

Medial Models of Populations of Nearly Tubular Objects*

Rohit R. Saboo¹, Joshua H. Levy³, Edward M. Chaney², and Stephen M. Pizer^{1,2}

¹ Computer Science and ² Radiation Oncology, The University of North Carolina at Chapel Hill
³ Formerly ¹, presently Morphormics Inc.

Abstract. Many structures in the world and several in the human body are nearly tubular in shape, i.e., have approximately circular cross-sections. Examples are portions of blood vessels, the colon and the whole head and neck. The method of modeling populations of slab-like objects as medial objects and segmentation using statistical shape and appearance models has been shown to be successful for several structures. However, the medial surface of a tubular object degenerates to a curve, and the statistics of even nearly tubular objects represented as slabs will typically be unstable. In this paper, we detail the representation, geometry and means of computing statistics of a population of nearly tubular medial models. We test our method on CTs of real rectums.

1 Introduction

In the human body, the blood vessels, the bronchi and the colon are examples of nearly tubular objects. Segmenting these structures is an important task in medical imaging and learning probability distributions on their populations is useful to segmentations [1]. Most of them can be thought of as a tube at the large scale with smaller scale changes understood as deviations from the tube. Some of these are shown in Fig. 1.

There are several definitions of tubular objects in the literature. Koenderink [2] defines a tube as the envelope of a set of spheres centered on a space curve. This is the medial definition of a tube. Swept surfaces and generalized cylinders model tubes as skeletal structures. A tubular generalized cylinder has a circular cross-section that may vary in size and have a possibly bent axis. In this paper, we discuss geometry and statistics for the former definition of tubes and then extend it to support deviations from the tubular structure. In Section 1.1 we discuss prior work done on modeling tubular objects. Section 1.2 presents the segmentation objective that serves as the driving problem.

1.1 Prior Work

Generalized cylinders, also known as generalized cones, were first proposed by Binford [3] with special instances studied extensively in computer vision. A straight homogenous generalized cone [4] is the surface obtained by sweeping a fixed cross-section along a straight axis while possibly scaling it, whereas a straight homogenous generalized cylinder may have a cross-section that can change shape. Huang et al. [5,6] discuss

* This work was supported under NIH grant number P01 EB02779. We would also like to thank other members of the Medical Image Display and Analysis Group at UNC Chapel Hill.

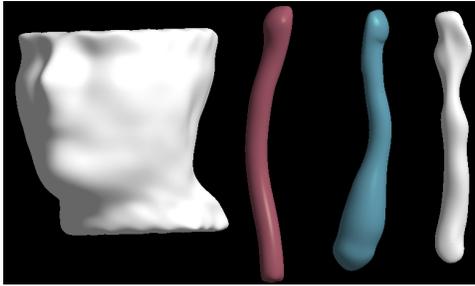


Fig. 1. Renderings of quasi-tube models fitted to different structures. From the left to right the structures are sections of the skin surface extracted from a 3D CT scan, carotid artery, internal jugular vein and the pharynx.

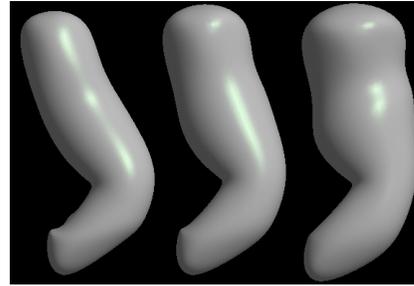


Fig. 2. mean model of a rectum (center) deformed by ± 1.5 standard deviations along the first mode of variation (left and right), which resembles the anatomical shape change due to gas.

generalized tubes that are constructed by sweeping a fixed cross-section along an axis with certain constraints. Terzopoulos et al’s [7] physically motivated deformable model uses image-based and regularity forces to deform the model. In our applications, it is difficult to get image-based forces to work due to poor contrast and noise.

Several of these center-line based methods are agnostic to the choice of the center-line. However, when modeling populations, there needs to be a principled way of finding the center-line so that no unwanted variation introduced due to the modeling process is reflected in the statistics. O’Donnell et al. [8] discuss a novel method of generalized cylinders that works around part of this issue by starting with a base cross-section that may be anisotropically scaled. Although they use only two scaling parameters, their method can be extended to produce arbitrary scaling. Further, they allow for local deformations of the cross-section by a spline function on the surface. However, they have not discussed any method to compute statistics of their structures.

The subtle differences between a skeletal and a medial axis are sometimes overlooked; the two are used interchangeably in some of the cited work. Having a true medial axis representation overcomes the issue of finding a unique center-line. The class of generalized cylinders whose medial axis is a curve is restricted to those with a circular cross-section. A non-circular cross-section results in a 2D medial surface.

A generalized cylinder and a structure with a well-defined medial axis are closely related. When we sweep a constant circular cross-section along a curved axis, the curved axis is the medial axis. However, if we sweep a non-constant circular cross-section along a curved axis, then such a structure may not have a curve as its medial locus, though some such objects (generated by sweeping spheres of varying sizes) will have a curve as the medial locus. Even when a generalized cylinder does not have a curve as its medial axis, it is useful to find an approximate medial axis for that object.

A quasi-tubular object can be thought of as a structure that is modeled as deviations from a tubular object. In the general case, it is a structure with a cross-section that may not be close to circular but does not vary much along the axis.

A sweep of a cross-section along an axis may result in two adjacent cross-sections crossing each other near a sharp bend on the axis. Such instances of the generalized cylinder are illegal, restricting the range of permissible cross-sections and axes. Damon [9] has described a method in the swept surface paradigm using a shape operator that can be used to detect these illegal generalized cylinders.

Several alternative methods exist that focus on extraction of the center-line from image data. Examples include the cores methods (height ridges of medial strength) developed by Aylward et al. [10] and by Fridman et al. [11].

No specialized means of performing statistics have been developed for the generalized cylinders and swept surface models discussed above, so they are best suited for modeling *individual* quasi-tubular objects versus *populations* of them. Such statistical descriptions on populations are useful if objects are to be segmented from images in which they have low contrast at their boundaries. There has been some work on modeling tubes with the help of a statistical shape model. The generalized stochastic tubes developed by Huang et al. [12] aid in the segmentation of blood vessels but are specialized for this application. De Bruijne et al. [13] have adapted the method of Active Shape Models with center-line based methods.

With statistical shape models a special concern is their robustness against the number of training samples, since in medicine these training samples can be very expensive. As mentioned by Joshi et al. [14], the orientation of the narrow medial sheet of objects with a nearly circular cross-section is sensitive to small changes in the boundary and will result in a population with broad variation. By avoiding this variability, the method we describe uses statistical shape models for which the probability estimation is particularly robust against the number of training samples.

1.2 The Driving Problem: Segmentation of Quasi-tubes

The method of segmentation via posterior optimization of m-reps developed by Pizer et al. [15] has been successful in dealing with slab-shaped objects with a lot of variability and poor contrast. We develop a new method that draws on the ideas from these methods but represents a tube-like object with a discretely sampled medial space curve and then models quasi-tubes as deviations from these tubes.

The segmentation method can be divided into two parts: training and the actual segmentation itself. During training, a rough m-rep model of the object is allowed to vary inside an optimizer that favors smooth models with a regularly spaced discrete medial mesh and that match well with the image data. The resulting models are known as training or fitted models.

These training models are then statistically analyzed. The variation in the shape space of the models is studied using Principal Geodesic Analysis (PGA), developed by Fletcher et al. [16], which is a variation of Principal Component Analysis (PCA) suited for non-linear spaces. The result is a mean shape $\bar{\mathbf{m}}$ and a prior $p(\mathbf{m})$ for segmentation. At the same time, the region around the object is divided into small parts and the distribution of intensities in each region is studied with the help of local region intensity quantile functions, developed by Broadhurst [17] and Stough [18]. We then apply PCA on these quantile functions to produce a likelihood function $p(\mathbf{I}|\mathbf{m})$ for segmentation.

When an image is to be segmented, the mean model is placed close to the real organ with the help of landmarks or manually. The model is allowed to deform along its principal modes of variation in an optimizer that favors likely shapes and intensity distributions around and within the object. The objective function maximized is the weighted sum of $\log p(\mathbf{m})$ and $\log p(\mathbf{I}|\mathbf{m})$ with the weights chosen to make the two terms have equal variance. This is a variant of the method of posterior optimization.

The remainder of the paper is organized as follows. In Section 2 we describe the representation and geometry for tubular medial models. In Section 3, we describe the way in which we estimate probability distributions on these models. We then describe the modeling of the deviations from a tubular to a quasi-tubular model in Section 4. Section 5 gives more details of our training and segmentation approaches. Finally in Section 6, we test our method on real data obtained from CTs of rectums and provide both quantitative and qualitative results on the same.

2 Medial Models for Tubes

A first order tube m-rep is a continuous space curve with a cone placed at every point along the curve. The axis of the cone is tangential to the space curve at the tip of the cone. Sweeping the edges of the cone bases gives the boundary of the modeled object, which is orthogonal to the rays from the cone tip to the cone base. The cones may have a half cone angle greater than $\pi/2$ but less than π . They are not allowed to intersect each other. Damon [19] has provided us with tools that can be used to measure local self-intersection (folding) of the object implied by the medial surface of a slab m-rep. In Section 2.1 we adapt these tools to do the same for tubular m-reps.

In practice, we represent the medial model of a tube by discretely sampling the space curve of cone tips. Each sample, shown in Fig. 3, is called an atom. Associated with each sample is its position in space, $\mathbf{p} = (x, y, z)$, and a cone with its tip positioned on the sample. The cone in turn is represented by its bisector, $\hat{\mathbf{U}}_0 = (U_{0,1}, U_{0,2})$, the half cone angle, θ , and the length of its inclined surface rays, r . The bisector of the cone always points along increasing arc length. Thus, the bisector points in the same direction when θ changes across $\pi/2$. To keep the discrete samples regularly spaced, while developing the models, we impose a penalty, called *irregularity penalty*, on the model that penalizes atoms moving away from the average of its neighbors.

The atoms at the two ends of the chain have an additional parameter describing the curvature of the cap at that end. However, when we are modeling open tubes, the end-atoms don't have any special properties and are just like any other atom.

A continuous medial curve $\gamma(u)$ is interpolated from these atom positions and cone bisector vectors with the help of piecewise cubic Hermite splines. The cone bisector vectors, scaled by the mean of the distance between the position of the atom and its two neighbors, are used as the tangents in the Hermite interpolation.

To resolve the rotational symmetry, we have a parameter ϕ that rotationally orients the entire tube along its length. One of the atoms in the tube is designated as a base atom. Usually this atom is close to a feature in the object that can help fix the rotational orientation. This feature can be an anatomic entity such as a part of a bone, a certain neighboring organ or tissue that can be easily identified in the entire population.

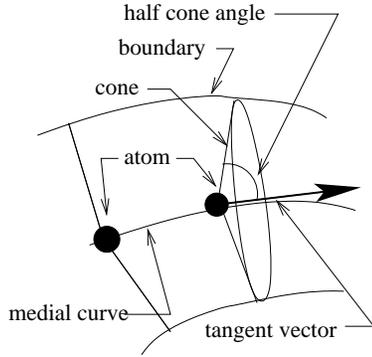


Fig. 3. Representation of a tube atom

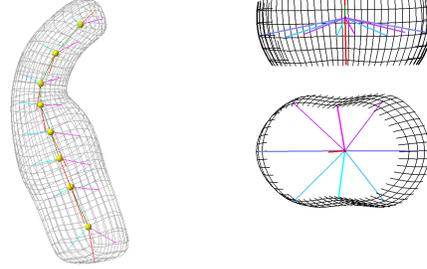


Fig. 4. A mean model of a rectum from one of our studies showing the medially implied surface as a wireframe.

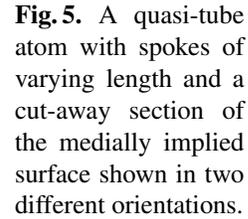


Fig. 5. A quasi-tube atom with spokes of varying length and a cut-away section of the medially implied surface shown in two different orientations.

Whenever the cone for this atom is rotated around its bisector, all the other cones are sympathetically rotated. This is needed for correspondences that depend upon the position along the circumference of the tube such as those required for quasi-tube statistics.

The surface implied by the tubular medial model is called as the *medially implied surface* and generated as follows. The circles at the base of the cones are discretely sampled. The medially implied surface is the fine and smooth mesh obtained by interpolating the set of surface points generated from the cones of all the atoms. The medially implied surface for a tube model of a rectum is shown in Fig. 4.

In Section 2.1, we highlight a geometric operator that is useful for detecting models with self-intersecting surfaces and interpolating between atoms in the medial axis. In Section 2.2, we describe a geometric measure, which we call curviness, that is useful in avoiding wavy medial axes.

2.1 Geometry on tubular models

Consider the medial atom cone to be composed of a set of spokes U extending from the tip to the base. These spokes are parameterized by the arc angle $\phi \in [0, 2\pi)$ on the circumference of the base made with respect to a zero point on the circumference. Define the 1D radial shape operator for a tube as

$$S_{rad}(\phi) = -proj_U \left(\frac{\partial \mathbf{U}(\phi)}{\partial s} \right), \quad (1)$$

where $proj_U$ is projection along the spoke vector U , \mathbf{U} is the corresponding unit spoke vector and s is an arc-length parametrization of the medial curve. Note that S_{rad} and r are also functions of u but we will ignore this for brevity. The derivative may be written in the form

$$\frac{\partial \mathbf{U}(\phi)}{\partial s} = a \cdot \mathbf{U} - \kappa_{rad} \cdot \gamma'(s), \quad (2)$$

where κ_{rad} is the principal radial curvature. Thus $S_{rad}(\phi) = \kappa_{rad}(\phi)$.

An important use of this shape operator is in detecting models that are *illegal*, i.e., some of the spokes are crossing each other and the surface has folded onto itself. Damon [19] shows us that spokes will cross each other if and only if $r \times \kappa_{rad} > 1$. Unlike Damon’s shape operator, our shape operator is a function of the angle ϕ . However, it suffices to evaluate this for the angle corresponding to the direction of the curve normal.

The model is illegal if $\exists u$ s.t. $r \times S_{rad} \geq 1$. This condition can be relaxed as a penalty more suitable for an optimizer, which expects a continuous objective function, that is the p -norm of the individual measurements $\max(0, r \times S_{rad} - \beta)$ along the tube for a certain threshold $\beta \in [0, 1)$. Experimentally, we have seen that $p = 6 - 8$ and a threshold of $0.8 - 0.9$ produces good results. Larger values of p are sometimes useful when we wish to make the aggregate measure more sensitive to local problems.

The S_{rad} operator can also be used to interpolate between two consecutive atoms on the same medial manifold. Han et al. [20] have used this operator to interpolate atoms in 2D medial manifolds. We have adapted the interpolation method to generate interpolated atoms for tubes. An important use is in interpolating atoms to improve correspondence between models.

2.2 Geometric Penalty - Curviness

A wavy medial sheet results in crooked-looking models. It also necessitates the use of significantly more samples in the computation of the illegality penalty. Penalizing fitted m-reps according to an aggregate measure of curviness can alleviate these problems. Apart from being anatomically improbable, there is nothing wrong with crooked-looking models. We define *curviness* by the p -norm of the total curvature over the entire length of the medial curve.

$$C = \left(\frac{1}{n-1} \int_0^{n-1} T^p du \right)^{1/p}, \quad (3)$$

where the total curvature T is related to the geometric curvature κ and the geometric torsion τ of the interpolated medial axis $\gamma(u)$ by $T = \sqrt{\kappa^2 + \tau^2}$. Different values of p between 2 and 10 are appropriate depending upon how much we need to emphasize individual sharp bends. However, larger values of p create steeper gradients that may not be good in an optimizer.

The curviness measure is rotation and translation invariant. Dividing by the arc length or the average inter-atom distance makes it scale-invariant too.

3 Shape Space and Statistics

A tube atom \mathbf{m} can be represented by the tuple $\mathbf{M} = \langle \mathbf{P}, \hat{\mathbf{U}}_0, \theta, r \rangle$. A tube consisting of n atoms can be represented by n such tuples concatenated together. Here the cone vertex \mathbf{P} belongs to the group \mathbb{R}^3 , the cone axis $\hat{\mathbf{U}}_0$ is a point on the two dimensional sphere S^2 , the half cone angle θ ranges from 0 to π and is related to the group \mathbb{RP}^1 , and

the cone length r belongs to \mathbb{R}^+ . All of the groups are Lie groups and except for \mathbb{R}^3 , they are not Euclidean manifolds. A Lie group has a differentiable group operator, an inverse element and the identity element. Using the group operators and by the action of the inverse on an element of the group on the element itself, we obtain the identity element. The other things needed are a distance metric in this space and the ability to project back and forth between this shape space and a tangent plane. Several of the results presented in this section are similar to those worked out by Fletcher et al. [16] for slabular m-reps.

The path with the shortest distance between two points in a manifold is known as the geodesic between them. The length of this path is called the geodesic distance. If we have a suitable mapping between the manifold and a tangent space, then this mapping can be used to measure the geodesic distance on the tangent plane. The map that takes us from the manifold to the tangent space is called as the Logarithmic chart and the reverse map is known as the Exponential chart. The maps of the individual components are given by the following equations. The map for the atom is simply the direct product of these maps. The maps for the position are identity functions and that for the radius are logarithm and exponential functions. In the following equations, $\theta \in (0, \pi)$, $\mathbf{U} = \langle u_1, u_2, u_3 \rangle \in \mathbb{S}^2$ and $\mathbf{U}' = \langle u'_1, u'_2 \rangle \in \mathbf{T}_{(0,0,1)}\mathbb{S}^2$.

$$\text{Log}(\theta) = \tan\left(\theta - \frac{\pi}{2}\right), \quad (4)$$

$$\text{Exp}(\theta') = \tan^{-1}(\theta) + \frac{\pi}{2}, \quad (5)$$

$$\text{Log}(\mathbf{U}) = \left(u_1 \frac{\alpha}{\sin \alpha}, u_2 \frac{\alpha}{\sin \alpha}\right), \quad (6)$$

$$\text{Exp}(\mathbf{U}') = \left(u'_1 \frac{\sin |u'|}{|u'|}, u'_2 \frac{\sin |u'|}{|u'|}, \cos |u'|\right), \quad (7)$$

$$(8)$$

where $\alpha = \cos^{-1}(u_3)$ and $|u'| = \sqrt{u'^2_1 + u'^2_2}$. All the above maps are taken centered at the identity element. The identity element for the group \mathbb{R}^3 is 0, for \mathbb{R}^+ it is 1 and for \mathbb{S}^2 it is the point $(0, 0, 1)$. For the group of θ , the identity element is $\pi/2$. To obtain the chart for a tangent plane centered at a point \mathbf{m} different from the identity element, we need to apply the inverse of that point to the element in order to move the tangent plane to the identity element. We use $\text{Log}_{\mathbf{m}}$ and $\text{Exp}_{\mathbf{m}}$ as the notation in this case.

To make the units of all the components commensurate in the Log map, we multiply the unitless quantities with the mean radius taken over all the corresponding atoms in the population. The geodesic distance is then defined as the norm of the difference of these normalized atoms projected into the tangent space.

We define the Fréchet mean μ of a set of atoms as the one that minimizes the sum of squared geodesic distances from all the corresponding atoms from the population.

$$\mu = \underset{\mathbf{m} \in \mathbf{M}}{\text{argmin}} \sum_i |\text{Log}_{\mathbf{m}}(\mathbf{m}_i)|^2. \quad (9)$$

We then compute the Log map of all these atoms and project them on the tangent plane centered on the mean. We do PCA on these projected atoms and keep the first

few modes that represent more than 90% of the total variation. Fig. 2 shows the shape variation along the first mode of variation of the rectum from our study.

4 Quasi-tubes

In this section, we show how we model the deviations from a tube to a quasi-tube.

Several objects in the real world can be modeled as deviations from a tubular structure. Take the head and neck for example. Start with a cylinder, make the cylinder bulge out in certain regions to produce features such as the nose and the lips, and make it cave in to produce the eye sockets and other cavities. Sections of the colon, blood vessels, the bronchial tree and many other organs in the body can be thought of as quasi-tubes.

A tube atom is represented by a cone of spokes. All of the spoke ends lie in a single plane, α , forming a cross-section of the tube. Our approach is to change the cross-section in this plane. This can be accomplished by inclining each spoke in the plane between it and the axis in such a way that the tips of all the spokes continue to lie in the plane α . Changing the cross-sectional shape in this way makes the computation of the shape operator in the circumferential direction straightforward and the component along the axis is still given by the 1D shape operator defined by equation 1 as \mathbf{U} is a function of the spoke angle ϕ .

The above method results in varying length spokes in an atom. Therefore the structure is no longer Blum medial but skeletal. Damon's proofs concerning the use of the shape operator S_{rad} in illegality measurements are valid for skeletal structures too.

The change in the length of the spokes is represented as a multiplicative parameter that belongs to the group \mathbb{R}^+ . The Log and Exp maps for this parameter have the same form as that for the radius r .

The segmentation step is divided into two scales. At the large scale, a mean of each quasi-tube atom across all population samples is used. This gives a large scale model with each cross-section having a different shape. In the small scale, the individual quasi-tube atoms are allowed to vary to allow for small changes in each cross-sectional shape.

5 Training and Segmentation

We have discussed how we geometrically and statistically model tubes and quasi-tubes. We now discuss their training and the appearance model we use.

To fit quasi-tube models to training cases and derive a probability distribution on that object we follow the general approach described in Merck [21]. Expert outlines designating the target object in several 3D images are converted into binary images. An initial model that somewhat resembles the object is created. Sometimes landmarks are added to initialize our models. The parameters for these models are then varied inside an optimizer. The objective function is set to a sum of a combination of image match and geometric penalty terms. The image match we use is the average of the sum of squared distances between a point on the object's surface and the closest point on the expertly outlined object. In certain places such as the crest, the reverse distance is used as the original distance is artificially low. When landmarks are used, the distance

between them and the corresponding positions on the model are taken into account. The geometric penalty terms consist of a combination of the irregularity, S_{rad} and the curviness penalties mentioned in Sections 2, 2.1 and 2.2 respectively. The weights for the geometric penalties are relaxed proportionally to the quality of the fit of the model.

We then compute the statistical shape model from all the models in the population. The log map of the atoms is computed. The spoke deviations are represented as a multiplicative parameter that belongs to the group \mathbb{R}^+ . The Log map of the deviations is the Log map for the group \mathbb{R}^+ as described in Section 3. We then compute the mean across the entire population to give us the mean shape. The atoms are projected at the tangent plane centered on the mean. The statistical shape model is computed as described in Section 3 except that the spoke deviations are only used in the computation of the mean and not in the computation of the modes of variation of the shape. These comprise the large scale shape variation. The spokes are allowed to vary in the smaller scale. This separation of variation at two different scales gives more stable statistics.

At segmentation time, the appearance model we use has been developed by Broadhurst [17] and Stough [18]. The object’s surface is divided into several regions. An intensity histogram for each region is computed and converted into quantile distributions. The distribution of these quantile functions is then analyzed with the help of Principal Component Analysis (PCA). As tubes can be arbitrarily cut-off, we optionally allow for the ends of the tubes to not contribute to these appearance models.

During segmentation, we use the method of posterior optimization where the statistical shape model gives the prior and the appearance model gives the likelihood.

6 Application and Results

Segmenting rectums from 3D CT scans is important for adaptive radiotherapy treatment for prostate cancer. It is important for the patient’s health and quality of life that the rectum does not receive too much radiation. The rectum changes shape a lot from day to day due to the presence of gas and faeces pushing the abutting prostate around. The presence of gas also creates a large variation in the intensity distributions and necessitates special handling of the interior intensity distributions in the rectum.

In our experiment the data comes from different days of several patients. For each patient, the statistical shape model is built with the help of data from just that patient. The study is done in a leave-one-out manner, i.e., we build the statistical shape model for each day of a patient by using data from the other days. This is clinically not possible but it suffices for the purposes of our study. In the availability of more training data, one can incorporate cross-patient statistics to get over this limitation.

We obtained 3D CT images with a resolution of $0.98 \times 0.98 \times 3 \text{ mm}^3$ of 5 patients with 13–18 images for each patient for a total of 79 images. We then trained quasi-tubular medial models on manual segmentations of rectums from this data. For our training, we used a hierarchical training process where we first fit a coarsely sampled tubular medial model to the data. We then allowed the individual spokes to vary, generating quasi-tubular models. After this we subdivided the model by generating interpolated atoms halfway between previously existing atoms and then refined this finely

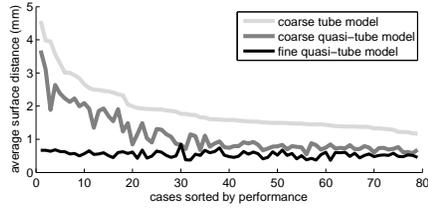


Fig. 6. Average surface distance of the training models at different stages.

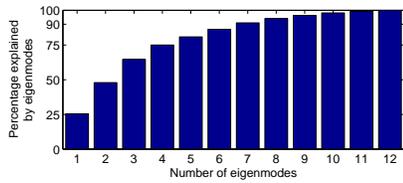


Fig. 7. Graph showing cumulative shape variation captured by the first few eigenmodes: the first 7 modes capture about 90% of the shape variation.

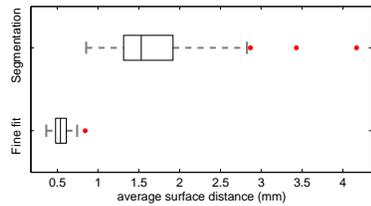


Fig. 8. Boxplot of distribution of segmentation results versus training results. The box goes from 25% to 75% quartile with a line for the median. The two lines at the end go 1.5 times the intra-quartile range (except when there is no more data). The red dots are outliers.

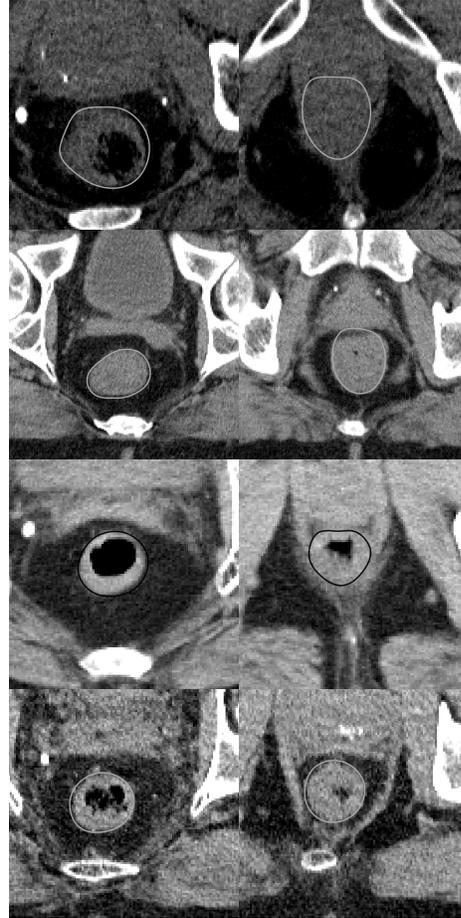


Fig. 9. Each row shows the outline (white or black) of our segmentation on two different axial slices of the same image. Note the poor contrast in the slices in the right column. The slices on the right are inferior with respect to those on the left. The first three rows are typical results and the last row is one of the better segmentations.

sampled model to fit the binary image better. Further subdivision yielded only marginal improvement.

The quantitative results of our training process are shown in Fig. 6. The median average surface distances for the coarse tube, coarse quasi-tube and the fine quasi-tube

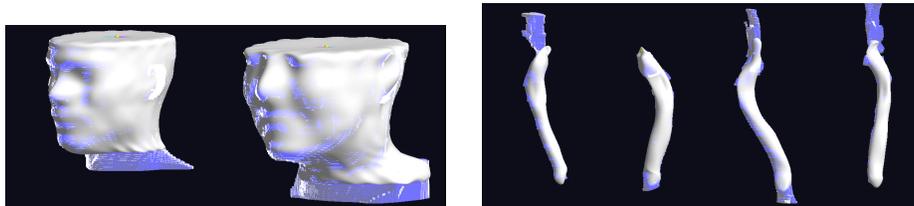


Fig. 10. Quasi-tubular medial models (in white) fit to a section of the head and neck skin surface (left) and the pharynx (right) vs. the manually segmented structures (in translucent blue)

models are 1.58, 0.84 and 0.54 mm with standard deviations of 0.70, 0.62 and 0.09 mm. We can conclude that we can model rectums fairly well and in cases where the tube model is lacking, the quasi-tubular model does a much better job.

In Fig. 7, we show the cumulative shape variation captured by the first few eigenmodes for a quasi-tubular model trained over the data of one of the patients. In about 7 modes, we can effectively capture most of the shape variation. In Fig. 2, we show the variation of the rectum along the first principal direction, which resembles anatomical shape changes due to filling of gas.

During segmentation, we used a similar hierarchical approach. The model was first initialized semi-automatically. We then segmented the model in the shape space of the coarse models. The result was then used to drive the segmentation in the shape space of the finely sampled trained medial models. In Fig. 8, we show the segmentation results versus the quality of fine training fits for all the patients. The median of the segmentation results is 1.53 mm. The segmentation results for a few of these is shown in Fig. 9. Notice the complete lack of contrast in some of the inferior slices.

We have also fit quasi-tubular models to various structures of the head and neck. In Fig. 10, we show quasi-tubular models fit to sections of the skin surface and the pharynx. The expertly contoured outline is shown in translucent blue. The average surface distance for the fitted models of the skin surface, common carotid artery, internal jugular vein and the pharynx are 1.22, 0.8, 1.3 and 1.13 voxels.

7 Conclusion

We have developed a new method for modeling populations of nearly tubular objects as a tubular medial model with deviations from the perfect tubularity described by local changes and demonstrated the same over real world rectum data. We have seen that this representation is also effective for elongated objects with distinctly noncircular but slowly varying cross-sections.

Further, we show how we can study population variations by doing statistics in the non-linear space in which these quasi-tubular medial models lie. We also discuss the radial shape operator that is needed for studying medial geometry on these models.

The rectum is a challenging organ due to immensely varying shape and poor image contrast. Several structures in the head and neck – skin surface, pharynx, jugular veins

and carotid arteries – provide difficult modeling challenges. We have shown that our models can be trained to within sub-voxel accuracy and give reasonable segmentations.

References

1. Cootes, T.F., Hill, A., Taylor, C.J., Haslam, J.: The use of active shape models for locating structures in medical images. In: IPMI. (1993) 33–47 [1](#)
2. Koenderink, J.J.: Solid Shape. MIT Press (1990) [1](#)
3. Binford, T.O.: Visual perception by computer. In: Proceedings of the IEEE Conference on Systems and Control. (1971) [1](#)
4. Ulupinar, F., Nevatia, R.: Shape from contour: straight homogeneous generalized cones. In: Proceedings of Computer Vision. (Dec 1990) 582–586 [1](#)
5. Huang, Q., Stockman, G.: Model-based automatic recognition of blood vessels from mr images and its 3D visualization. In: Proceedings of the IEEE International Conference on Image Processing. Volume 3. (Nov 1994) 691–695 [1](#)
6. Huang, Q., Stockman, G.C.: Generalized tube model: recognizing 3D elongated objects from 2D intensity images. In: CVPR. (Jun 1993) 104–109 [1](#)
7. Terzopoulos, D., Witkin, A., Kass, M.: Constraints on deformable models: recovering 3D shape and nongrid motion. *Artif. Intell.* **36**(1) (1988) 91–123 [2](#)
8. O’Donnell, T., Boulton, T., Fang, X.S., Gupta, A.: The extruded generalized cylinder: a deformable model for object recovery. In: CVPR. (21–23 Jun 1994) 174–181 [2](#)
9. Damon, J.: Swept regions and surfaces: Modeling and volumetric properties. *Theoretical Computer Science* **In Press, Accepted Manuscript** [3](#)
10. Aylward, S., Pizer, S., Eberly, D., Bullitt, E.: Intensity ridge and widths for tubular object segmentation and description. *MMBIA* (1996) 131 [3](#)
11. Fridman, Y., Pizer, S.M., Aylward, S., Bullitt, E.: Segmenting 3D branching tubular structures using cores. In: MICCAI. (November 2003) 570–577 [3](#)
12. Huang, Q., Stockman, G.: Generalized stochastic tube model: tracking 3D blood vessels in mr images. In: ICPR. Volume 2. (9–13 Oct 1994) 156–160 [3](#)
13. de Bruijne, M., van Ginneken, B., Viergever, M., Niessen, W.: Adapting active shape models for 3D segmentation of tubular structures in medical images. *IPMI* **2732** (2003) 136–147 [3](#)
14. Joshi, S., Pizer, S., Fletcher, P., Yushkevich, P., Thall, A., Marron, J.: Multiscale deformable model segmentation and statistical shape analysis using medial descriptions. *IEEE Transactions on Medical Imaging* **21**(5) (May 2002) 538–550 [3](#)
15. Pizer, S., Fletcher, T., Fridman, Y., Fritsch, D., Gash, A., Glotzer, J., Joshi, S., Thall, A., Tracton, G., Yushkevich, P., Chaney, E.: Deformable m-reps for 3D medical image segmentation. *IJCV - Special UNC-MIDAG issue* **55**(2) (Nov-Dec 2003) 85–106 [3](#)
16. Fletcher, P.T., Lu, C., Pizer, S.M., Joshi, S.: Principal geodesic analysis for the study of nonlinear statistics of shape. *Medical Imaging, IEEE Transactions on* **23** (2004) 995–1005 [3, 7](#)
17. Broadhurst, R., Stough, J., Pizer, S., Chaney, E.: A statistical appearance model based on intensity quantiles histograms. *ISBI* (2006) [3, 9](#)
18. Stough, J., Broadhurst, R.E., Pizer, S.M., Chaney, E.L.: Regional appearance in deformable model segmentation. In: IPMI. (2007) [3, 9](#)
19. Damon, J.: Determining the geometry of boundaries of objects from medial data. *International Journal of Computer Vision* **63** (June 2005) 45–64 [4, 6](#)
20. Han, Q., Pizer, S.M., Damon, J.N.: Interpolation in discrete single figure medial objects. In: Proceedings of Computer Vision and Pattern Recognition. (2006) 85 [6](#)
21. Merck, D., Tracton, G., Saboo, R., Levy, J., Chaney, E., Pizer, S., Joshi, S.: Training models of anatomic shape variability. *Medical Physics* **35**(8) (2008) 3584–3596 [8](#)