S. Sommer, F. Lauze, S. Hauberg, and M. Nielsen, "Manifold valued statistics, exact principal geodesic analysis and the effect of linear approximations," *European Conference on Computer Vision*, pp. 43–56, 2010.
- [4] Z. Peng and R. Laramée, "Higher dimensional vector field visualization: A survey," *Theory and Practice of Computer Graphics*, pp. 149–163, 2009.
- [5] D. L. Pham, C. Xu, and J. L. Prince, "A survey of current methods in medical image segmentation," in *Annual Review of Biomedical Engineering*. Annual Reviews, 2000, vol. 2, pp. 315–338.
- [6] P. C. Wong and R. D. Bergeron, "30 years of multidimensional multivariate visualization," in *Scientific Visualization, Overviews, Methodologies, and Techniques*. Washington, DC, USA: IEEE Computer Society, 1997, pp. 3–33.
- [7] D. Keim, "Designing pixel-oriented visualization techniques: theory and applications," *IEEE Transactions on Visualization and Computer Graphics*, vol. 6, no. 1, pp. 59–78, 2000.
- [8] N. Jacobson and M. Gupta, "Design goals and solutions for display of hyperspectral images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 43, no. 11, pp. 2684–2692, 2005.
- [9] B. Wandell and L. Silverstein, *Foundations of vision*. Sinauer Associates Sunderland, MA, 1995.
- [10] G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae (Wiley Series in Pure and Applied Optics)*, 2nd ed. Wiley-Interscience, 2000.
- [11] N. Jacobson, M. Gupta, and J. Cole, "Linear fusion of image sets for display," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 10, pp. 3277–3288, Oct. 2007.
- [12] P. Ready and P. Wintz, "Information extraction, SNR improvement, and data compression in multispectral imagery," *IEEE Transactions on Communications*, vol. 21, no. 10, pp. 1123–1131, 1973.
- [13] J. Tyo, A. Konsolakis, D. Diersen, and R. Olsen, "Principal-components-based display strategy for spectral imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 41, no. 3, pp. 708–718, 2003.
- [14] M. R. Gupta and N. N. P. Jacobson, "Wavelet principal component analysis and its application to hyperspectral images," *IEEE International Conference on Image Processing*, pp. 1585–1588, 2006.
- [15] S. Kaewpjit, J. Le Moigne, and T. El-Ghazawi, "Automatic reduction of hyperspectral imagery using wavelet spectral analysis," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 41, no. 4, pp. 863–871, 2003.
- [16] M. S. Drew and G. Hamarneh, "Visualizing diffusion tensor dissimilarity using an ICA based perceptual color metric," in *Color Imaging*, 2007, pp. 42–47.
- [17] J. Wang and C.-I. Chang, "Independent component analysis-based dimensionality reduction with applications in hyperspectral image analysis," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 44, no. 6, pp. 1586–1600, 2006.
- [18] R. F. Erbacher, D. L. Gonthier, and H. Levkowitz, "Color icon: A new design and a parallel implementation," in *SPIE*, G. G. Grinstein and R. F. Erbacher, Eds., vol. 2410, 1995, pp. 302–312.
- [19] C. G. Healey and J. T. Enns, "Large datasets at a glance: combining textures and colors in scientific visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 5, no. 2, pp. 145–167, 1999.
- [20] V. Tsagaris and V. Anastassopoulos, "Multispectral image fusion for improved RGB representation based on perceptual attributes," *International Journal of Remote Sensing*, vol. 26, pp. 3241–3254(14), 2005.
- [21] Z. Fang, T. Möller, G. Hamarneh, and A. Celler, "Visualization and exploration of time-varying medical image data sets," in *Graphics Interface (GI)*, 2007, pp. 281–288.
- [22] I. Borg and P. J. F. Groenen, *Modern Multidimensional Scaling: Theory and Applications (Springer Series in Statistics)*, 2nd ed. Springer, Berlin, 2005.
- [23] K. Rasche, R. Geist, and J. Westall, "Detail preserving reproduction of color images for monochromats and dichromats," *Computer Graphics and Applications, IEEE*, vol. 25, no. 3, pp. 22–30, 2005.
- [24] —, "Re-coloring images for gamuts of lower dimension," *Computer Graphics Forum*, vol. 24, pp. 423–432(10), 2005.
- [25] A. Brun, H.-J. Park, H. Knutsson, and C.-F. Westin, "Coloring of DT-MRI fiber traces using Laplacian eigenmaps," in *Computer Aided Systems Theory (EUROCAST'03), Lecture Notes in Computer Science*, R. M. Diaz and A. Q. Arencibia, Eds., vol. 2809. Springer Verlag, 2003, pp. 564–572.
- [26] N. Neophytou and K. Mueller, "Color-space CAD: Direct gamut editing in 3D," *IEEE Computer Graphics and Applications*, vol. 28, no. 3, pp. 88–98, 2008.
- [27] L. Wang, J. Giesen, K. T. McDonnell, P. Zolliker, and K. Mueller, "Color design for illustrative visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 6, pp. 1739–1754, 2008.
- [28] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, "A Global Geometric Framework for Nonlinear Dimensionality Reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [29] S. T. Roweis and L. K. Saul, "Nonlinear Dimensionality Reduction by Locally Linear Embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, 2000.
- [30] P. J. Basser, J. Mattiello, and D. LeBihan, "MR diffusion tensor spectroscopy and imaging," *Biophysical Journal*, vol. 66, no. 1, pp. 259–267, 1994.
- [31] E. Stejskal and J. Tanner, "Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient," *Journal of Chemical Physics*, vol. 42, pp. 288–292, Jan. 1965.
- [32] X. Pennec, P. Fillard, and N. Ayache, "A Riemannian framework for tensor computing," *International Journal of Computer Vision*, vol. 66, no. 1, pp. 41–66, 2006.
- [33] Y. Assaf and O. Pasternak, "Diffusion tensor imaging (DTI)-based white matter mapping in brain research: a review," *Journal of Molecular Neuroscience*, vol. 34, no. 1, pp. 51–61, 2008.
- [34] C.-F. Westin, S. E. Maier, H. Mamata, A. Nabavi, F. A. Jolesz, and R. Kikinis, "Processing and visualization for diffusion tensor MRI," *Medical Image Analysis*, vol. 6, no. 2, pp. 93–108, 2002.
- [35] D. L. Bailey, D. W. Townsend, P. E. Valk, and M. N. Maisey, Eds., *Positron Emission Tomography: Basic Sciences*. Springer, 2005.
- [36] M. N. Wernick and J. N. Aarsvold, Eds., *Emission tomography: the fundamentals of PET and SPECT*. Academic Press, 2004.

⁴Available from <http://www.math.umn.edu/~wittman/mani/>

- [37] H. Zaidi, Ed., *Quantitative Analysis in Nuclear Medicine Imaging*. Springer, 2006.
- [38] R. E. Carson, "Tracer kinetic modeling in PET," *Positron Emission Tomography*, pp. 127–159, 2005.
- [39] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache, "Log-Euclidean metrics for fast and simple calculus on diffusion tensors," *Magnetic Resonance in Medicine*, vol. 56, no. 2, pp. 411–421, 2006.
- [40] G. Kindlmann, R. S. J. Estepar, M. Niethammer, S. Haker, and C.-F. Westin, "Geodesic-loxodromes for diffusion tensor interpolation and difference measurement," in *Medical Image Computing and Computer-Assisted Intervention*, ser. Lecture Notes in Computer Science 4791, Brisbane, Australia, 2007, pp. 1–9.
- [41] Z. Wang and B. Vemuri, "An affine invariant tensor dissimilarity measure and its applications to tensor-valued image segmentation," *Computer Vision and Pattern Recognition*, vol. 1, pp. 1–228–1–233 Vol.1, 2004.
- [42] E. Yorukk, B. Acar, and R. Bammer, "A physical model for DT-MRI based connectivity map computation," *Medical Image Computing and Computer-Assisted Intervention*, pp. 213–220, 2005.
- [43] K. Pearson, "On lines and planes of closest fit to systems of points in space," *Philosophical Magazine*, vol. 2, no. 6, pp. 559–572, 1901.
- [44] P. Comon, "Independent component analysis, a new concept?" *Signal Process.*, vol. 36, no. 3, pp. 287–314, 1994.
- [45] R. Ramanath and M. S. Drew, "Color perception," *Wiley Encyclopedia of Computer Science and Engineering*, 2008.
- [46] M. Stokes, M. Fairchild, and R. Berns, "Precision requirements for digital color reproduction," *ACM Transactions on Graphics*, vol. 11, no. 4, pp. 406–422, 1992.
- [47] S. Umeyama, "Least-squares estimation of transformation parameters between two point patterns," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 4, pp. 376–380, 1991.
- [48] J. Morovic and M. R. Luo, "The fundamentals of gamut mapping: A survey," *Journal of Imaging Science and Technology*, vol. 45, no. 3, pp. 283–290, 2001.
- [49] P. J. Rousseeuw and A. M. Leroy, *Robust Regression and Outlier Detection*. Wiley, 1987.
- [50] F. Jourdan and G. Melançon, "Multiscale hybrid MDS," *Information Visualisation*, pp. 388–393, July 2004.
- [51] S. Wakana, H. Jiang, L. M. Nagae-Poetscher, P. C. M. van Zijl, and S. Mori, "Fiber Tract-based Atlas of Human White Matter Anatomy," *Radiology*, vol. 230, no. 1, pp. 77–87, 2004.
- [52] J. Fitzsimmons, M. Kubicki, K. Smith, G. Bushell, R. S. J. Estepar, C. Westin, P. Nestor, M. Niznikiewicz, R. Kikinis, R. McCarley, and M. Shenton, "Diffusion tractography of the fornix in schizophrenia," *Schizophrenia Research*, vol. 107, no. 1, pp. 39–46, 2009.
- [53] G. Hamarneh and J. Hradsky, "Bilateral filtering of diffusion tensor magnetic resonance images," *IEEE Transactions on Image Processing*, vol. 16, no. 10, pp. 2463–2475, 2007.
- [54] W. Tseng, V. Wedeen, T. Reese, R. Smith, and E. Halpern, "Diffusion tensor MRI of myocardial fibers and sheets: Correspondence with visible cut-face texture," *Journal of Magnetic Resonance Imaging*, vol. 17, no. 1, pp. 31–42, 2003.
- [55] I. J. Legrice, P. J. Hunter, and B. H. Smaill, "Laminar structure of the heart: a mathematical model," *Am J Physiol Heart Circ Physiol*, vol. 272, no. 5, pp. H2466–2476, 1997.
- [56] D. F. Scollan, A. Holmes, R. Winslow, and J. Forder, "Histological validation of myocardial microstructure obtained from diffusion tensor magnetic resonance imaging," *American Journal of Physiology Heart and Circulatory Physiology*, vol. 275, no. 6, pp. H2308–2318, 1998.
- [57] K. C. Hongbin Guo, Rosemary Renaut and E. Reiman, "Clustering huge data sets for parametric PET imaging," *Biosystems*, vol. 71, no. 1–2, pp. 81–92, 2003.
- [58] A. Saad, G. Hamarneh, T. Möller, and B. Smith, "Kinetic modeling based probabilistic segmentation for molecular images," in *Medical Image Computing and Computer-Assisted Intervention*, 2008, pp. 244–252.
- [59] I. Elshafiey, "Diffusion tensor magnetic resonance imaging of lesions in multiple sclerosis patients," *Proceedings of the Nineteenth National Radio Science Conference*, pp. 626–633, 2002.
- [60] H. Kitzler, W. Bengler, 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.
- [61] M. Schlüter, B. Stieltjes, H. K. Hahn, J. Rexilius, O. Konrad-Verse, and H. Peitgen, "Detection of tumour infiltration in axonal fibre bundles using diffusion tensor imaging," *The International Journal of Medical Robotics and Computer Assisted Surgery*, vol. 1, no. 3, pp. 80–86, 2006.
- [62] S. Price, N. Burnet, T. Donovan, H. Green, A. Pena, N. Antoun, J. Pickard, T. Carpenter, and J. Gillard, "Diffusion tensor imaging of brain tumours at 3T: A potential tool for assessing white matter tract invasion?" *Clinical Radiology*, vol. 58, pp. 455–462(8), 2003.
- [63] B. Stieltjes, M. Schlter, B. Didinger, M.-A. Weber, H. K. Hahn, P. Parzer, J. Rexilius, O. Konrad-Verse, H. Peitgen, and M. Essig, "Diffusion tensor imaging in primary brain tumors: Reproducible quantitative analysis of corpus callosum infiltration and contralateral involvement using a probabilistic mixture model," *NeuroImage*, vol. 31, no. 2, pp. 531 – 542, 2006.
- [64] M. Prastawa, E. Bullitt, and G. Gerig, "Simulation of brain tumors in MR images for evaluation of segmentation efficacy," *Medical Image Analysis*, vol. 13, no. 2, pp. 297 – 311, 2009, includes Special Section on Functional Imaging and Modelling of the Heart.
- [65] J. Corso, E. Sharon, S. Dube, S. El-Saden, U. Sinha, and A. Yuille, "Efficient multilevel brain tumor segmentation with integrated bayesian model classification," *IEEE Transactions on Medical Imaging*, vol. 27, no. 5, pp. 629–640, 2008.
- [66] S. Lu, D. Ahn, G. Johnson, and S. Cha, "Peritumoral Diffusion Tensor Imaging of High-Grade Gliomas and Metastatic Brain Tumors," *American Journal of Neuroradiology*, vol. 24, no. 5, pp. 937–941, 2003.
- [67] T. Humphries, A. Saad, A. Celler, G. Hamarneh, T. Möller, and M. Trummer, "Segmentation-based regularization of dynamic SPECT reconstructions," in *IEEE Medical Imaging / Nuclear Science Conference*, 2009, pp. 2849–2852.
- [68] C. Brewer, "Color use guidelines for data representation (<http://www.colorbrewer.org>)," in *Proc. Section on Statistical Graphics*, 1999, pp. 55–60.
- [69] M. Wijffelaars, R. Vliegen, J. van Wijk, and E.-J. van der Linden, "Generating color palettes using intuitive parameters," *Computer Graphics Forum*, vol. 27, no. 4, pp. 743–750, 2008.
- [70] A. Reilhac, G. Batan, C. Michel, N. Costes, and A. Evans, "PET-SORTEO: a platform for simulating realistic PET studies," in *Nuclear Science Symposium*, vol. 7, Oct. 2004, pp. 4053–4057.