Published in final edited form as:

J Am Stat Assoc. 2008 December 1; 103(484): 1470–1480. doi:10.1198/016214508000000832.

A Case Study in Pharmacologic Colon Imaging Using Principal Curves in Single Photon Emission Computed Tomography

Brian S Caffo, Ciprian M Crainiceanu, Lijuan Deng, and Craig W Hendrix

Abstract

In this manuscript we are concerned with functional imaging of the colon to assess the kinetics of microbicide lubricants. The overarching goal is to understand the distribution of the lubricants in the colon. Such information is crucial for understanding the potential impact of the microbicide on HIV viral transmission. The experiment was conducted by imaging a radiolabeled lubricant distributed in the subject's colon. The tracer imaging was conducted via single photon emission computed tomography (SPECT), a non-invasive, in-vivo functional imaging technique. We develop a novel principal curve algorithm to construct a three dimensional curve through the colon images. The developed algorithm is tested and debugged on several difficult two-dimensional images of familiar curves where the original principal curve algorithm does not apply. The final curve fit to the colon data is compared with experimental sigmoidoscope collection.

1 Introduction

Anti-microbial lubricants could have a profound impact in the spread of sexually transmitted diseases, such as the human immunodeficiency virus (HIV). Understanding the efficacy of these treatments requires understanding the anatomic distribution and the kinetics of viral-mixed semen along with those of the lubricant within the lumen of sexually receptive organs, ideally under typical forces of vaginal or anal intercourse. If the lubricant has a low penetration, or its distribution does not comport with the location of the viral-mixed semen, then its efficacy would be in question. However, the complexity of the physical systems involved prohibit the analytic study of such questions. Instead they must be evaluated experimentally.

This manuscript considers data arising from one the first experiments to empirically evaluate anti-microbial lubricants. Being foundational research, the aims are focused and hence we solely consider the problem of estimating the distributional penetrance of the lubricant after anal intercourse.

We now describe the experiment in greater detail. An over-the-counter lubricant was mixed with a radioactive tracer; this served as a surrogate marker for the anti-microbial treatment. This lubricant was inserted into a subject's colon and physical forces similar to those of anal coitus were subsequently used to distribute the liquid. The distribution of the tracer/lubricant mixture was then imaged using a single photon emission computed tomography (SPECT) scanner.

To elaborate on the imaging procedure, SPECT is a non-invasive, in-vivo functional imaging technique. SPECT images arise by the application of computed tomography techniques to projections obtained by counting emitted photons from a radioactive tracer placed in the body. SPECT images are of lower resolution than those obtained from other modalities, such as X-Ray CT, MRI and PET. However, SPECT represents a relatively low-cost method for obtaining functional imaging; that is to say it offers the ability to image the body as it functions through biological interactions with the tracer.

After scanning, a followup procedure was performed whereby a sigmoidoscope was used to sample the tracer/lubricant mixture at various positions within the colon. For our purposes, the sigmoidoscope was a lubricated tube with an optical fiber and an additional channel for a mechanical sampling device, which in this experiment was a brush inside a casing.

Before discussing the statistical goals of this manuscript, we briefly discuss relevant colon anatomy. A diagram of the colon is given in Figure 1. From the anus, the next structure is the anal canal, then rectum. The sigmoid colon follows by traveling a highly variable course, anteriorly and slightly inferiorly, from right to left, where it transitions into the descending colon, which travels up the left side of the abdomen.

A few sample transverse slices of the raw (reconstructed) image data are given in Figure 2. (A transverse plane divides the body into upper and lower regions.) As can be seen, the raw data is difficult to interpret, or to get any sort of anatomical bearings. Transverse slices of the data having been thresholded and overlaid on the accompanying X-ray CT image are given in Figure 3. Here the X-ray CT and SPECT are collected at the same time and registered in the same space by the scanner software. In the top four plots, one can see the hip bones on the X-ray CT image and the tracer distribution in the descending colon. The middle four plots show the tracer distribution around the lower portion of the descending colon and the sigmoid colon, whereas the lower four plots display the tracer near the rectum.

The primary goal of this investigation is to provide a semi-automated procedure to estimate the concentration of the lubricant by distance in the colon using the SPECT image. With reliable information from the imaging data, the sigmoidoscope collection would not be necessary in future studies. In addition, the sigmoidoscope itself displaces the liquid, hence has limitations for measuring the tracer distribution.

We focus our attention entirely on the SPECT images, as interest lies only in the distributional penetrance of the tracer/lubricant mixture. The accompanying X-ray CT images convey very little relevant additional information, since they were collected primarily as attenuation maps for the SPECT reconstruction. (Attenuation maps are used to perform the reconstruction algorithm that converts the raw data to images.) As such, they display almost no colon anatomy and largely highlight skeletal structures. In principle, other imaging modalities, such as high-resolution magnetic resonance imaging, could be used as anatomical complements to the SPECT images. However, it would be difficult to register these anatomical scans with the SPECT images. Moreover, they would have to be collected at a different time, hence the distribution of the tracer/lubricant mixture, as well as the colon shape itself, may have changed dramatically.

To solve the problem, we develop an algorithm based on principal curves (Hastie and Stuetzle, 1989) to perform the fitting. The resulting algorithm incorporates constrained endpoints (of the curve), constrained interior points and the image intensities. Moreover, a novel “molding” procedure is proposed that greatly improves the ability of the algorithm to fit complex data structures. The algorithm is tested and debugged on a collection of difficult two-dimensional (2-D) images that the original, unmodified, principal curve algorithm can not fit.

With one subject, the scientific contribution of this work is largely a proof of concept. However, it will be shown that the algorithm appears to work quite well and requires little user input. The excellent results of the algorithm have lead to the potential policy change of eliminating the costly and invasive sigmoidoscope collection.

The article is organized as follows. Section 2 overviews the data. Section 3 covers curve fitting algorithms, beginning with a literature review. Subsection 3.1 covers curve characterizations while 3.2 discusses principal curves while our modified algorithm is discussed in 3.3. Section

4 tests these algorithms on constructed 2-D data, while the algorithms are applied to the real SPECT data in Section 5. The manuscript concludes with a discussion in Section 6.

2 Data

Ten milliliters of radiolabeled lubricant (99 *mTc*-sulfur colloid mixed with K-Y Jelly® Johnson and Johnson, New Brunswick, NJ) were injected into the subject's colon. Following rectal administration of the radiolabeled gel, the subject underwent simulated receptive anal intercourse. The experimental paradigm was designed to mimic the typical forces that would influence the lubricant's distribution.

The patient was imaged on a dual-head VG SPECT-CT imaging system (GE Medical Systems, Waukesha, WI) equipped with a low-end computed tomography (CT) unit (Hawkeye). Accompanying each SPECT image, an X-ray computed tomography image was also collected for anatomical reference and reconstruction of the SPECT image. The image was reconstructed using the ordered subsets EM algorithm (Hudson and Larkin, 1994) and filtered as provided with the scanner software (General Electric eNTEGRA workstation, version 1.04, GE Medical Systems, Waukesha, WI).

The SPECT data is represented as a three dimensional array. In our application the dimension of the array is $128 \times 128 \times 128$. Each voxel (three dimensional pixel) represents a 3.45 mm^3 physical area. The image intensity values are proportional to the concentration of the tracer at that location. The intensity values range from 0 to 187, though we note that the absolute scale is somewhat arbitrary, as the image was rescaled during reconstruction.

After imaging, a sigmoidoscope was used to collect physical concentration measurements. Samples were taken at 5 *cm* intervals up to 40 *cm*. The study was approved by the Johns Hopkins Institutional Review Board and informed written consent was given.

3 Curve fitting algorithms

Calculating centerlines for anatomical structures such as blood vessels, neurons or colons has a rich history in the computer vision and medical image processing literature. Much of the research involving colons are applied to X-ray CT images for the purposes of finding polyps. For example, McFarland et al. (1997) presented a semi-automated centerline extraction algorithm. Samara et al. (1998) proposed a semi-automated voxel search algorithm for centerline construction.

Another class of methods employ Dijkstra's algorithm (see Dijkstra, 1959), where voxels and intensity values are treated as a networked graph. Search algorithms are used to find minimal distance paths through the graph (Chiou et al., 1998; Bitter et al., 2001; Hong et al., 1997). Wan et al. (2001) used distance fields to compute central paths, while Chaudhuri et al. (2004) used similar concepts, but with a different distance measure. Cohen and Kimmel (1997) derived a path tracking routine in two-dimensional images by calculating a minimal path between two fixed end points.

Deschamps and Cohen (2001) extended the so-called fast marching algorithm to three dimensional objects to extract a minimal path through the colon. Ge et al. (1999) used a fast topological thinning algorithm to generate a three dimensional skeleton of a binary colon volume, which was subsequently pruned. Bouix et al. (2003) used a technique called medial surface extraction to compute a centerline curve, which was then pruned. Finally, Telea and Vilanova (2003) gave a level-set algorithm for building a colon centerline.

After having mentioned only a subset of the related work on calculating colon centerlines, we emphasize that our problem differs markedly from these approaches in six key aspects: *i*) this work considers SPECT images rather than high resolution anatomical X-ray CT images; *ii*), the image is of the tracer/lubricant mixture, not of colon anatomy; the tracer may be at lower concentrations at different areas of the colon, where the centerline is still desired, so that image intensities are not the primary concern (as opposed to the analysis of X-ray CT images); *iii*), unlike the colon anatomy, the tracer distribution can be discontinuous, interrupted by stool and gas; therefore, techniques requiring connected graphs would not apply; *iv*) the scientific application is extremely novel, with no comparable experiments; *v*) there is a direct comparison measurement available in the sigmoidoscopy data; *vi*) our approach and characterization of the problem is more statistically oriented than existing algorithms.

Statistical approaches in the area of curve fitting are few - as opposed to the embarrassment of riches available for fitting proper functions. Below we discuss the principal curve algorithm, a fundamental algorithm in the area of curve fitting. However, we found that, unmodified, this algorithm could not handle complex images. Moreover, it does not incorporate image intensities. Therefore we propose a modified algorithm with several notable benefits.

3.1 Characterization of curves

We characterize the problem as follows. Let $f: \mathbb{R} \rightarrow \mathbb{R}^3$ be defined so that $f(t) = \{f^x(t), f^y(t), f^z(t)\} = \{X(t), Y(t), Z(t)\}$. Here, f is the three dimensional position of a curve through the colon at a latent argument $t \in [0, 1]$. The value of $f(t)$ represents the coordinate points in the image; hence, the curve in three dimensional space is then the projection of $f(t)$ over t . As such, this is a standard representation of a curve in three dimensional space (see Thorpe, 1979, for example). The constraint that t resides in $[0, 1]$ is arbitrary, and used for identifiability. Throughout we conceptualize $f(t)$ as the position of a curve at a “time point” t . We emphasize that this is only for pedagogy; the image has no temporal component. Conceptually considering time highlights the identifiability issue that two functions traveling the same path at different rates yield the same curve.

The general requirements for the fitted curve are: *i*) the curve should follow a smooth path through the center of the tracer distribution, *ii*) the curve must be able to traverse empty spaces, where the continuity of the tracer is interrupted, *iii*) user specified starting and ending points should be incorporated, *iv*) the algorithm should be flexible enough to allow for other user specified points that the curve must travel through and *v*) the curve should prefer traversing higher intensity points to lower. We note that point *v*) must be factored in carefully, as excessively emphasizing the image intensities can result in poor fits. Consider the example data, where the highest intensity points are aggregated near the anus. Any algorithm that minimizes the curve integral of the intensities through the image will focus on spending as much time as possible in these areas, contrary to goals *i*) and *ii*).

To achieve these requirements we build on the principal curve algorithm of Hastie and Stuetzle (1989). Our modifications grew out of difficulty in getting reasonable fits from the original algorithm for challenging imaging problems and the fact that it was not designed to incorporate image intensity values. Our approach does require user-specified endpoints for the curve, mostly because of the nature of the scientific problem. However, a novel molding procedure automates the selection of a starting curve, starting at a line connecting the two specified endpoints and gradually capturing finer details of the curve. Moreover, we provide an elegant solution to accounting for differences in image intensity that does not suffer from the problems outlined above.

We describe the principal curve algorithm and our modifications in the next two sections. However, before a discussion of the algorithms, we discuss necessary preprocessing and some basic starting points for curve fitting through images.

A first step is to threshold the image. This is done both to remove background noise and to reduce the number of points included in the fitting. Further acceleration is obtained by randomly sampling points that survive the threshold. Points should be sampled uniformly among those that are above the threshold. Sampling points with probabilities relative to image intensities gives poor results, because areas of the colon with low intensities are omitted.

Notationally, let $\{(X_i, Y_i, Z_i)\}_{i=1}^n$ be the sampled locations. To be specific, (X_i, Y_i, Z_i) represents the lattice value of a single sampled point above the original threshold. The index ordering, i , is arbitrary; that is, the points can be selected in any order. Figure 4 shows the sampled colon data. Let $\{t_i\}_{i=1}^n$ represent a collection of (unknown) associated time variables for the sampled points. Throughout, dropping the subscript will represent the vector of the collection of variables, such as $t = (t_1, \dots, t_n)'$.

Much of the challenge of this problem is to appropriately assign values for the time points. Tabling this issue for the moment, presume that these points were known. Then a smooth curve through the data could be easily fit. Our approach uses the three spline equations:

$$\begin{aligned} E[X_i] = f^x(t_i) &= \beta_0^x + \beta_1^x t_i + \beta_2^x t_i^2 + \beta_3^x t_i^3 + \sum_{k=1}^K b_k^x (t_i - \xi_k)_+^3, \\ E[Y_i] = f^y(t_i) &= \beta_0^y + \beta_1^y t_i + \beta_2^y t_i^2 + \beta_3^y t_i^3 + \sum_{k=1}^K b_k^y (t_i - \xi_k)_+^3, \\ E[Z_i] = f^z(t_i) &= \beta_0^z + \beta_1^z t_i + \beta_2^z t_i^2 + \beta_3^z t_i^3 + \sum_{k=1}^K b_k^z (t_i - \xi_k)_+^3, \end{aligned} \quad (1)$$

where $\{\xi_k\}_{k=1}^K$ are knots placed at equally spaced quantiles of $\{t_i\}_{i=1}^n$. Here $K + 4$ is the degrees of freedom of the smoother for each dimension. The benefits of using regression splines are many, including the easy specification of the basis and the easily derived derivatives, which are required later.

3.2 Principal curves

The principal curve algorithm is a general method for fitting a curve through data residing in an arbitrary dimensional space. A curve, $f(t)$, is said to be a principal curve if for each data value, say (X_i, Y_i, Z_i) , the curve at t_i satisfies

$$f(t_i) = E[(X_i, Y_i, Z_i) | \text{closest point of } f \text{ to } (X_i, Y_i, Z_i) \text{ occurred at time } t_i].$$

This recursive form of self-consistency motivates an algorithm. Suppose that a starting collection of time points is given.

1. Approximate the principal curve by a scatterplot smoother, regressing $\{(X_i, Y_i, Z_i)\}_{i=1}^n$ on $\{t_i\}_{i=1}^n$ [as in Equation (1)].
2. Update the time points by redefining t_i as the time point on the curve closest to (X_i, Y_i, Z_i) for $i = 1, \dots, n$. That is, define

$$t_i = \operatorname{argmin}_{t \in [0,1]} \sqrt{\left\{X_i - \tilde{f}^x(t)\right\}^2 + \left\{Y_i - \tilde{f}^y(t)\right\}^2 + \left\{Z_i - \tilde{f}^z(t)\right\}^2}$$

for $i = 1, \dots, n$, where $\tilde{f}^x(t)$, for example, represents the current estimate of $f^x(t)$.

These steps are then iterated until convergence. Of course, instead of starting at a collection of time points, a starting curve could be given, in which case the algorithm simply starts at Step 2. For example, one could start the algorithm at the first principal component line through the data. In our setting, we start the algorithm at a line connecting user-specified endpoints.

Conceptually the steps can be thought of as the following. First the data is projected onto the three planes considering time and each spatial dimension. That is, the data (t, X) , (t, Y) and (t, Z) are considered; where, for example, (t, X) refers to the collection of $\{t_i\}_{i=1}^n$ and $\{X_i\}_{i=1}^n$ data points. Secondly, a scatterplot smoother is fit in those three planes to calculate an updated curve. Next, the orthogonal projections of the data points onto the curve are calculated. The time points associated with the projections onto the curve are used to then update the latent time variable. While this strategy has considerable intuitive appeal, a formal proof of convergence is not available. However, it has been used successfully in (unrelated) image processing settings (Banfield and Raftery, 1992).

3.3 A modified principal curve algorithm

In this section we discuss generalizations of the principal curve algorithm that allow it to be viable in our scientific setting. These modifications are: allowing for user-specified endpoints and interior points of the function, incorporating the image intensities, warming-up, or molding, the algorithm to achieve better fit, a grid search to perform the minimization in the second step of the algorithm and a stopping rule based on relative mean squared error.

Consider the incorporation of user-specified endpoints. Notationally suppose that (x_0, y_0, z_0) and $t_0 = 0$ and $(x_{n+1}, y_{n+1}, z_{n+1})$ and $t_{n+1} = 1$ are given. A modification of the algorithm that forces the curve to start and end at these points (respectively) simply forces the constraint $0 \leq t_i \leq 1$ for $i = 1, \dots, n$ and adds the relevant Lagrange multiplier terms to Equation (1). That is, the multiplier terms force $f^x(0) = x_0, f^x(1) = x_{n+1}, f^y(0) = y_0, f^y(1) = y_{n+1}, f^z(0) = z_0, f^z(1) = z_{n+1}$. Specifically, let W be the basis matrix for Equation (1) (note the same basis is used for all three dimensions). As in Equation (1), we used regression splines with knots at equally spaced points between 0 and 1, though any linear smoother applies. Let $\tilde{W}$ be the basis evaluated at the constrained values of t and let $\tilde{x}, \tilde{y}$ and $\tilde{z}$ be vectors of the constrained values. Then the goal is to maximize the least squares equations for the models

$$E[X] = W\beta^x \quad E[Y] = W\beta^y \quad E[Z] = W\beta^z$$

subject to the constraints

$$\tilde{W}\beta^x = \tilde{x} \quad \tilde{W}\beta^y = \tilde{y} \quad \tilde{W}\beta^z = \tilde{z}$$

The fitted values are then (see Searle, 1971, for example)

$$\begin{aligned}
\widehat{\beta}_c^x &= \widehat{\beta}^x - (W'W)^{-1} \widetilde{W} \left\{ \widetilde{W}'(W'W)^{-1} \right\}^{-1} (\widetilde{W} \widehat{\beta}^x - \widetilde{x}), \\
\widehat{\beta}_c^y &= \widehat{\beta}^y - (W'W)^{-1} \widetilde{W} \left\{ \widetilde{W}'(W'W)^{-1} \right\}^{-1} (\widetilde{W} \widehat{\beta}^y - \widetilde{y}), \\
\widehat{\beta}_c^z &= \widehat{\beta}^z - (W'W)^{-1} \widetilde{W} \left\{ \widetilde{W}'(W'W)^{-1} \right\}^{-1} (\widetilde{W} \widehat{\beta}^z - \widetilde{z}),
\end{aligned} \tag{2}$$

where $\widehat{\beta}^x$, $\widehat{\beta}^y$ and $\widehat{\beta}^z$ are the unconstrained fitted coefficients.

It is worth noting that constraining the endpoints as such implies that the final fitted curve will not satisfy the self-consistency property. In particular, points near the fixed endpoints will *not* have their values of t updated by the orthogonal projection to the curve extended in perpetuity, but instead by the closest value to the constrained ends.

This solution can be adapted to incorporate other constrained points along the curve. However, unlike the endpoints, the corresponding values of t_i are not known. Therefore, these associated time points must be estimated in Step 2 of the principal curve algorithm and the matrix of constrained time points, $\widetilde{W}$, must be updated. We note that this procedure must be used with care, as identifiability problems can occur with constrained interior points. A minimum requirement is that there generally needs to be more parameters than constrained points. For example, one cannot constrain a line to traverse three points (unless those points fall exactly on a line).

Consider incorporating the image intensities. Specifically, define Σ^{-1} to be a matrix with some function of the image intensities corresponding to the points $\{(X_i, Y_i, Z_i)\}_{i=1}^n$ along the main diagonal and zeros elsewhere. Then consider the weighted regression version of Equation 2

$$\begin{aligned}
\widehat{\beta}_c^x &= \widehat{\beta}^x - (W'\Sigma^{-1}W)^{-1} \widetilde{W} \left\{ \widetilde{W}'(W'\Sigma^{-1}W)^{-1} \right\}^{-1} (\widetilde{W} \widehat{\beta}^x - \widetilde{x}), \\
\widehat{\beta}_c^y &= \widehat{\beta}^y - (W'\Sigma^{-1}W)^{-1} \widetilde{W} \left\{ \widetilde{W}'(W'\Sigma^{-1}W)^{-1} \right\}^{-1} (\widetilde{W} \widehat{\beta}^y - \widetilde{y}), \\
\widehat{\beta}_c^z &= \widehat{\beta}^z - (W'\Sigma^{-1}W)^{-1} \widetilde{W} \left\{ \widetilde{W}'(W'\Sigma^{-1}W)^{-1} \right\}^{-1} (\widetilde{W} \widehat{\beta}^z - \widetilde{z}),
\end{aligned}$$

where now $\widehat{\beta}^x$, $\widehat{\beta}^y$ and $\widehat{\beta}^z$ are the unconstrained weighted fitted coefficients. For example,

$$\widehat{\beta}^x = (W'\Sigma^{-1}W)^{-1} W'\Sigma^{-1}X$$

Note that the fit is invariant to scalings of the image intensities. Though we used the raw image intensities to define Σ^{-1} , other definitions could be used to adapt the fit. For example, using the square of the intensities will put greater emphasis on high intensity points while using the square root will put less. Users of the algorithm should be advised that it may be useful to trim the intensities to avoid excessive impact for outlying points. This was not necessary in our application.

The most important modification of the principal curve algorithm lies in choosing an appropriate starting curve. We have found that getting the curve in the correct neighborhood is crucial for fitting complicated structures. Therefore, we employ a series of warm-up runs, using few degrees of freedom (small K) to obtain starting values that correctly model the gross features of the data. We start the warm-up runs at a line connecting the two specified endpoints.

For each value of K , the modified principal curve algorithm could then be run to convergence. After convergence, the estimated curve is used as a starting value for a subsequent warm-up run of the algorithm with K increased. A final run with the desired value of K uses the result of the warm-up runs as the starting value.

This method tends to mold the curve to the gross features of the data before moving on to the finer ones. Therefore we refer to this procedure as “molding”. We have found that molding is the single most important aspect of obtaining reasonable fits. Furthermore, the most critical aspect of using these warm-up runs is systematically increasing the degrees of freedom. Whether or not the algorithm is run until convergence within each warm-up value of K does not seem to impact results, other than slowing the algorithm down. Therefore, only one iteration for each value of K was used to warm-up to the final run.

Two algorithms were compared to perform the maximization required in Step 2 of the principal curve algorithm. First, a modified BFGS algorithm was used that can accommodate constrained endpoints (Byrd et al., 1995). Secondly, a simple grid search was also used. Both techniques were effective, though the grid search was quite a bit faster. Therefore, the results presented employed that technique using 1, 000 or 10, 000 grid points between 0 and 1. All calculations were performed in the R programming language (R Development Core Team, 2006).

With regard to stopping the algorithm, we suggest the relative change in the mean squared error of the estimated function summed across the three dimensions. Specifically, define

$$SMSE = \frac{1}{n} \sum_{i=1}^n \left[\{X_i - f^x(t_i)\}^2 + \{Y_i - f^y(t_i)\}^2 + \{Z_i - f^z(t_i)\}^2 \right]$$

and hence the stopping criteria required

$$\frac{|SMSE^{current} - SMSE^{old}|}{SMSE^{old}}$$

to be less than a desired tolerance.

4 Idealized 2-D curves

To build intuition, the algorithm was applied to highly idealized two-dimensional images created using image processing software. Figure 5 displays the raw images. Note they are of constant intensity, hence the weight matrix, Σ^{-1} , was set to an identity matrix. These images were selected as they possess several interesting features, such as sharp turns, and are quite a bit more difficult to fit than the actual data set. Many of these test images were motivated by Kegl et al. (2000). The 512×512 pixel images were created in the GIMP (GNU Image Manipulation Program, www.gimp.org) by freehand drawing with a mouse.

Figure 6 shows the sample fitted curves to the spiral using $K = 20$ to define the degrees of freedom (bottom plots), as well as the starting line (top plots). The fixed endpoints are shown as blue points. In each row, the leftmost plots show the fitted curve. The next plots show the (X, Y, t) data and the fitted curve in three dimensional space with the (X, Y) data shown projected onto the horizontal plane. The final two plots to the right show the (t, X) and (t, Y) data and the fitted curve projections. Recall that the smoothing portion of the algorithm fits a smoother to these two projections.

The same plots for the other examples are given in Figure 12. In all of the cases, the fit is obviously good. Because of the sampling and the fast grid-search maximization, the algorithms take only roughly thirty seconds to run (on a Pentium dual core, 2.16 *GHz* laptop with 2 *GB* of ram). Also, the process of molding the fit tends to get the curve in a very close neighborhood of the desired limit, thus requiring very few iterations for the final run.

We also explored the use of constrained interior points. Figure 7 shows the results of the fitted models employing constrained interior points for the “3” and spiral images (shown in green). Here $K = 20$ was used to define the final degrees of freedom of the smoothers. For the “3” image, the constrained interior points were nicely incorporated into the smooth curve. In contrast, for the spiral, the constrained interior points interfered with the fit, forcing the curve to make unnatural bends to incorporate them. We find that unless constrained interior points are carefully chosen and implemented, their use can greatly reduce the quality of the resulting fit. This can arise from a conflict between our opinion of where the curve should traverse and that of the principal curve algorithm's. Also, constrained interior points can prevent the molding procedure from exploring poorly fitting functions that only model gross features but lead the algorithm to a space of curves that fit the finer ones.

4.1 Comparison with alternative methods

In this section we briefly compare the modified principal curve approach with related methods and particularly the principal curve algorithm without the molding procedure, but still constraining endpoints. We note that the original principal curve algorithm without constrained endpoints performs very poorly, as it was designed for more variable data.

We first consider an alternative algorithm for curve fitting discussed in Deng (2007). Here the time points were estimated using an objective function that sought to maximize the curve integral of the fitted curve through the image, subject to a length penalty. A stochastic search algorithm was employed that made assumptions on the latent time variables. Specifically, a “constant speed” assumption was made that equally spaced the values of t_i between 0 and 1. Moreover it was presumed that the values of t_i were monotonic in either the sagittal, axial or transverse directions. Though the algorithm performs quite well when these assumptions are met, it failed with mild departures from the monotonicity assumption. Figure 8 shows the problem in an idealized two-dimensional colon-type image created using image processing software. Specifically around the sigmoid-curve, the monotonicity assumption fails, and hence a poor fit is obtained. For figures such as the spiral, where monotonicity fails regardless of the rotation of the image, this method does not apply.

To further emphasize the importance of the molding procedure, Figure 9 shows fitted curves obtained using a fixed degrees of freedom starting from a line for $K = 3, 5, 10$ and 20, for the four idealized images. In all cases, the curve is struggling to hit as many points as possible, to minimize the orthogonal projections onto the curve. The curve does not have the ability to fit gross features and so immediately takes as complex a shape as possible. This problem could be overcome by starting the curve using several constrained interior points, as described in Section 3.3, then perhaps one could quickly discard them, as a few iterations would place the curve in the right neighborhood. However, this would require excessive user input and the molding procedure automatically avoids this problem.

5 Application to the microbicide imaging study

After building intuition with the idealized two-dimensional data, we applied the methods to the SPECT data. The SPECT data was processed (reconstructed, filtered, thresholded and sampled) as outlined in the previous sections. The algorithm was run using $K = 5$ to define the smoother degrees of freedom. The low number was used because the tracer does not extend

past the descending colon, and follows a very smooth function, with only mild complexity near the sigmoid colon.

One thousand points were sampled from the image and a grid search using ten thousand equally spaced points between zero and one were employed in the second step of the principal curve algorithm. The modified BFGS algorithm was also employed, though it demonstrated no difference in results. The constrained endpoints were selected by comparisons with bone structure from the X-ray CT image. No constrained interior points were necessary, as the lubricant distribution is fairly continuous.

Figure 10 shows the fit in three orientations, two of which display the sampled data, one showing the intensities and the other displaying the orthogonal projections onto the curve. The shape of the curve was evaluated by physicians, who claimed it closely followed prior knowledge regarding colon anatomy.

The fitted curve was compared with the sigmoidoscope results. The distance along the curve was calculated as

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

where $\widehat{f}^x(t)$, for example, denotes the final estimate for $f^x(t)$. Because of the simple form for the basis used, f has easily calculated derivatives. We did not, however, find a closed form for the resulting integral, which was evaluated numerically. Image intensities were then calculated along the curve using the original, unsampled, image in neighborhoods of size 1, 2, 4 and 8 voxels. Here a neighborhood of size c is defined as the box centered at the current voxel with sides $(2c + 1)$. The concentration was estimated as the sum of the image intensities in the neighborhood, divided by the number of voxels.

Figure 11 displays the results as well as the sigmoidoscope data for comparison. Since image processing arbitrarily scales the reconstructed SPECT image, the intensities used in the concentration calculation are only proportional to the actual tracer concentration. Therefore, all of the curves were divided by their maximum value, to obtain a unit-free comparison.

Because the SPECT image does not indicate the exact location of the start of the anus, the curves were calibrated so that their maxima match that of the sigmoidoscope data. However, there is some degree of error as inspection of the SPECT image clearly indicates that the maximal concentration is located in the sigmoid colon. In contrast, the maximum for the sigmoidoscope data, 35 cm (1.15 feet), is too far into the colon for the sigmoid. In addition to the maximum being too distal, the sigmoidoscope data also appears more diffusely distributed than the SPECT data. This is consistent with the hypothesis that the repeated application of the scope distributes the tracer/lubricant mixture, thus giving an inaccurate picture of penetrance. Also, unlike the sigmoidoscope, the SPECT image can reasonably detect the tracer distribution well into the colon. As such, the image processing tools presented offer an extremely valuable source of quantification of the tracer/lubricant distribution without this source of error.

Currently the authors are working with scientific collaborators to obtain accurate measures of the starting distance of the tracer distribution in the SPECT image.

6 Discussion

In this manuscript we consider the difficult problem of three dimensional curve fitting in a novel, scientifically important study. The success of the algorithm will have a direct benefit

on the ability to estimate distributional penetrance in SPECT colon imaging. The performance of the algorithm has inspired the possibility that the costly and invasive sigmoidoscope collection can be abandoned and replaced by statistical image processing tools.

With regard to the future of this scientific application, researchers are now considering dual isotope studies, where both a surrogate for the lubricant and a surrogate for the HIV-infected semen are simultaneously imaged by being tagged with tracers at different energy levels. If successful, such experimental technique will offer accurate experimental validation of anti-microbial lubricants.

We also note that the algorithm has potential further scientific application in unrelated areas. Specifically, it is currently being evaluated for its viability for estimating white matter tracts from magnetic resonance images of the brain.

The developed algorithm proved to be a near ideal candidate for obtaining a centerline through the distribution of the SPECT tracer. Techniques for constraining endpoints and interior points were given. Furthermore, in the challenging 2-D test images, the single most important modification was clearly the molding procedure, where lower smoothing degrees of freedom were used to capture gross features of the data before moving on to the finer ones.

While this research has provided scientific collaborators with a set of easily implementable tools for processing their images and obtaining concentration/distance curves, the larger, more general, problem of arbitrary dimensional curve fitting leaves much room for further methodological development. We have found that this collection of problems, while considered in the image processing literature, has received insufficient attention in the statistics literature.

As for potential future directions, note that the algorithm is not completely automated. For example, starting and ending points of the curve must be given and a final degrees of freedom for the smoothers must be settled on. For the latter point, given that the algorithm runs very quickly, we recommend trying several values of K and using visual inspection and prior anatomical knowledge to decide on a final value. In particular, as K gets too large, the curve will make visually unnatural bends to minimize the orthogonal projections of the data onto the curve. For future research, we believe that penalized spline approaches for fitting may more completely automate the procedure. In addition, it is possible that separate degrees of freedom, and separate bases, could be used for f^x , f^y and f^z respectively. Finally, given the variation in complexity of some of two-dimensional shapes, smoothers that can accommodate variable degrees of freedom across t (adaptive smoothers) may be able to fit more complex structures.

Finally, though the developed algorithm proved to be very successful in this application, other methods for allocating the time points could be considered in other settings. For example, Deng (2007) considered a stochastic search algorithm under a user specified objective function. In addition, as a referee has pointed out, the latent time variable t could be treated in a Bayesian manner. That is a prior on the time variable could be set, along with priors on all of the smoothing parameters. An added benefit of such an approach is the built-in ability to quantify uncertainty in the curve fits.

References

- Banfield J, Raftery A. Ice Floe Identification in Satellite Images Using Mathematical Morphology and Clustering about Principal Curves. *Journal of the American Statistical Association* 1992;87(417)
- Bitter I, Kaufman A, Sato M. Penalized-distance volumetric skeleton algorithm. *IEEE Trans Vis Comp Graph* 2001;7:195–206.
- Bouix, S.; Siddiqi, K.; Tannenbaum, A. Flux driven fly throughs. In 2003 conference on Computer Vision and Pattern Recognition (CVPR 2003); 2003. p. 449-454.

- Byrd R, Lu P, Nocedal J, Zhu C. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing* 1995;16(5):1190–1208.
- Chaudhuri P, Khandekar R, Sethi D, Kalra P. An efficient central path algorithm for virtual navigation. *Proceedings of the Computer Graphics International (CGI'04)* 2004;00:188–195.
- Chiou, R.; Kaufman, A.; Liang, Z.; Hong, L.; Achniotou, M. Interactive path planning for virtual endoscopy. In *Proc. of the IEEE Nuclear Science and Medical Imaging Conference*; 1998.
- Cohen L, Kimmel R. Global minimum for active contour models: A minimal path approach. *Journal of Computer Vision* 1997;24(1):57–78.
- Deng, L. Master's thesis. Department of Biostatistics; Johns Hopkins University: 2007. Spline-based curve fitting with applications to kinetic imaging.
- Deschamps T, Cohen L. Fast extraction of minimal paths in 3d images and applications to virtual endoscopy. *Medical Image Analysis* 2001;5:281–299. [PubMed: 11731307]
- Dijkstra E. A note on two problems in connexion with graphs. *Numerische Mathematik* 1959;1:269–271.
- Ge Y, Stels D, Wang J, Vining D. Computing centerline of a colon: A robust and efficient method based on 3d skeletons. *Journal of Computer Assisted Tomography* 1999;23(5):786–794. [PubMed: 10524868]
- Hastie T, Stuetzle W. principal curves. *Journal of the American Statistical Association*. 1989
- Hong L, Muraki S, Kaufman A, Bartz D, He T. Virtual voyage: Interactive navigation in the human colon. In *Proc. of SIGGRAPH'97*. 1997
- Hudson HM, Larkin RS. Accelerated image reconstruction using ordered subsets projection data. *IEEE Transactions on Medical Imaging* 1994;13:601–609. [PubMed: 18218538]
- Kegl B, Krzyzak A, Linder T, Zeger K. Learning and design of principal curves. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 2000;22(3):281–297.
- McFarland E, Wang G, Brink J, Balfe D, Heiken J, Vannier M. Spiral computed tomographic colonography: Determination of the central axis and digital unraveling of the colon. *Acad Radiol* 1997;4:367–373. [PubMed: 9156234]
- R Development Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing; Vienna, Austria: 2006. ISBN 3-900051-07-0
- Samara, Y.; Fiebich, M.; Dachman, A.; Doi, K.; K., H. Automated centerline tracking of the human colon. In *Proc. of the SPIE Conference on Image Processing*; 1998.
- Searle, S. *Linear Models*. Wiley; 1971.
- Telea A, Vilanova A. A robust level-set algorithm for centerline extraction. In *Proceedings of the symposium on Data visualisation 2003* 2003;185–194.
- Thorpe, J. *Elementary Topics in Differential Geometry*. Springer-Verlag; New York: 1979.
- Wan M, Dachille F, A. K. Distance-field based skeletons for virtual navigation. In *IEEE Visualization* 2001 2001:239–245.

Figure 1.
Anatomical diagram of the colon.

Figure 2.

Example raw image data axial slices. The left two show the tracer distribution in axial slices of the descending colon while the third from the left shows distribution around the sigmoid colon. The final plot shows the distribution near the anus.

Figure 3.

Processed SPECT images by transverse slices; images proceeding from the upper left to the lower right proceed inferiorly (toward the feet). The selected images are spaced roughly 10 *cm* apart with each image representing a 3.45 *mm* thick slice. For anatomical reference, the hip bones are clearly visible in the upper left hand plots.

Figure 4.
Sampled colon images in two orientations. Colors represent image intensity.

Figure 5.
Sample images.

Figure 6.

Fitted images for the spiral. In the top panel the line interpolating the two (blue) specified endpoints is given where the bottom plots the fits after being run to twenty degrees of freedom. The first plots on the left show the fitted curve. The second shows the data and the fitted curve, both projected onto the (X, Y) plane as well as the full data with the unobserved variable t . The latter two plots show the (t, X) and (t, Y) projections of the data and the fitted curve.

Figure 7. Results of model fits to the “3” and spiral using constrained interior points. The constrained interior points are in green while the constrained endpoints are shown in blue.

Figure 8. Stochastic search curve fitting algorithm of Deng (2007) used in a idealized two-dimensional colon image.

Figure 9. Results of principal curve fits not employing the molding procedure. From left to right, the plots use 3, 5, 10 and 20 degrees of freedom. From top to bottom they are the spiral, sawtooth, “3” and “a” images from Figure 5.

Figure 10.

Fitted colon curve in three orientations. The first two plots also show the sampled data points, one with image intensities and another with the end orthogonal projections. The final plot shows the curve and the (blue) constrained endpoints.

Figure 11.

Concentration estimates by distance (in centimeters) from the curve beginning (near the anus) using various neighborhood sizes around the curve. The sigmoidoscope data is shown in black. All curves are normalized relative to their maximum value.

Figure 12.
Fitted values for three example data sets at 20 degrees of freedom.