4-21-2009

NONLINEAR TUBE-FITTING FOR THE ANALYSIS OF ANATOMICAL AND FUNCTIONAL STRUCTURES

Jeff Goldsmith
Johns Hopkins Bloomberg School of Public Health, Department of Biostatistics, jgoldsmi@jhsph.edu

Brian S. Caffo
Johns Hopkins Bloomberg School of Public Health, Department of Biostatistics

Ciprian Crainiceanu
Johns Hopkins Bloomberg School of Public Health, Department of Biostatistics

Daniel Reich
Johns Hopkins Hospital, School of Medicine, Radiology and Radiological Science

Yong Du
Johns Hopkins University, School of Medicine, Radiology-Medical Imaging Physics

See next page for additional authors

Authors
Jeff Goldsmith, Brian S. Caffo, Ciprian Crainiceanu, Daniel Reich, Yong Du, and Craig Hendrix

Nonlinear Tube-Fitting for the Analysis of Anatomical and Functional Structures

Jeff Goldsmith, Brian Caffo, Ciprian Crainiceanu,
Daniel Reich, Yong Du, Craig Hendrix

20 April 2009

Abstract

We are concerned with the estimation of the exterior surface of tube-shaped anatomical structures. This interest is motivated by two distinct scientific goals, one dealing with the distribution of HIV microbicide in the colon and the other with measuring degradation in white-matter tracts in the brain. Our problem is posed as the estimation of the support of a distribution in three dimensions from a sample from that distribution, possibly measured with error. We propose a novel tube-fitting algorithm to construct such estimators. Further, we conduct a simulation study to aid in the choice of a key parameter of the algorithm, and we test our algorithm with validation study tailored to the motivating data sets. Finally, we apply the tube-fitting algorithm to a colon image produced by single photon emission computed tomography (SPECT) and to a white-matter tract image produced using diffusion tensor imaging (DTI).

1 Introduction

A common problem in biomedical imaging research is to mathematically model anatomical structures and to summarize them in an appropriate space. In this manuscript, we focus on modeling tube-like anatomical structures, such as the colon or white matter fiber bundles in the brain. In our setting, the objects of measurement are measured by biological signals represented in a two- or three-dimensional array obtained via imaging or some other indirect measurement of anatomy or function. Finding the best mathematical representation of the tube to quantify the anatomical or functional image remains a difficult - and neglected - problem in statistics. In this manuscript, we develop an algorithm for fitting tubes to collections of points and apply this algorithm to data from two motivating examples based on different medical imaging modalities.

Figure 1: Example SPECT colon image shown from two angles. Sizes represent image intensity.

Our first application is to single-photon-emission computed tomography (SPECT) images from an experiment to evaluate the distributional penetrance of anti-human immunodeficiency virus (HIV) microbicidal lubricants in the colon. SPECT images are produced by applying computed-tomography techniques to projections of photons emitted by a radioactive tracer. In this experiment, a radiolabelled over-the-counter lubricant was distributed in a subject's colon. Some example data is shown in Figure 1, where the size corresponds to the concentration of the tracer at that location.

Knowledge of the distributional penetrance of the tracer, along with knowledge of the distribution of HIV-infected semen after intercourse with an infected partner, would give crucial information regarding efficacy of the treatment for preventing transmission. Currently, we study only the distributional penetrance of the lubricant via SPECT imaging. Our goals are to obtain an accurate tube through the tracer to outline the colon, along with a metric to measure the tube's extent at various orthogonal cross-sections.

Our second application is to diffusion tensor imaging (DTI) tractography. DTI is a magnetic resonance imaging (MRI) technique that measures the diffusivity of water in the brain along several gradients. Such information is useful as water diffuses anisotropically along white-matter tracts. Hence, DTI gives more detailed images of white-matter anatomy compared to standard MRI techniques. In fact, anisotropic diffusion can be used to reconstruct bundles of white-matter tracts, a process called tractography [3, 4, 31, 28]. An example tract is given in Figure 2.

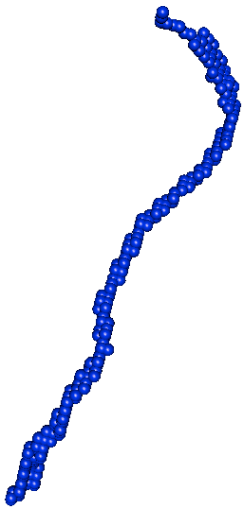

Figure 2: Example DT tract image.

Tractography holds great promise as a quantitative measure of white-matter health, though tractography methods are still in development. An example of potential application of tractography is to the study of multiple sclerosis (MS), which causes demyelination in the white-matter. (Myelin is the fatty insulation that surrounds axons and that is found primarily in the white-matter.) Individuals with MS can suffer profound disability, such as loss of vision and motor function. The ability to quantify tissue damage using DTI tractography has important clinical and research implications. Several parameters of the tracts, including shape, volume, and anisotropy may be useful for monitoring the progression of MS.

In both of these applications we seek a method of mathematically representing the tracer or anatomical structure with an envelope or tube that "represents" the object in imaging space. Here, what is meant by "represents" is context- and modality-specific, as different imaging techniques and settings can result in vastly different goals for estimating the tube. A strength of our proposed method is its ability to accommodate different settings.

We distinguish the tube-fitting problem from the volume of excellent work on simultaneous confidence bounds around estimated functions. In our case, the tube is not a measure of uncertainty but is instead is the estimand of interest. Our procedure follows these steps: *i*. fit a curve through the center of the object; *ii*. calculate tangent lines along the curve; *iii*. calculate planes normal to the tangent lines; *iv*. project nearby points onto these planes; *v*. fit a circle, ellipse or convex hull to the points in the plane. The resulting object represents an estimate of the structure. The curve from *i*. represents the "spine" of the tube. It is also useful in our applications to represent the metric by which measures of extent

of the tube are taken. This component relies on existing methodology; the remaining four and the application to two imaging modalities represent the methodological advances of this manuscript.

We apply the tube-fitting algorithm to an example of each type of image. The results indicate that the procedure could be used in the SPECT application as a replacement for the an invasive sigmoidoscope procedure, which is currently used instead of image processing. In the second application, our algorithm is an improvement over the current approach that examines tracts by looking at a sequence of axial slices along the image. Taking measurements along axial slices does not produce satisfactory results when the long axis of the tube is not normal to the axial plane (see Section 6).

The manuscript is structured as follows. Section 2 describes the data sets in detail. Section 3 outlines the modified-principal-curve algorithm that serves as the basis of the tube-fitting procedure. Section 4 gives the detailed tube-fitting algorithm. Section 5 provides a validation study of the algorithm. Section 6 presents the application to SPECT and DT images. Section 7 is a discussion.

2 Motivating Data Sets

2.1 SPECT Colon Imaging

Our SPECT colon data arise from an experiment designed to study the viability of microbicide lubricant for HIV transmission during anal intercourse. SPECT imaging uses a radioactive isotope as the source signal. Projections of emitted photons are collected via gamma cameras mounted on a gantry that are rotated around the subject. Computed tomography algorithms are used to convert the projection images into a three dimensional image. The principal benefit of SPECT imaging is the ability to image biological interactions with the tracer, rather than simply imaging anatomy.

The experiment was designed to simultaneously image the distribution of surrogates for the microbicide lubricant and the viral-mixed semen to assess if the coverage of the lubricant is sufficient. The experiment used an over-the-counter lubricant as a surrogate for the microbicide, which was mixed with a radioactive tracer (TC-sulfur colloid). A radiolabeled surrogate for the viral-mixed semen is being used for extensions of the experiment, though the data considered here contains only the lubricant.

Ten milliliters of the radiolabeled lubricant was injected into the subject's colon, who subsequently underwent receptive anal intercourse. The subject was then imaged using a dual-head VG SPECT imaging system (GE Medical Systems, Waukesha, WI) equipped with a low-end X-Ray computed tomography system (Hawkeye). The image was reconstructed

using an ordered subsets EM algorithm and filtered as provided with the scanner software (GE eNTEGRA workstation, version 1.04). We present analysis of the reconstructed SPECT image of the distributed lubricant.

2.2 Quantification of DTI Tractography

As mentioned above, DTI [3] has two major values as an imaging modality its sensitivity to tissue microstructure [6] and its ability to guide tractography of the major white matter tracts [32]. This is due to DTIs sensitivity to diffusion anisotropy [5] the tendency of water to diffuse in a particular direction, which, in the brain and spinal cord, is often along the course of an axonal tract. By combining analysis of tissue microstructure with tractography, we can limit our focus to one or several tracts with specific functional correlates for example, motion, vision, and language. Within these tracts, we can analyze not only quantities derived from DTI including anisotropy, absolute and directional diffusivity of water, and tract volume but also quantities derived from other MRI sequences that have been coregistered to the DTI [37]. This offers the possibility of a comprehensive, multimodal approach to analyzing the structure-dysfunction relationship in the central nervous system.

To compare tract-specific imaging results across individuals, a normalization procedure is required. There are two general approaches: whole-brain normalization, which involves warping brains to match one another or some canonical atlas, and tract-specific normalization, which focuses specifically on the tract of interest. The former approach is computationally intensive and may require sacrificing optimal registration of the tract of interest to achieve acceptable registration of the whole brain. We have introduced an approximate tract-specific normalization approach that samples tracts in a slice-by-slice manner [37] this approach has yielded promising correlations between tract-specific MRI quantities and clinical disability scores [38, 35]. However, because white matter tracts in the brain do not typically run perpendicular to the cardinal imaging planes (axial, coronal, and sagittal), a parametric approach that accounts for each tracts specific shape and anatomical course would reduce noise and might increase sensitivity for detection of relevant abnormalities. The parameterization would be different for each tract but would ideally be generated by an algorithm that could be applied automatically to any tract.

Details of our MRI acquisition protocol have been described [36]. On a 3-Tesla Philips Intera scanner, we obtained whole-brain DTI images in the axial plane (2.2 x 2.2 x 2.2 mm voxels interpolated on the scanner to 0.8 x 0.8 x 2.2 mm; parallel imaging with a sensitivity-encoding reduction factor of 2; 2 averages; 32 non-coplanar gradient directions with a nominal b-value of 700 s/mm^2; and a scanner average of 5 minimally diffusion-weighted scans with b33 s/mm^2). We coregistered all images to the first minimally diffusion-weighted scan using the Automatic Image Registration (AIR) algorithm [42] with a 6-parameter rigid-

body transformation, and we corrected the gradient directions for the rotational component of the transformation. We then estimated the diffusion tensor in the standard fashion [2], diagonalized the tensor to obtain its eigenvalues and eigenvectors, and calculated maps of anisotropy and diffusivity. These analyses were performed in DtiStudio [25], as well as with custom software purpose-written in Matlab (The Mathworks, Natick, MA).

We used the DTI datasets to obtain 3D reconstructions of the corticospinal tracts using the fiber association by continuous tractography method [32, 33, 34]. We reconstructed the tracts using fractional anisotropy and the principal eigenvector of the diffusion tensor. We used every voxel in the brain with fractional anisotropy > 0.13 as a potential starting point for tractography and halted individual tracts once a voxel with fractional anisotropy < 0.13 was encountered or when the reconstructed tract turned at an angle steeper than 40 degrees from one voxel to the next. We chose multiple restrictive regions of interest to limit the reconstructed corticospinal tracts to their known anatomical course; these regions of interest have been described previously [36] and were drawn in the rostral medulla, rostral pons, and subcortical white matter. We manually excluded the rare spurious fibers that were included in this reconstructions but that clearly fell outside the major portion of the corticospinal tract.

3 Modified Principal Curve Algorithm

To construct our tube, we first need a fitted curve that acts as a centerline for our data. Statistical analysis for three-dimensional curve-fitting and centerline calculation has received little attention in the statistical community. (We emphasize the difference between fitting non-parametric functions, a process well studied in the smoothing literature and non-parametric curve-fitting.) However, curve fitting has received a great deal of attention in the computer-vision and medical-image-processing literature. Relevant literature exists in the field of virtual colonoscopy and localization of polyps (see for example, [30, 39, 40, 23, 12, 18]). These approaches generally require connected curves, and are not directly applicable to the range of problems that we consider, which may have interrupted structures or may be a voxel-wise reduction of connected-curve data. Also relevant from the image processing literature is the tremendous volume of work on Bezier curves and B-splines (see the review in [14]). Though we have explored Bezier approaches, we do not utilize them because of the large amount of user input required to appropriately place knots.

Another popular collection of techniques treats the points of the image as a networked graph and uses combinatorial algorithms to find globally optimal paths [15, 16, 11, 13, 7, 8]. Dijkstra's algorithm is often used to find solutions [19].

Perhaps the most statistical approach that we have encountered relies on the use of

principal curves [20, 22, 21]. These generalizations of principal components find a curve achieving a local minimum for the sum of the orthogonal distances of the points onto the curve. This approach is useful in statistical methods of image analysis [1, 9]. Numerous modifications of the principal-curve idea have been published (see the discussion in [26]). In addition there are related stochastic search algorithms for finding centerlines, as considered in [17]. Our approach in this manuscript utilizes the modified-principal-curve algorithm presented in [9]. This procedure can accommodate interrupted curves, constrained points and can fit low variation curves that the original algorithm could not. We briefly describe the procedure below.

To start, we need a method for representing a curve. The study of differential geometry has revealed several equivalent methods for representing real-valued curves [14, 41, 27], including implicit representations, the set of points $\{(x, y, z) \in \mathbb{R}^3 \mid F(x, y, z) = G(x, y, z) = 0\}$, for surfaces F and G and parametric representations. We focus entirely on parametric representations, of which implicit representations are a special case [27]. An allowable parametric representation sets $f(t) = \{f^x(t), f^y(t), f^Z(t)\} : [a, b] \to \mathbb{R}^3$, where $[a, b]$ is an interval in $\mathbb{R}$ and at least one of $df^x(t)/dt$, $df^y(t)/dt$ or $df^z(t)/dt$ is non zero. We assume the constraint $t \in [0, 1]$ for identifiability. However, this assumption alone does not uniquely specify a curve. Indeed, if the curve is considered to be the location of a particle at time t, then the same curve can arise from particles following the same path at different rates.

Given this parametric curve representation, we view the process of fitting a curve through three-dimensional coordinates as inherently a missing-data problem. That is, let $\{(X_i, Y_i, Z_i)\}_{i=1}^n$ be a collection of realized values for the coordinate functions. The process of finding a curve through them largely amounts to finding a reasonable estimates for the missing data $\{t_i\}_{i=1}^n$. However, estimating the missing time data, $\{t_i\}_{i=1}^n$, is a difficult process.

Here $\{X_i, Y_i, Z_i\}_{i=1}^n$ are lattice values of points in the image surviving a thresholding procedure for noise reduction. In addition, to improve computing times, we often work with a subset of the points, sampled uniformly, as the curve is often well defined with much fewer points. This is not necessary for the DTI tractography example, but it speeds up computing substantially at no loss of quality-of-fit for the SPECT colon data.

The basic principal curve algorithm is a blocked-maximization algorithm that iterates between two steps: calculating the time points and fitting curves to the coordinate data: $\{(X_i, t_i)\}$, $\{(Y_i, t_i)\}$, $\{(Z_i, t_i)\}$. Suppose that an initial estimate of f, say $\tilde{f}$, is given. Then, the t_i are calculated as:

$$t_i = \text{argmin}_{t \in [0,1]} \mid\mid \tilde{f}(t) - (X_i, Y_i, Z_i) \mid\mid . \tag{1}$$

The estimate, $\tilde{f}$ is then updated by fitting a smoother between the $\{X_i\}$ and the $\{t_i\}$, the $\{Y_i\}$ and the $\{t_i\}$, and the $\{Z_i\}$ and the $\{t_i\}$, separately. We use cubic regression splines for this

portion of the algorithm, though other smoothers could be used. However, regression splines allow for easily calculated derivatives on the coordinate function. The steps of updating the $\{t_i\}$ and $\tilde{f}$ are iterated until the change in $\tilde{f}$ between successive steps is sufficiently small.

Several modifications to the principal curve algorithm outlined above are proposed by [9]. First, the modified-principal-curve approach allows for user-specified end points. Second, it molds the curve by gradually increasing the degrees of freedom in the regression splines, so that gross features of the curve are captured before fitting finer details. This provides for better fits to complex curves. Third, the modified approach incorporates image intensities to adjust the emphasis placed on high- and low-intensity points in the curve-fitting. Fourth, a grid search is used to perform the minimization in the second step of the algorithm to speed up convergence. Finally, the stopping criterion is based on relative change in mean square error. As described originally, the modified-principal-curve-fitting algorithm also allows for user-specified interior points, though constrained interior points did not lead to better fits in our applications.

The modified principal curve algorithm is semiautomated, requiring user defined end-points and, in some cases, adjustment of the final number of degrees of freedom used in the regression splines. This algorithm provides a differentiable curve that acts as a centerline through the data. We emphasize, however, that the algorithm used to construct the center-line curve functions independently of the algorithm used to fit the tube. So, for example, other less automated procedures, such as using B-splines with user-selected knot points, could be used for this step.

4 Tube-Fitting Algorithm

4.1 Background

In the previous section, we outlined the curve-fitting algorithm, which provides a centerline to serve as the basis for our tube-fitting algorithm. Briefly, our tube-fitting algorithm consists of the following steps: *i*. select a starting point on the fitted curve; *ii*. identify nearby image points; *iii*. project nearby image points onto the plane orthogonal to the fitted curve at the starting point; *iv*. fit a bivariate normal distribution to the (now two dimensional) points in the orthogonal plane; *v*. use a level set of the bivariate normal to define the tube at the chosen starting point on the fitted curve. We repeat these steps as often as desired. We now examine these steps in greater detail.

Let $f(t)$ be a centerline curve at t. We define the tube around the curve f at t to be an ellipse in the plane orthonormal to $f(t)$ such that $f(t) \in G(t)$. To elaborate, the function $g(t) = \nabla f(t)/||f(t)||$, where ∇ is the gradient function and $||\cdot||$ is the Eu-

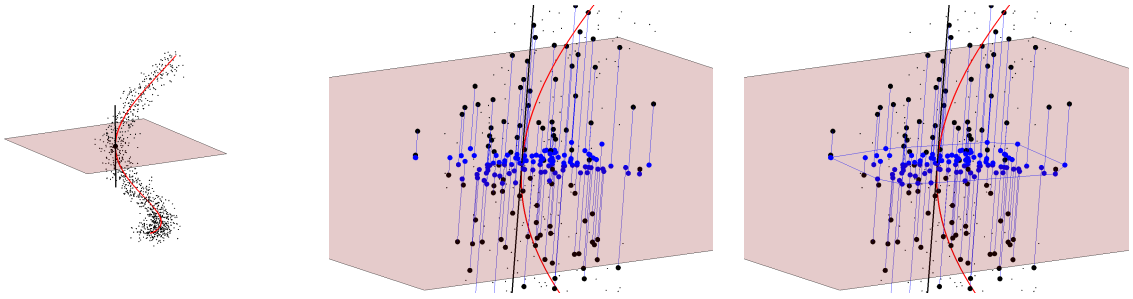

Figure 3: The first plot shows an example fitted curve, tangent line at a point t_0 and the normal plane to the tangent line at the point. The second plot zooms in on the normal plane and displays the local orthogonal of points onto the normal plane. The third plot shows the convex hull fit to the data projected onto the normal plane.

clidean norm, gives the unit tangent vector to the curve f at point t. Then the set $G'(t) = \{w \in \mathbb{R}^3 \mid g(t)'\{w - f(t)\} = 0\}$ is the normal plane to $g(t)$ passing through the point $f(t)$, and, regarding the plane $G'(t)$ as the cartesian plane for convenience, we take $G(t)$ to be a set of the form $\{x, y \in \mathbb{R} \mid Ax^2 + Bxy + Cy^2 + Dx + Ey \geq F\}$. The object of interest is then the collection of points $\{d(t), G(t)\}$, where

$$d(t) = \int_0^T \sqrt{\left\{\frac{d}{dt}\hat{f}^x(t)\right\}^2 + \left\{\frac{d}{dt}\hat{f}^y(t)\right\}^2 + \left\{\frac{d}{dt}\hat{f}^z(t)\right\}^2}\, dt, \tag{2}$$

is the distance along the curve and $G(t)$ is the ellipsoidal cross section of the tube at time t.

By choosing to restrict the shape of the compact set $G(t)$ to an ellipse, we necessarily restrict our algorithm to those structures with ellipsoidal cross sections. A more general definition is possible: define tube around the curve f at t to be a compact collection of points $G(t)$ in the plane orthonormal to $f(t)$ so that $f(t) \in G(t)$. However, such a generalization was not needed in either of our applications.

In this section, we are concerned with the construction of an estimate of $G(t)$ from imaging data. This task is complicated by the limitations of our imaging modalities. For the SPECT image, the location of the radioactive tracer is measured with error, meaning that some of the image points $\{X_i, Y_i, Z_i\}_{i=1}^n$ do not lie inside the object of interest $\{d(t), G(t)\}$. For the specific DTI tract in question, there is less noise. Below we outline a procedure for estimating $G(t)$ accounting for noise in the image.

4.2 The Algorithm

We first overview our approach before discussing details. For any point on the curve with associated latent point, we project nearby points onto the orthogonal plane described above.

We then rotate the projected points to a two-dimensional coordinate system. In that coordinate system we then estimate $G(t)$. Finally, we rotate the results back to the original coordinate system.

We begin with a point on the fitted curve, say P_0, and corresponding latent point $t_0 \in [0, 1]$. The point, P_0, does not have to be, and will typically not be, the projection of an observed point onto the curve, nor does it have to lie on the lattice defined by the imaging coordinates. We then select "nearby" points to estimate $G(t_0)$. Let $\{P_i\}_{i=1}^n$ be the points used in the curve-fitting procedure, $P_i = (X_i, Y_i, Z_i)$. These are the image points surviving the thresholding and sampling procedure, and depending on the imaging procedure may be a noisy version of the underlying structure. We define nearby points from our collection of coordinates based on the proximity of their associated $\{t_i\}_{i=1}^n$ to t_0, where t_i is the value of the latent variable t such that minimizes the distance between $f(t)$ and P_i. Specifically, we project every observed point onto the final estimated curve and calculate its associated t_i via (1). If the modified principal-curve algorithm is used to fit the centerline $f(t)$, the collection $\{t_i\}_{i=1}^n$ is calculated as a byproduct of the final steps of the algorithm. Thus proximity in terms of t is a measure of the nearness of the projections of the P_i onto $f(t)$. The neighborhood of points to be projected onto the orthogonal plane is defined as

$$\{P_{i_j}\} = \{P_i \mid |t_0 - t_i| < t_r\}. \tag{3}$$

where t_r is the range of the time window. We use the points $\{P_{i_j}\}$ to estimate $G(t_0)$.

Intuitively, measuring proximity in terms of the latent variable t allows us to select the nearest neighbors of P_0, defining "nearness" as distance on the curve $f(t)$. This strategy for defining a neighborhood of observed points around P_0 has major benefits over competing methods, such as using a standard neighborhood based on the Euclidean distance between the observed points and P_0. Specifically, for P_0 near high curvature in the fitted curve, observed points can then overly contribute to the fitted tube at multiple locations. See Figure 4 for a two-dimensional illustration. The blue points in the left panel are in the Euclidean neighborhood structure of P_0, and include points that lack face validity for contributing to the estimate of $G(t_0)$. In contrast, the right panel shows that the neighborhood defined as (3) has much better behavior. We note that in areas of low curvature, our method for choosing $\{P_{i_j}\}$ coincides with the method using Euclidean distance.

For our applications, we have found that choosing $.05 \leq t_r \leq .2$, depending on the total number of image points, includes enough locations to estimate the tube's shape without employing locations that are very distant (recall that our curve-fitting algorithm specifies $0 \leq t \leq 1$). However, we emphasize that truncating points in this way is done primarily for computational purposes. In estimating $G(t_0)$ we weight points (see below) by their distance in t, so that further away points contribute less to the current estimate of $G(t_0)$. Hence all

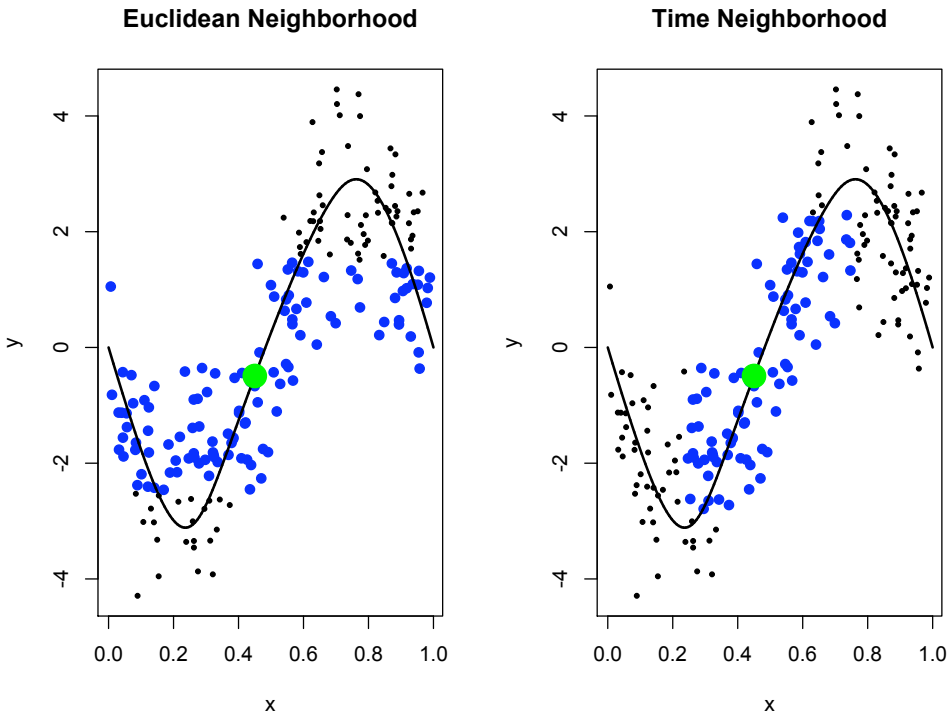

Figure 4: Comparison of methods for defining the neighborhood around P_0. In the left, the neighborhood is defined as though points with a Euclidean distance from P_0 less than 2; in the right, we use a t-window with $t_r = .2$. In both, the point P_0 is shown in green, and the points in the set $\{P_{i_j}\}$ are shown in blue.

points could be used; however, as for common kernel-smoothing procedures, the above filter can eliminate unnecessary computing for points with extremely low weights.

Next, we project the image points close to P_0 into the plane orthogonal to the fitted curve at time t_0. We define the orthogonal plane as above, using the normalized derivative of the fitted curve as a normal vector, say $\mathbf{n}$, and the point P_0 as a point in the plane, so that the plane consists of all points $\mathbf{r}$ satisfying $\mathbf{n} \cdot (P_0 - \mathbf{r}) = 0$, where $\cdot$ represents a vector dot product. An intuitive approach is to project each point onto the orthogonal plane in the standard sense, taking the projection of P_i to be that point in the plane with minimum Euclidean distance. However, we find that in areas of high curvature this approach gives poorly performing projections. This stems from the fact that these projections skew toward the interior of the curve, rather than remaining centered around the fitted curve. Figure 5 illustrates this point in two dimensions, showing standard projections and projections using our novel approach for a sample of points in a time window. A three dimensional illustration appears in Section 5.

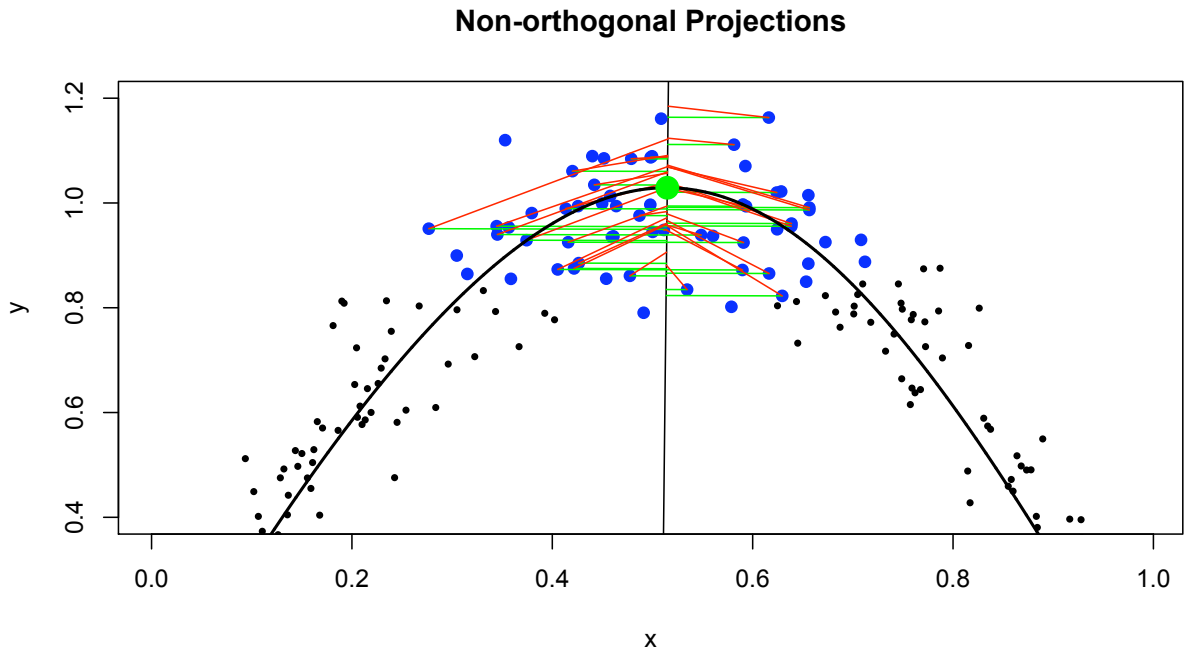

Figure 5: Two projection methods shown in an area of high curvature: in green, standard projections minimizing the distance between a point and the line; in red, our modified projections that maintain a point's distance from the fitted curve. Again, P_0 is highlighted in green and the $\{P_{i_j}\}$ are shown in blue.

Instead, we use a method that maintains a point's distance and direction from the centerline in its projection into the orthogonal plane. Conceptually, our method stretches the space containing the centerline $f(t)$ and the points P_i around it so that $f(t)$ is linear. Standard orthogonal projections are used in the transformed space to project the $\{P_{i_j}\}$ into the plane normal to $f(t_0)$. This approach avoids the issues described above that arise when standard projections are used in the observed, non-transformed space. Our method is carried out in the following steps. Let $\tilde{P}_i = \tilde{f}(t_i)$ be the projection of the point P_i onto the fitted curve. Consider the plane orthogonal to $\tilde{f}(t_i)$. By definition, both P_i and $\tilde{P}_i$ lie in this plane. We rotate and translate this plane so that: the plane is parallel to the axial plane of the image (i.e. is horizontal), the height of the plane is $Z = 0$, and the point $\tilde{P}_i$ is at the origin. To accomplish this, we let $\mathbf{n} = g(t_i) = \nabla f(t_i)/||f(t_i)||$ and express the plane orthogonal to $\tilde{P}_i$, $\mathbf{n} \cdot (\tilde{P}_i - \mathbf{r}) = 0$, as points $\mathbf{r} = \{x, y, z\}$ satisfying $n_x x + n_y y + n_z z + d = 0$, where $d = -\mathbf{n} \cdot \tilde{P}_i$.

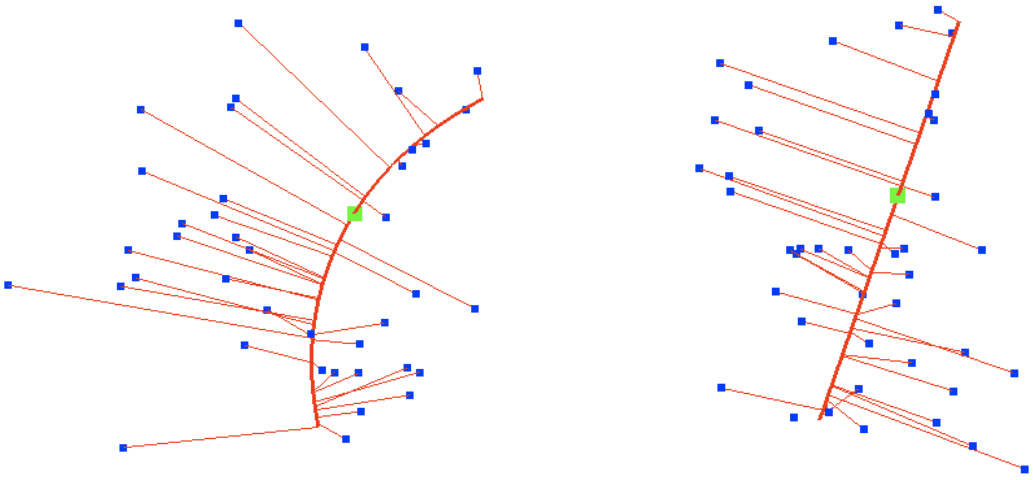

Figure 6: Local linearization in an area of high curvature. Keeping with previous figures, P_0 is shown in green and $\{P_{i_j}\}$ are shown in blue.

Next, we sequentially apply the rotation matrices

$$\mathbf{Q}_x = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta_x) & \sin(\theta_x) \\ 0 & \sin(\theta_x) & -\cos(\theta_x) \end{bmatrix} \text{ and } \mathbf{Q}_y = \begin{bmatrix} \cos(\theta_y) & 0 & \sin(\theta_y) \\ 0 & 1 & 0 \\ \sin(\theta_y) & 0 & -\cos(\theta_y) \end{bmatrix} \tag{4}$$

where $\theta_x = \arctan(\frac{-n_y}{n_z})$ and $\theta_y = \frac{n_x}{n_z}\cos(\theta_x)$ to the points $\left[P_i, \tilde{P}_i\right]$ in the orthogonal plane first with respect to the X-axis and then the Y. The resulting rotated plane is parallel to the XY-plane. Finally, we subtract the Z coordinate from this plane and translate so that $\tilde{P}_i$ is at the origin. Let P_i' be this affine transformation of the observed point P_i. Applying these steps to P_0 projects the plane orthogonal to the fitted curve at t_0 into the XY plane and places P_0' at the origin.

Notice that P_i' has Z coordinate 0 and its distance from the origin is equal to the distance between P_i and its projection onto the fitted curve, $\tilde{P}_i$. We perform this process for all points in the neighborhood of P_0, $\{P_{i_j}\}$, to obtain a set of rotated and translated points, $\{P_{i_j}'\}$. These points in two dimensional space have distance and direction from the origin that is the same as their distance and direction from the fitted centerline. In effect, we have locally linearized our fitted curve and collapsed the locations in the current t-window into a single plane. This is demonstrated in Figure 6.

We use the $\{P_{i_j}'\}$ to estimate the tube $G(t_0)$. Hence, our task is to choose a suitable shape surrounding these points to serve as our estimate of the tube in the plane orthogonal to the curve at P_0. This choice necessarily restricts the class of estimable structures to those with the specified cross sectional shape. Our first inclination was to use the convex hull

– the smallest closed set containing the points – because it is flexible and comparatively unrestrictive (this approach is analogous to using $\hat{a} = X_{(1)}$ and $\hat{b} = X_{(n)}$ to estimate the endpoints of the support of s uniform distribution). However, such estimates do not account for any noise in the measurements inherent in some imaging techniques and would only be acceptable for very high resolution images without noise. Moreover, convex hulls can not accommodate tube cross sections that are not closed, nor can a tube constructed from convex hulls change smoothly along the centerline $f(t)$. A circle centered at the origin, i.e. the level set of a bivariate normal distribution with no correlation, was too restrictive for the shapes seen in practice. Figure 7 shows an example set $\{P'_{i_j}\}$ for which a circle is a poor shape for constructing $G(t_0)$. Therefore, we use a level set of a bivariate normal distribution,

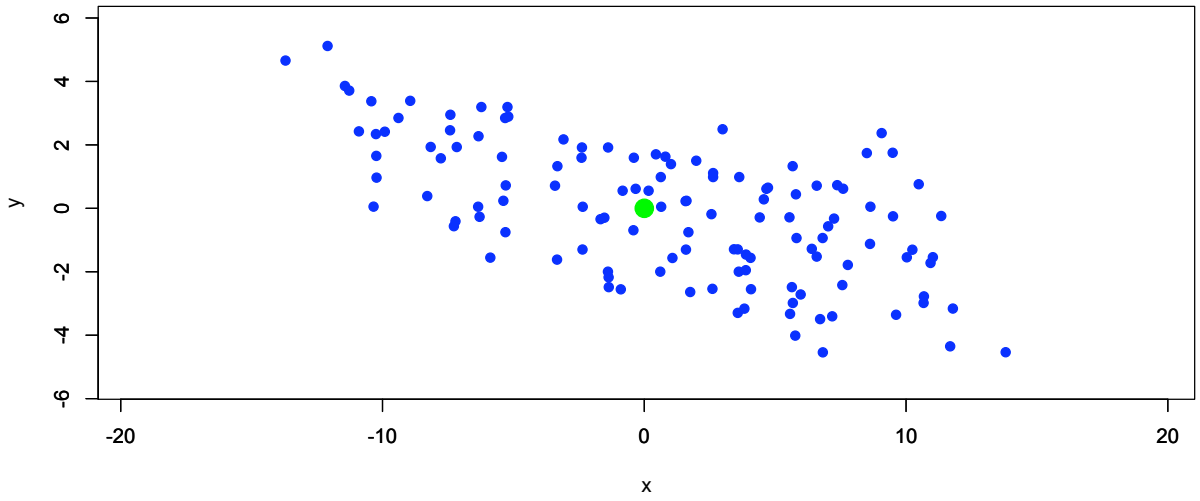

Figure 7: $\{P_{i_j}\}$ used to estimate the tube $G(t_0)$. This example comes from the descending colon in our SPECT application. Here, P'_0 is shown in green and the set $\{P'_{i_j}\}$ used to estimate $G(t_0)$ is shown in blue.

an ellipse, to define $G(t_0)$. This choice coincides with observed sets $\{P_{i_j}\}$ as well as our scientific collaborators' knowledge of the anatomical structures in our motivating data sets; for other applications, a different choice for the shape of $G(t_0)$ may be needed. Let $\{w_{i_j}\}$ be the collection of normalized weights. Then the estimated bivariate normal has mean and variance

$$\tilde{\mu} = \sum_{j=1}^{J} w_{i_j} P'_{i_j} \text{ and } \tilde{\Sigma} = \sum_{j=1}^{J} w_{i_j} (P'_{i_j} - \tilde{\mu})(P'_{i_j} - \tilde{\mu})^T \tag{5}$$

where $1 \leq j \leq J$ indexes i_j.

We choose weights $\{w_{i_j}\}$ so that more distant points have a smaller impact on the density estimation than nearer points. Specifically, points further away from P_0, as defined by the

distance from t_{i_j} to t_0, influence the fit less. We use a cosine-transformed distance as the weight:

$$w_{i_j} = \frac{\cos\left[\frac{(t_{i_j}-t_0)\pi}{r}\right]+1}{\sum_{j=1}^{J}\cos\left[\frac{(t_{i_j}-t_0)\pi}{r}\right]+1} \tag{6}$$

with r the half width of the t-window. This weighting scheme has the desired effect of emphasizing nearby points while smoothly decreasing to zero for more distant points. An example of the weighting scheme is shown in Figure 8. We note that other weighting schemes that decrease to zero at the tails, specifically kernel weighting schemes, give very similar results. Schemes that do not decrease to zero, like one that gives uniform weight to all points in the t window, are less desirable because the resulting tube is not necessarily smooth.

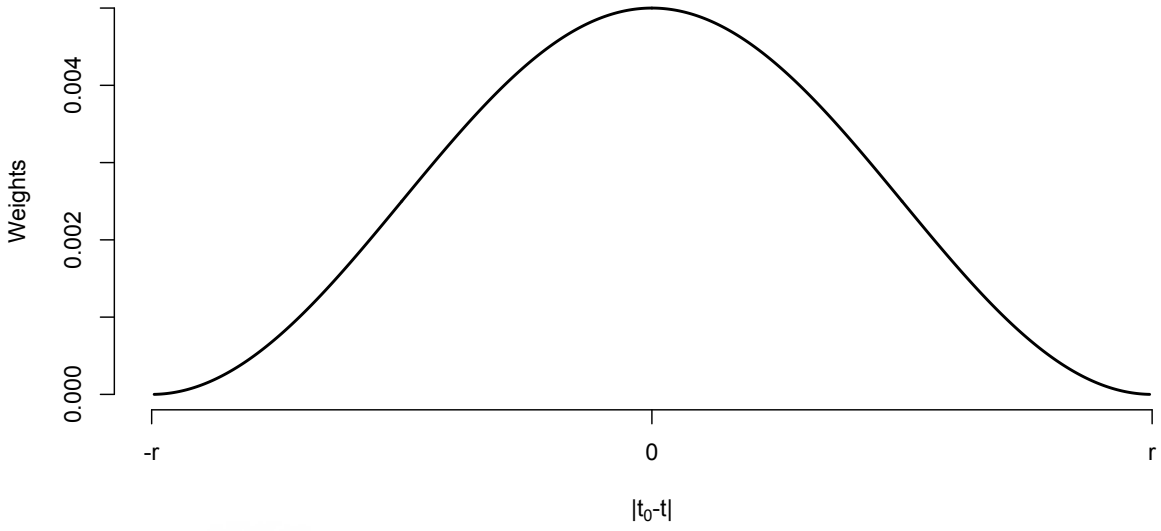

Figure 8: Weights for $t_0 = .1875$ and $r = .1$.

The tube at point t_0 in the rotated space, $G'(t_0)$ say, is then estimated by the level set:

$$\hat{G}'(t_0) = \left\{ P \in \mathbb{R}^2 \;\middle|\; |2\pi\tilde{\Sigma}|^{-1/2}\exp\{-(P-\tilde{\mu})'\tilde{\Sigma}^{-1}(P-\tilde{\mu})\}/2\} > l \right\} \tag{7}$$

where l is chosen so that

$$\int_{P\in\hat{G}'(t_0)} |2\pi\tilde{\Sigma}|^{-1/2}\exp\{-(P-\tilde{\mu})'\tilde{\Sigma}^{-1}(P-\tilde{\mu})\}/2\}dP = 1-\alpha. \tag{8}$$

Typically the choice of α will be context-specific, depending on the shape of the ellipse $G(t_0)$ and the measurement error variance (if any exists). In Section 4.3, we explore the effect of α on the resulting estimate $\hat{G}(t_0)$.

As a final step, we project the ellipse (and the P_j', if necessary) into the plane orthogonal to the fitted curve at t_0. This transformation, from $\hat{G}'(t_0)$ to $\hat{G}(t_0)$, is achieved by inverting the steps that take P_0 to the origin of the XY plane. Specifically, we translate the ellipse horizontally by the amounts that place P_0 at the origin and vertically by adding in the Z coordinate of the rotated plane. Finally, we multiply by $\mathbf{Q}_y^{-1}$ and $\mathbf{Q}_x^{-1}$, where the matrices are defined in equation 4 with $\mathbf{n} = g(t_0)$, to rotate the ellipse back into the plane orthogonal to $f(t)$ at t_0.

In some cases, the image data contain additional information regarding the points $\{P_i\}$. For instance, the SPECT data incorporate concentrations $\{C_i\}$ of the radioactive tracer used to generate the images. In these situations, accurate cross-sectional values of the additional image parameters can be calculated from the $\{C_{i_j}\}$.

Though starting points can be chosen arbitrarily, we generally choose points equally spaced and a relatively small distance apart on the fitted curve. The tube-fitting algorithm proceeds iteratively along the fitted curve, repeating in full the steps detailed above for each of the selected points. This results in a fitted tube that gives an adequate level of detail about the shape of the imaged object without being computationally prohibitive.

4.3 Choosing α

As previously mentioned, the underlying modality and noise characteristics of the image impacts how one selects the cross-sectional ellipse covering the structure. In other words, how to select α in the level set of the bivariate normal. Our two examples highlight the difficulty in obtaining a universal rule. The SPECT image is clearly very noisy, as is required by the underlying Poisson decay of the tracer and the other sources of noise imposed during image acquisition and reconstruction. The DTI tract, on the other hand, *appears* nearly noise free. However, there is noise in the underlying DTI image and potential noise and bias from the tractography algorithm. However, without repeat scans, it is impossible to characterize this variability in the DTI image. Therefore, we seek the most accurate representation of the tract image, acknowledging that there are sources of noise and bias that are not represented or quantified. Thus, the choice of α differs greatly in these two applications.

To elaborate on this choice, we have two competing goals: *i.* to maximize the coverage of the true ellipse $G(t_0)$ by our estimated ellipse $\hat{G}(t_0)$; and *ii.* to avoid choosing $\hat{G}(t_0)$ excessively large through the inclusion of points not in $G(t_0)$. To characterize these goals, we examine the quantities

$$TP = \frac{A\{G(t_0) \cap \hat{G}(t_0)\}}{A\{G(t_0)\}} \text{ and } FP = \frac{A\{G(t_0)^c \cap \hat{G}(t_0)\}}{A\{G(t_0)\}} \tag{9}$$

where $A(\cdot)$ gives the area of the designated shape. These quantities, TP and FP, can be

thought of respectively as the true and false positive rates normalized to the area of $G(t_0)$, so that $0 \leq TP \leq 1$ and $0 \leq FP$. These quantities are analogs of the true and false positive rates from the analysis of classification data.

As discussed above, because it depends highly on the distribution of measurement errors and other factors, the choice of α will be context-specific. We therefore advise a validation study tailored to application at hand, if such a study is possible. Indeed, in Section 5 we present validation data both to confirm the tube-fitting algorithm and to aid in selecting α for our SPECT imaging application. However, a study of this kind is not always possible, so here we present a brief simulation designed to provide a basis for evaluating the interplay between the choice of α. and noise levels in the image.

We posit an underlying collection of true points from an ellipse and add spherical noise. The goal is to estimate the ellipse. Thus two sources of variation are considered, the sampling of underlying true points and noise. The simulation consisted of the following steps:

1. Points are sampled uniformly from the interior of the underlying ellipse G;
2. Normal errors with variance matrix $\Sigma = \sigma^2 I_{2x2}$ are added to the sampled points to give observed points;
3. From the observed points, a bivariate normal is estimated and used to construct $\hat{G}$ for a range of α values;
4. TP and FP are calculated for each of the α values.

These steps are iterated 100 times each for a variety of ellipse shapes and measurement error variances. Figure 9 shows some of the ellipses G used in our simulation, as well as a representative collection of sample points and observed points.

We found that two main factors should contribute to the choice of α: the measurement error variance and the eccentricity of the ellipse. The eccentricity of an ellipse with semi-major and -minor axes A and B, respectively, is defined as $e = \sqrt{\frac{A^2-B^2}{A^2}}$. For low measurement error variance, say $\sigma = .1 * B$, the eccentricity of the ellipse is irrelevant: taking $\alpha = .12$ gives $TP = .95$ and $FP = .1$. For large measurement error variance, the eccentricity of the ellipse is quite important. Indeed, for an ellipse with $A = B$ and $\sigma = B$, $\alpha = .62$ yields $TP = .95$ and $FP = .2$, while for an ellipse with $A = 4B$ and $\sigma = B$, the same choice of α gives $TP = .55$ and $FP = .05$. Figure 10 shows the results of our simulation study. We present these results in two ways. First, keeping the eccentricity of the ellipse constant, we examine the effect of varying σ on TP and FP. Second, we keep σ constant and allow the shape of the ellipse to vary. Finally, we note that the results presented here hold for ellipses in other scales; that is, TP and FP as a function of α are the same for $A = B = 10$ and $\sigma = 5$ and for $A = B = 100$ and $\sigma = 50$.

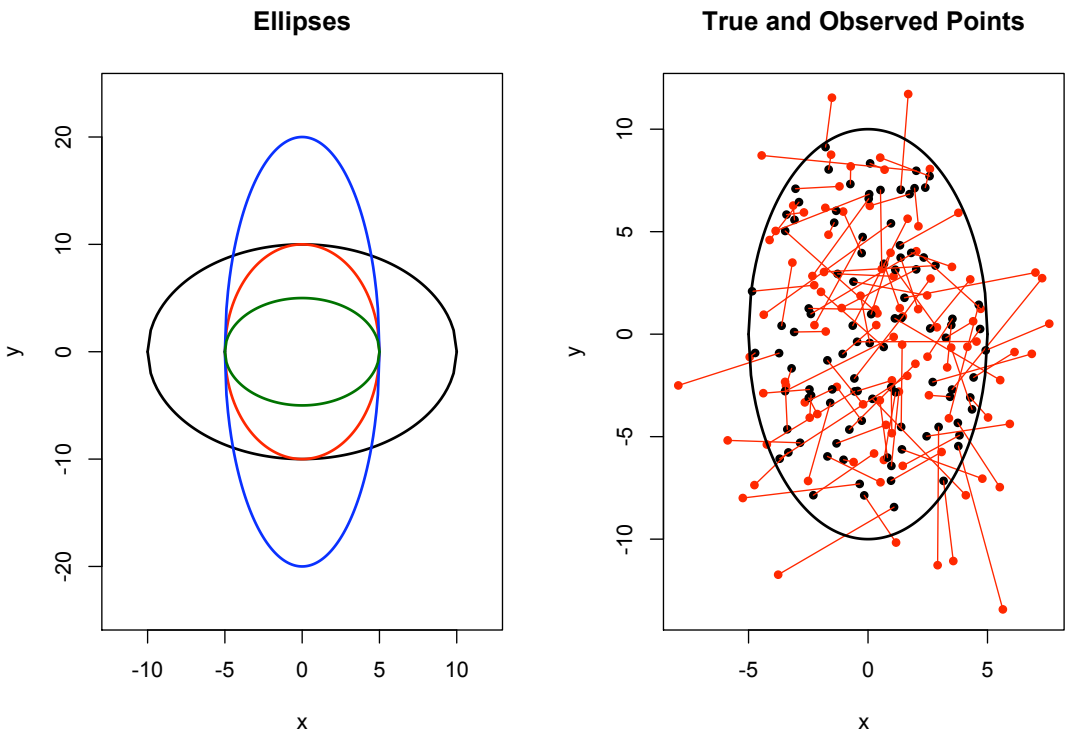

Figure 9: The left panel shows several of the true ellipses used in our simulation. On the right, we show the points sampled uniformly from an ellipse in black and the points observed with measurement error in red.

5 Validation

Before applying our tube-fitting algorithm to image data, we pursued a brief validation study using mathematical phantoms. A mathematical phantom is simply a shape, created digitally, which is then passed through a computational model of the imaging process. Accurate computational models of diffusion imaging are not available due to the inherent complexity of nuclear spin systems and water diffusion. On the other hand, very accurate models for some transmission and emission imaging, such as X-rays, planar scintigraphy, SPECT, PET (positron emission tomography) and CT (X-ray computed tomography) are available. In these cases the imaging process is perhaps simpler to model than in MRI and highly accurate models of the imaging physics have been created. To generate the SPECT images for validation, we used system models implemented in the Division of Medical Imaging Physics in the Department of Radiology at the Johns Hopkins University. In this method, the projection data of the phantom were obtain using analytical projector that models all of the important components of the images physics, including photon interactions inside both

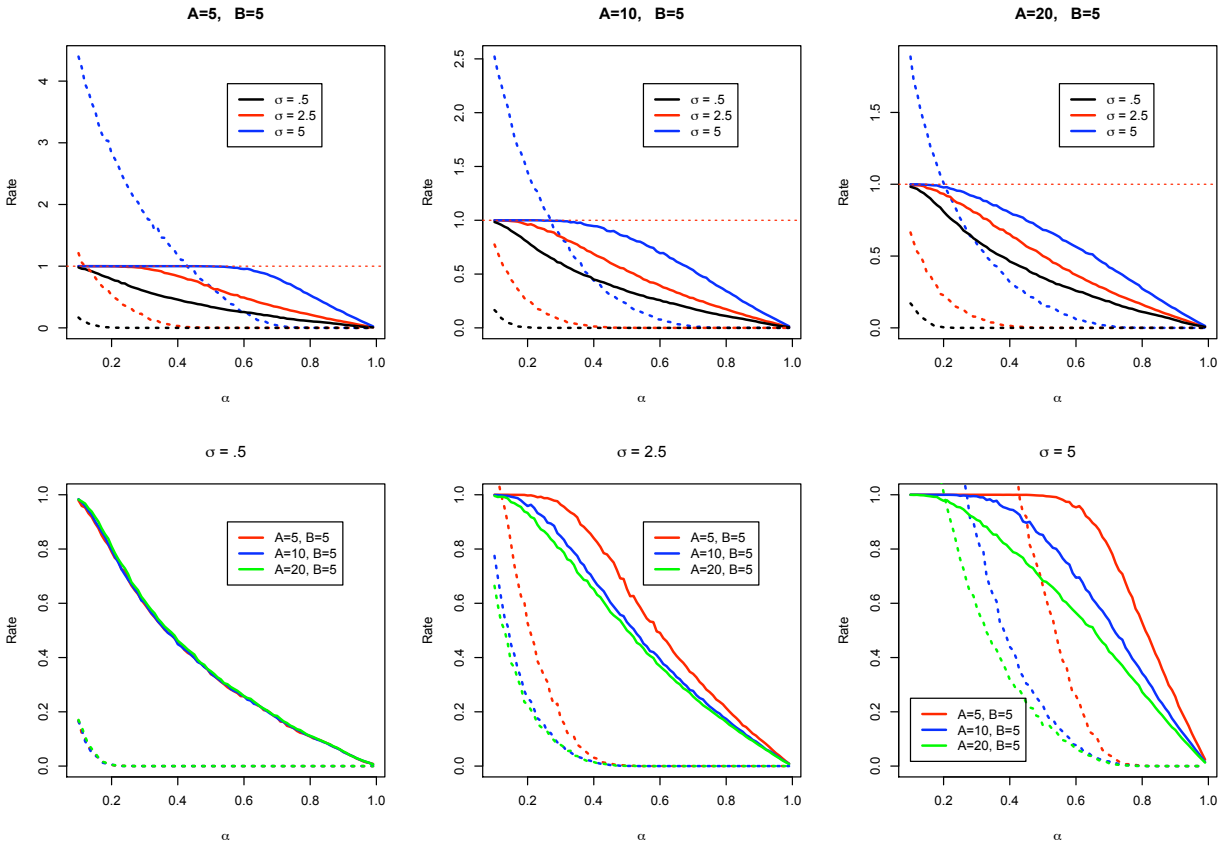

Figure 10: Results of the simulation for choosing α. For all panels, the solid line represents TP and the dashed line FP. In the top row, each panel fixes the shape of G and varies σ, while in the bottom row σ is fixed for each panel and the shape of the ellipse changes.

the phantom and the detector system. The 3D SPECT images were then reconstructed from the projection data using an iterative statistical algorithm.

Our phantom is a three-dimensional coil of fixed diameter; we also investigated coils with monotonically increasing or decreasing diameters, with very similar results. The phantom was projected using the analytical projector described above with effects of attenuation and detector resolution blur. Poisson noise with data-derived means was also added to the projection data. Several noise levels were investigated to mimic images taken at baseline, three hours, 10 hours and 24 hours after introduction of the tracer. (Later images are noisier than earlier images.) The SPECT images were reconstructed using OS-EM algorithm [24].

Each image had an imaging space of over 20,000 non-background voxels. To speed up curve-fitting and tube-fitting algorithms, we randomly sampled 1,000 locations among these, separately for each validation image. The algorithms were run on the sampled locations, and the resulting fitted tube was compared to the true underlying anatomical structure.

Additionally, we varied the level set of the bivariate normal used to define the tube at each point along the fitted curve, which is equivalent to varying the choice of α. Particularly, we were interested in the proportion of points included in the tube that were indeed in the anatomical structure (true positives), the proportion of points included in the tube that are not in the structure (false positives), and the effect of α – the level of the bivariate normal used in constructing $\hat{G}'(t_0)$ – on these rates. Our goals are to maximize the true positive rate while minimizing the false positive rate; that is, we want our tube to be large enough to capture the structure but not so large as to include extraneous points.

There are two important differences between our current validation study and the simulation study in Section 4.3. The first is that our current study is tailored to the SPECT application, and is therefore preferable for selecting α in this setting. Second, the true and false positive rates discussed here are taken over the entire fitted tube, rather than at a single fitted ellipse as in our previous simulations.

It is worth noting that the fitted tube captured the shape of the anatomical structure quite well, even in noisier images. As seen in the left panel of Figure 11, the false positives and false negatives occurred primarily in a thin layer on the outer surface of the anatomical structure. These errors are at least in part due to variations in the fitted curve and tube

Figure 11: False positives (black) and false negatives (red) in tubes fitted to a SPECT scan at 10 hours. Tubes were constructed using our local linearization method of projecting (left) and standard orthogonal projections (right).

induced by randomly sampling 1,000 points from the more that 20,000 non-background voxels rather than to a general deficiency in the tube-fitting algorithm. All other errors occurred at the endpoints of the tube, due to the placement of the user specified endpoints in the

the curve-fitting algorithm. We include in the right panel of Figure 11 a similar image for a tube constructed using standard orthogonal projections. Here, the false positives occur almost exclusively on the interior side of the structure and the false negatives occur almost exclusively on the exterior side; this is consistent with our concerns above, namely that orthogonal projections skew toward the interior of the fitted curve. These observations reinforce our projection method and give us confidence in the ability of the tube-fitting algorithm to accurately reproduce an imaged structure.

Figure 12 shows the true and false positive rates as a function of $1 - \alpha$. From these graphs we see that for a fixed α, noisier images contain greater rates of both true and false positives; the noise in the image leads to a wider fitted tube. We also note that the α level set used to construct the tube does not correspond to the true positive rate. Hence, it must be viewed as a tuning parameter used to balance the true and false positive rates. Based on these figures, we select $.8 \leq 1 - \alpha \leq .9$, depending on the amount of noise in the image. For noiseless images (shown in Figure 12 as "Truth"), choosing $1 - \alpha = .9$ gives a true positive rate $\approx .95$ and a false positive rate $\approx .15$; for our noisiest image (24 hours after baseline), choosing $1 - \alpha = .8$ gives similar rates. We note that our selection is based on visual inspection rather than a well-defined optimizing procedure.

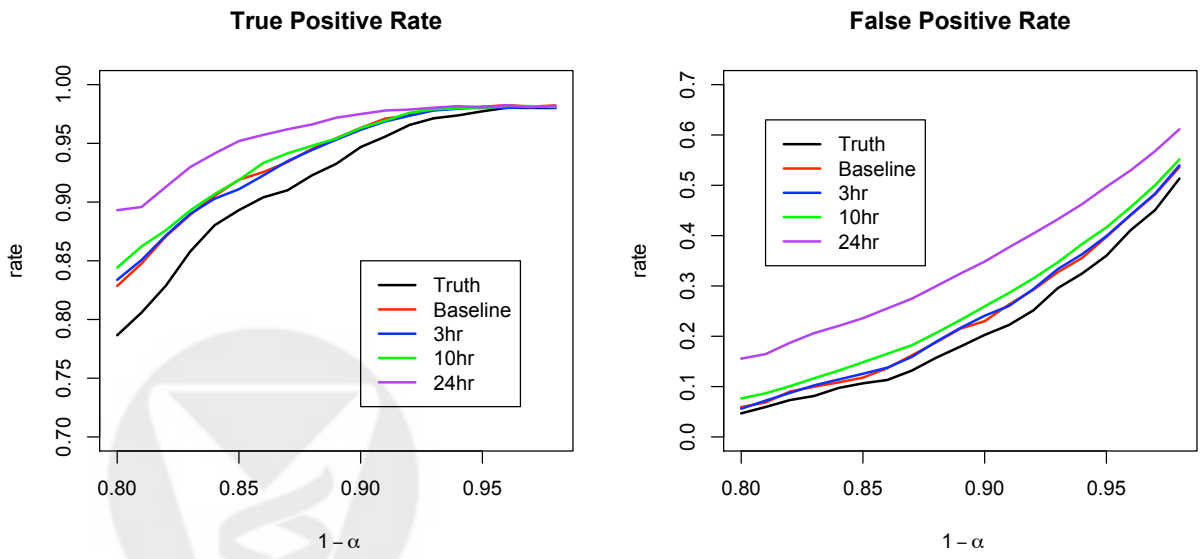

Figure 12: True positive and false negative rates for each of the five validation images as a function of the level used to determine the fitted tube

6 Applications

6.1 SPECT Images

We first consider the SPECT colon image, which was taken shortly after introduction of the tracer. We first filtered the image with a simple histogram filter to remove the background. Next we sampled a subset of the non-background points to both fit the curve and the tube. We used the modified principal curve-fitting algorithm with $K = 5$ final degrees of freedom to find the centerline. Next, we employed the tube-fitting algorithm with a time window width $r = .2$ and $\alpha = .15$. Other time windows produced generally similar results. However, shorter windows are more sensitive to local variations in the density of sampled points, whereas longer windows oversmooth and lose some gross anatomical features.

Figure 13 shows the sampled colon data, the fitted curve, and the fitted tube colored according to tracer concentration. As described above, the tracer concentration at each point along the curve was taken to be the summed concentration of those points used to define the tube at that point. Though the fitted tube plausibly recreates colon anatomy in terms of shape and width, we are unable to make a comparison between the fitted tube and the subject's colon. SPECT-CT scanners typically produce poor CT scans; therefore we lack good anatomical images that could be used to make this comparison. However, a benefit of the tube-fitting method is that it allows us to recreate the colon without radiating participants unnecessarily or requiring additional expensive equipment.

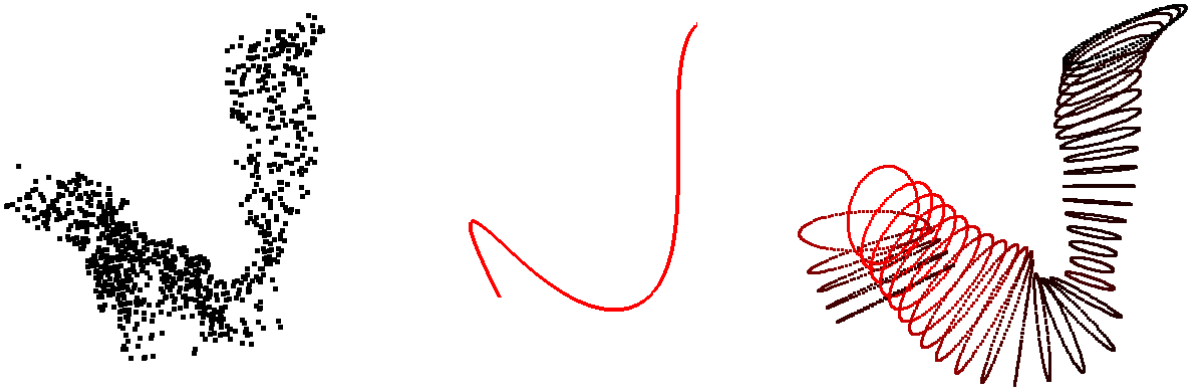

Figure 13: Three steps of the tube-fitting process. Farthest left is the sampled data from the SPECT image; center is the centerline produced by the curve-fitting algorithm; right is the fitted tube, shaded by tracer concentration (red is higher concentration, black is lower).

Figure 14 shows the concentration-by-distance curve. To find distance along the curve, we initially employed the arc-length formula 2 using the final fitted curve. Though often the

gradient of the fitted curve is easy to calculate, a closed-form solution for the integral is not available. We have found that simply calculating distance using the function value along the fine grid of values of t used to create the tube is equally accurate. That is, we simply use linear approximation between equally spaced latent time points to measure distance along the curve.

Computing the concentration at each distance from the curve onset can be accomplished in a variety of ways. Using the the neighborhood of t_0 described in section 4, we can define for each ellipse the collection of intensities $\{C_{i_j}\}$ for those points $\{P_{i_j}\}$ that are used to estimate the tube $G(t_0)$. A straightforward approach defines a proxy for the concentration as $\sum_{j=1}^{J} C_{i_j}$. However, a more accurate measure of concentration is $\frac{\sum_{j=1}^{J} C_{i_j}}{A}$, where $A = \text{area}\{\hat{G}(t_0)\}$, which takes the cross-sectional area of the colon into account. We compare these methods for finding concentration to those using a voxel-wise squared neighborhood approach [9]. This approach consists of finding all image points that fall within a cube of a given edge length and summing the concentrations of those points. Three comparative drawbacks are apparent in this method: *i* as in the case of the projections above, points that are near in terms of Euclidean distance but not t-distance may be included; *ii* there is no way to account for the width of the colon at each point; and *iii* the voxel-neighborhood approach is significantly more computationally intensive, especially for larger cube sizes. Notice in figure 14 that the voxel-neighborhood approach potentially underestimates concentration by averaging background voxels along with non-background.

6.2 DTI images

Our second application is to a diffusion-tensor tractogram of the intracranial portion of the corticospinal tract; this tract runs from the cerebral cortex of the posterior frontal lobe to the spinal cord and consists primarily of motor axons.

As with the colon application, we begin by implementing the curve-fitting algorithm to the image data to find a centerline. Though the imaged tract has less apparent complexity than the colon, to achieve an optimal fit we use $K = 8$ as the final degrees of freedom. It is also worth noting that the image contains only 231 locations, so no sampling is necessary. (The point density of DTI-derived tractograms can be highly variable, and, as noted above in Section 2.2, there is substantial undersampling bias for tractograms of the corticospinal tract.) Next, we employ the tube-fitting algorithm with time-window width of $r = .4$ and $\alpha = .1$. The time-window width is much wider than in the case of the colon application due to the relative sparsity of points: a wider window is necessary to fit reasonable bivariate normals, though such a wide window may smooth some of the finer details of the tract. A lower α is chosen because of the low noise level in the tract image.

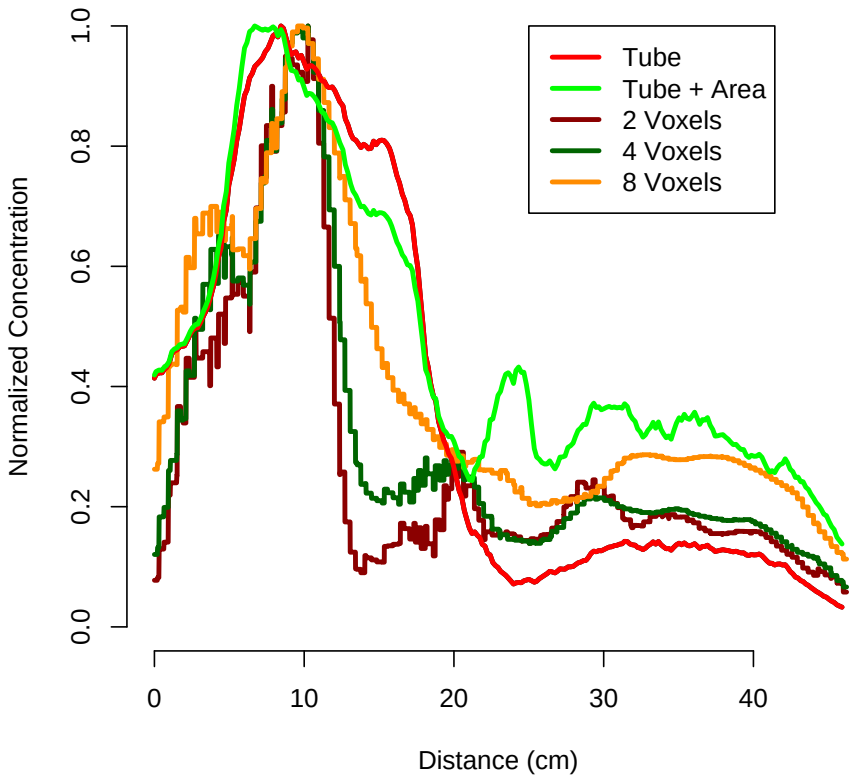

Figure 14: Concentration by distance from beginning of curve (near the anus). Concentration calculated using t-window and voxel-neighborhood approaches. Each curve is normalized to its maximum value.

Figure 15 superimposes the fitted tube on the anatomic image from which it is generated. The shading on this tube is constant because there is no intensity data associated with the tractography results (though see the discussion). However, useful plots may be formed by examining the area of orthogonal cross-section by distance along the tract. Additionally, the tract parameterization can enable location-specific measurement of underlying MRI quantities if the tube position is used as a region-of-interest.

We approximate the volume of the imaged tract as follows. Recalling that the level set of a bivariate normal is an ellipse, we calculate the cross-sectional area of the tube at many points along the curve using the formula $Area = \pi ab$, where a and b are one half the ellipse's major and minor axes respectively. Multiplying this area by the (very small) distance to the next point on the curve and summing the resulting volumes give a Reimann-sum-like

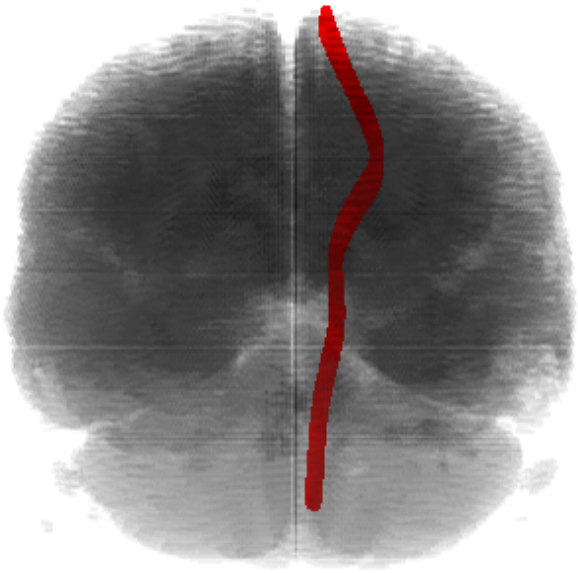

Figure 15: Fitted tube superimposed on the anatomic image.

approximation to the tract volume. Longitudinal measurements of tract volume or density for a multiple sclerosis patient is of clinical interest as a measure of disease progression.

7 Discussion

We have introduced a novel method for fitting three-dimensional tubes to imaging structures. Notably, we demonstrated the utility of the method on two very distinct imaging applications under different imaging modalities. In the colon SPECT imaging application, the tube-fitting method greatly improves upon previously used method of voxel-wise square neighborhoods. Moreover, our method produces an accurate mathematical model of the structure. Such characterization of the object of interest could be useful for subsequent shape analysis and for defining new measures of extent, volume, and other features of anatomical structures.

With regard to future applications, we note that the ongoing SPECT study is now collecting dual-isotope images, with the goal of comparing the relative distribution of microbicidal lubricants and HIV-infected semen in the colon. Surrogates for both are tagged with tracers emitting gamma photons with different are injected at the same time and are simultaneously imaged. Such dual-isotope studies may lead to drastically improved fits – using image data from both tracers greatly increases the number of points available for our curve- and tube-fitting algorithms. However, these studies raise the problem of accurately distinguishing and characterizing two tracer distributions in the colon. Moreover, the study now collects

images serially at several time points, giving us the opportunity to study changes in the concentration-by-distance curves over time. We are currently investigating the use of the accompanying X-ray computed tomography image for registration across time. Also, determining an anatomical landmark to compare curves across subjects remains difficult. We anticipate that bone landmarks from the CT image could be used to solve this problem, though we acknowledge that the colon can be fairly mobile across time and it's relation to bones may not be straightforward.

In the case of DTI tractography, the application of the tube-fitting algorithm to longitudinal images of multiple sclerosis patients will provide measures of disease progression. Ideally, one could use these measures to provide clinical evidence for the effectiveness of treatments. Validating the prediction performance of such measures remains an important problem. Moreover, the curve-fitting technique may not be applicable to all tracts (see Figure 16), and without an accurate centerline or without additional assumptions such as spatial contiguity of all points within the tract (so that the sampled tracts shown here can be considered non-random undersamplings), the tube-fitting algorithm will not work. Also, we are currently only using tracts created after ample preprocessing. Quantities derived from the original DTI image, such as anisotropy or diffusivity measurements, may produce more informative summaries of the tube. It is also possible that tube-fitting for this problem is best integrated into the tractography algorithm, which we have currently treated as a completely separate preprocessing step. Alternatively, the tensor itself could potentially be used to derive the individual tubes, obviating the need for tractography. However, we note that the fact that our algorithm only relies on existing tractography algorithms is also a strength, as it can be immediately applied.

We also note that one of the most important white matter tracts, the corpus callosum, which connects the left and right hemispheres of the brain, is not a tube-like structure. Instead it is more of a surface, with no clear centerline. Clearly to analyze such structures a different approach is necessary. We are investigating the possibility of using principal surfaces for this task [29, 10].

A related problem germane to both application is the study of curves and tubes across individuals and across time. For example, the analysis of the volume/distance curves or the analysis of other features estimated by the tube remains an open question.

The curve-fitting algorithm itself could be improved. As seen above, it is not universally applicable. Moreover, a more automated algorithm with less user input is desirable. We are currently experimenting with a new stochastic search algorithm for finding centerlines, such as the use of genetic algorithms and simulated annealing. A benefit of these approaches is the wide range of objective functions can be constructed to force a desired curve fit.

The tube-fitting algorithm presented here is a novel approach for the estimation of the

Figure 16: Example corticospinal tracts for which the curve-fitting algorithm (and therefore the tube-fitting algorithm) fails.

the support of distributions in three dimensions. It is limited in that it requires the support to have a reasonable centerline and in that it uses ellipses to estimate the cross sectional extent. However, it has proved a useful algorithm in two applications and holds a good deal of potential to be utilized in the field of medical imaging.

Acknowledgements

The authors would like to acknowledge the National Multiple Sclerosis Society for funding the acquisition of the DTI data.

References

[1] J.D. Banfield and A.E. Raftery. Ice Floe Identification in Satellite Images Using Mathematical Morphology and Clustering about Principal Curves. *Journal of the American Statistical Association*, 87(417), 1992.

[2] P. Basser, J. Mattiello, and D. LeBihan. Estimation of the effective self-diffusion tensor from the nmr spin echo. *J Magn Reson B*, 103:247–254, 1994.

[3] P. Basser, J. Mattiello, and D. LeBihan. Mr diffusion tensor spectroscopy and imaging. *Biophysical Journal*, Jan 1994.

[4] P. Basser, S. Pajevic, C. Pierpaoli, and J. Duda. In vivo fiber tractography using dt-mri data. *Magnetic Resonance in Medicine*, Jan 2000.

[5] P. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor mri. *J Magn Reson B*, 111:209–219, 1996.

[6] C. Beaulieu. The basis of anisotropic water diffusion in the nervous system - a technical review. *NMR in Biomedicine*, 15:435–455, 2002.

[7] I. Bitter, M. Sato, M. Bender, A. Kaufman, M. Wan, and M. Wax. Automatic, accurate and robust colon centerline algorithm. *Radiology*, 217:Suppl:705, 2000.

[8] I. Bitter, M. Sato, M. Bender, K. McDonnell, A. Kaufman, and M. Wan. Ceasar: A smooth, accurate and robust centerline extraction algorithm. *In Proc. of IEEE Visualisation 2000*, 2000.

[9] B.S. Caffo, C.M Crainiceanu, L. Deng, and C.W Hendrix. A Case Study in Pharmacologic Colon Imaging Using Principal Curves in Single Photon Emission Computed Tomography. *JASA*, 103:1470–1480, 2009.

[10] K.Y. Chang, J. Ghosh, I. Inc, and CA Sunnyvale. A unified model for probabilistic principal surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(1):22–41, 2001.

[11] P. Chaudhuri, R. Khandekar, D. Sethi, and P. Kalra. An efficient central path algorithm for virtual navigation. *Proceedings of the Computer Graphics International (CGI'04)*, 00:188 – 195, 2004.

[12] R. Chiou, A. Kaufman, Z. Liang, L. Hong, and M. Achniotou. Interactive path planning for virtual endoscopy. *In Proceedings of the IEEE Nuclear Science and Medical Imaging Conference*, 1998.

[13] R. Chiou, A. Kaufman, Z. Liang, L. Hong, and M. Achniotou. An interactive fly-path planning using potential fields and cell decomposition for virtual endoscopy. *IEEE Transactions on Nuclear Science*, 46:1045–1049, 1999.

[14] E. Cohen, R.F. Riesenfeld, and G. Elber. *Geometric Modeling With Splines: An Introduction*. AK Peters, Ltd., 2001.

[15] L. Cohen and R. Kimmel. Fast marching the global minimum of active contours. *International Conference on Image Processing, ICIP' 96, Lausanne, Switzerland*, 1:473–476, 1996.

[16] L. Cohen and R. Kimmel. Global minimum for active contour models: A minimal path approach. *Journal of Computer Vision*, 24 (1):57–78, 1997.

[17] Lijuan Deng. Spline-based curve fitting with applications to kinetic imaging. Master's thesis, Department of Biostatistics, Johns Hopkins University, 2007.

[18] T. Deschamps and L. Cohen. Fast extraction of minimal paths in 3d images and applications to virtual endoscopy. *Medical Image Analysis*, 5:281–299, 2001.

[19] E. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.

[20] T. Hastie and W. Stuetzle. Principal curves. *Journal of the American Statistical Association*, 1989.

[21] T Hastie, R Tibshirani, and J Friedman. *The Elements of Statistical Learning*. Springer-Verlag, New York, 2001.

[22] T.J. Hastie. Principal curves and surfaces. *Laboratory for Computational Statistics Technical Report 11, Stanford University, Dept. of Statistics*, 1984.

[23] L. Hong, S. Muraki, A. Kaufman, D. Bartz, and T. He. Virtual voyage: Interactive navigation in the human colon. *In Proceedings of SIGGRAPH'97*, 1997.

[24] HM Hudson and RS Larkin. Accelerated image reconstruction using ordered subsets of projection data. *IEEE Transactions on Medical Imaging*, 13(4):601–609, 1994.

[25] H. Jiang, P.C. van Zijl, J. Kim, G.D. Pearlson, and S. Mori. Dtistudio: resource program for diffusion tensor computation and fiber bundle tracking. *Computer Methods and Programs in Biomedicine*, 81:106–116, 2006.

[26] B. Kegl, A. Krzyzak, T. Linder, and K. Zeger. Learning and design of principal curves. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(3):281–297, 2000.

[27] E. Kreyszig. *Differential Geometry*. Courier Dover Publications, 1991.

[28] D. LeBihan, J. Mangin, C. Poupon, and C. Clark. Diffusion tensor imaging: Concepts and applications. *Journal of Magnetic Resonance Imaging*, Jan 2001.

[29] M. Leblanc and R. Tibshirani. Adaptive principal surfaces. *Journal of the American Statistical Association*, 89(425):53–64, 1994.

[30] E. McFarland, G. Wang, J. Brink, D. Balfe, J. Heiken, and M. Vannier. Spiral computed tomographic colonography: Determination of the central axis and digital unraveling of the colon. *Academy Radiology*, 4:367–373, 1997.

[31] S. Mori and P. Barker. Diffusion magnetic resonance imaging: its principle and applications. *Anat Rec (New Anat)*, Jan 1999.

[32] S. Mori, B.J. Crain, V.P. Chacko, and P.C. van Zijl. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*, 45:265–269, 1999.

[33] S. Mori and P.C. van Zijl. Fiber tracking: principles and strategies - a technical review. *Annals of Neurology*, 15:468–480, 2002.

[34] S. Mori, S. Wakana, P.C.M. van Zijl, and L.M. Nagae-Poetscher. *MRI Atlas of Human White Matter*. Elsevier, 2005.

[35] D.S. Reich, S.A. Smith, E.M. Gordon-Lipkin, A. Ozturk, B.S. Caffo, L.J. Balcer, and P.A. Calabresi. Damage to the optic radiation in multiple sclerosis is associated with retinal injury and visual disability. *Archives of Neurology*, 2009.

[36] D.S. Reich, S.A. Smith, C.K. Jones, K.M. Zackowski, P.C. van Zijl, P.A. Calabresi, and S. Mori. Quantitative characterization of the corticospinal tract at 3t. *American Journal of Neuroradiology*, 27:2168–2178, 2006.

[37] D.S. Reich, S.A. Smith, K.M. Zackowski, E.M. Gordon-Lipkin, C.K. Jones, J.A. Farrell, S. Mori, P.C. van Zijl, and P.A. Calabresi. Multiparametric magnetic resonance imaging analysis of the corticospinal tract in multiple sclerosis. *Neuroimage*, 38:271–279, 2007.

[38] D.S. Reich, K.M. Zackowski, E.M. Gordon-Lipkin, S.A. Smith, B.A. Chodkowski, G.R. Cutter, and P.A. Calabresi. Corticospinal tract abnormalities are associated with weakness in multiple sclerosis. *American Journal of Neuroradiology*, 29:333–339, 2008.

[39] Y. Samara, M. Fiebich, A. Dachman, K. Doi, and Hoffmann K. Automated centerline tracking of the human colon. *In Proceedings of the SPIE Conference on Image Processing*, 1998.

[40] Y. Samara, M. Fiebich, A. Dachman, J. Kuniyoshi, K. Doi, and K. Hoffmann. Automated calculation of the centering of the human colon on ct images. *Academic Radiology*, 6:352–359, 1999.

[41] J.A. Thorpe. *Elementary Topics in Differential Geometry*. Springer, 1979.

[42] R.P. Woods, S.R. Cherry, and J.C. Mazziotta. Rapid automated algorithm for aligning and reslicing pet images. *Journal of Computer Assisted Tomography*, 16:620–633, 1992.