

# Segmentation of Tubular Objects by Posterior Optimization of M-reps

Rohit R. Saboo

rohit@cs.unc.edu

Christina Villarruel

crvillar@cs.unc.edu

Edward M. Chaney

chaney@unc.edu

Julian G. Rosenman

julian\_rosenman@med.unc.edu

Stephen M. Pizer

pizer@cs.unc.edu

## Abstract

*Many organs in the human body are nearly tubular in shape, i.e., have circular cross-sections. Examples are portions of blood vessels and the colon. The method of segmentation of slab-like organs with the help of statistics on medial models has been shown to be successful for several organs. However, the medial surface of a tubular object degenerates to a curve, and the statistics of even nearly tubular objects will typically be unstable. Traditional medial models cannot handle this degeneracy. In this paper, we detail the representation, geometry and means of computing statistics on tubular medial models. We test our method on synthetic tubes with shapes sampled from the distribution of a rectum and CTs on carotid arteries and real rectums.*

## 1. Introduction

Many objects in the human body are roughly in the shape of a tube. Some of these include sections of rectum, blood vessels and the bronchial tree. In our work, a tube refers to an object that has a circular cross-section that may vary in size and have a possibly bent axis. Segmenting these objects is an important task in medical image analysis. Most of these objects can be modeled as a set of small scale local deviations from a tube. In this paper, we tackle the problem of finding the larger scale tubular structure.

Segmenting the vascular tree has been a very important problem in medical imaging. Most existing work on segmenting nearly tubular objects uses a tube following paradigm, where the tube is segmented a piece at a time. Fridman et. al. [3] use a medial model for a tube. However, their model is restricted to tubes with cross-sections that do not vary too quickly. They use a model called ‘the core’, which is an m-rep extracted as a height ridge of a medial strength function on medial atoms. The medial atoms are characterized by a point and a set of equi-length spokes emanating from the point in the shape of a cone. They compute a medialness measure to follow the medial axis of the tube

after the core has been initialized. These methods, however, are tailored to suit the vascular and the bronchial structures, which have much better contrast than rectums and are also much narrower compared to rectums. Rectums on the other hand can be very variable in shape due to the presence of gas and other non-tissue material. Also the contrast in the image is not very good.

The method of segmentation via posterior optimization of m-reps developed by Pizer et. al. [5] has been very successful in dealing with shapes with a lot of variability and poor contrast. However, the m-reps used to date use a medial surface to represent a bent variable width slab. In the case of a tube with a perfectly circular cross-section, the medial surface degenerates to a space curve. For most real tubes, the medial surface is very narrow. The orientation of a narrow medial sheet is very sensitive to small changes in the boundary and will result in very variable statistics. We develop a new method that draws on the ideas from those used for slab m-reps but represents a tube-like object with a discretely sampled medial space curve.

The segmentation method can be divided into two parts: training and the actual segmentation itself. A rough model of an m-rep for the object is developed manually. An expert manually segments at least a dozen training images. The resulting manual segmentations are used to develop the m-rep model for the object. Optionally, a set of landmarks are also identified on the object. These correspond to places on the surface of the object that can be easily located. We then vary the parameters of the m-rep model inside an optimizer that minimizes the distance between the m-rep implied boundary and the actual segmentation as well as the distance between the actual landmarks and the corresponding positions on the surface of the m-reps. At the same time, we try to ensure that the m-rep model appears smooth and remains legal. The precise definitions of legality, smoothness and distance will be covered in the later sections of this paper. The resulting set of models are known as trained or fitted models.

These trained models are then analyzed statistically. The variation in the shape space of the models is studied using

Principal Geodesic Analysis (PGA), which is a variation of Principal Component Analysis (PCA) suited for non-linear spaces. At the same time, the region around the object is divided into smaller parts and the distribution of intensities in each region is studied with the help of local region intensity quantile functions, developed by Broadhurst [1] and Stough [7]. We then apply PCA on these quantile functions.

When a target 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 under an optimizer that penalizes improbable shapes and intensity distributions around the object. This is known as the method of posterior optimization

Similarly for tubes, we start with a tubular medial model. The model is trained with the help of manually segmented images. Shape and intensity statistics are obtained and then these probability distributions are used to drive the segmentation.

In section 2 we describe the representation, geometry and methods for computing shape statistics on tubes. Finally, in section 3 we present some results on the segmentation of rectums.

## 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 its tip. Sweeping the edges of the base of these cones defines the boundary of the modeled object. The cones may have a half cone angle greater than  $\pi/2$  but less than  $\pi$ . They are not allowed to intersect each other. Damon has developed a Shape penalty measure for slab m-reps that penalizes such intersecting cones. We have similarly developed a shape penalty measure for tubular m-reps.

In practice, we represent the medial model of a tube by discretely sampling the space curve of cones. Each sample is called an atom and is shown in Fig. 1. 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,  $\theta$ , and the length of its inclined surface,  $r$ . The bisector of the cone co-incides with the tangent to the medial curve. To keep the discrete samples from clustering or moving away, we impose a regularity penalty on the object that penalizes if an atom moves away from the average defined by its neighbors.

The atoms at the two end of the chain have an additional parameter describing the curvature of the cap at that end. These atoms are called end-atoms.

The continuous medial curve  $\gamma(u)$  is interpolated from these atom positions and cone bisector vectors with the help




Figure 1. Representation of a tube atom




Figure 2. A mean model of a rectum from one of our studies showing the medially implied surface as a wireframe.

of piecewise cubic Hermite splines. The cone bisector vector is scaled by the mean of the distance between the position of the atom and its two neighbors.

Further, to resolve the rotational symmetry, we have a parameter  $\phi$  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 organ that can fix the rotational orientation. Whenever the cone for this atom is rotated around its bisector, all the other cones are sympathetically rotated.

The circle at the base of the cone is discretely sampled. The set of surface points generated from the cones of all the atoms are interpolated to produce a fine and smooth mesh resulting in a medially implied surface. The medially implied surface for a tube model of a rectum is shown in Fig. 2. Each position on the surface of the tube is uniquely represented by a set of co-ordinates and this co-ordinate system provides correspondence between deformed versions of the tube.

### 2.1. Shape Penalty: $rS_{rad}$

We wish a penalty against locally self-intersecting shapes. Ideally this penalty should be computed over the entire space curve, but this would be computationally very expensive. Thus we compute the shape penalties only over the atom positions. Since this could still allow the model to become illegal, we impose a penalty on the curviness of the

space curve, as described later.

The penalty against local self-intersection is based on a one-dimensional shape operator for tube medial models. This operator is an adaptation of the shape operator for ordinary medial models developed by Damon et. al. [2]. Han et. al. [4] have used Damon's operator to detect shape illegalities in ordinary medial models.

Consider the medial atom cone to be composed of a set of spokes  $U$  extending from the tip to the base. We can parameterize these spokes by the arc angle  $\omega$  on the circumference of the base made with respect to a zero point on the circumference. It has a range of  $[0, 2\pi)$ . Then the radial shape operator is defined as

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

where  $proj_U$  is projection along the spoke vector  $U$ ,  $\mathbf{U}$  is the corresponding unit spoke vector and  $u$  is an arc-length parameterization of the medial curve. The derivative may be written in the form

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

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

Surfaces will be illegal and fold onto themselves if any of these spokes cross. Damon shows that spokes will not cross each other if  $r \times \kappa_{rad} < 1$ . Unlike Damon's shape operator, our shape operator is a function of the angle  $\omega$ . However, we need not evaluate this for every spoke. A simple heuristic based on vector mathematics can be used to find out the direction of the maximum bending, and we need evaluate  $S_{rad}$  for only this angle.

In practice, the derivatives are computed by taking finite differences of geodesically interpolated atoms. Interpolation of medial atoms is briefly described in the next section. Also this  $S_{rad}$  operator is evaluated for several positions along the medial axis and not just at the atoms. The norm of all of the  $S_{rad}$  values obtained is reported as the shape penalty. The ordinary 2-norm has been experimentally shown to produce good results.

## 2.2. Geometric Penalty - Curviness

To penalize medial axis curviness, we define a curviness measure,  $C$  := the  $p$ -norm of the total curvature sampled at several places on the medial axis,  $\gamma(u)$ . The total curvature is related to the curvature  $\kappa$  and the torsion  $\tau$  of the medial

axis, as follows.

$$T = \sqrt{\kappa^2 + \tau^2} \quad (3)$$

$$\kappa = \frac{|\gamma'(u) \times \gamma''(u)|}{|\gamma'(u)|^3} \quad (4)$$

$$\tau = \frac{[\gamma'(u) \cdot \gamma''(u) \cdot \gamma'''(u)]}{|\gamma'(u) \times \gamma''(u)|^2} \quad (5)$$

$$C = \left( \int_0^{n-1} T^p dt \right)^{1/p}, \quad (6)$$

where  $n$  is the number of atoms in the tube. Recall that  $\gamma(u)$  is obtained by Hermite interpolation. This relation for curviness may not have an analytical form, so we evaluate it with the help of numerical methods. Also, the value of  $p$  is usually chosen as 2. However, if one wishes to penalize the existence of any sharp bend or kink, it may be increased. Beyond a value of 6, the penalty becomes too sensitive for any meaningful use.

## 2.3. Shape Space and Statistics

Before computing statistics in the shape space of tubular  $m$ -reps, the somewhat arbitrary ends of a tube model need to be normalized. This action eliminates the possibility of this random cut-off from becoming a dominant mode of variation in the shape statistics. After this normalization, the surface of the tube is parameterized and divided into different regions for the purpose of computing RIQFs.

A tube atom  $\mathbf{m}$  can be represented by the tuple  $\langle \mathbf{P}, \hat{\mathbf{U}}_0, \theta, r \rangle$ . A tube consisting of  $n$  atoms can be represented by  $n$  such tuples concatenated together. Here  $\mathbf{P}$  belongs to the Lie group  $\mathbb{R}^3$ ,  $\hat{\mathbf{U}}_0$  is a point on the two dimensional sphere  $S^2$ ,  $\theta$  ranges from 0 to  $\pi$ , and  $r$  belongs to  $\mathbb{R}^+$ . Except for  $\mathbb{R}^3$ , all the other groups mentioned belong to non-Euclidean spaces. We need to define a distance metric in this space and also be able to project back and forth between this shape space and a tangent plane.

The path with the shortest distance between two points in a non-Euclidean manifold is known as the geodesic between them. The length of this path is called the geodesic distance. To be able to compute the geodesic distance, we need to define the mappings between the shape space and the tangent space. The map that takes us from the shape space to the tangent space is called as the Logarithmic map and the reverse map is known as the Exponential map. The maps of the individual components are given by the following equations. The map for the atom is simply the direct




Figure 3. The figure in the center shows the mean model of a rectum of a patient obtained by training over 17 days. The top and bottom figures show the same rectum deformed by  $\pm 1.5$  standard deviations along the first mode of variation.

product of these maps.

$$\begin{aligned}
 \text{Log}(\mathbf{x}) &= \mathbf{x}, \\
 \text{Log}(\phi) &= \left( y_1 \frac{\alpha}{\sin \alpha}, y_2 \frac{\alpha}{\sin \alpha} \right), \\
 \text{where } \alpha &= \cos^{-1}(y_3), \\
 \text{Log}(\theta) &= \tan \left( \theta - \frac{\pi}{2} \right), \text{ and} \\
 \text{Log}(r) &= \log(r).
 \end{aligned} \tag{7}$$

$$\begin{aligned}
 \text{Exp}(\mathbf{x}) &= \mathbf{x}, \\
 \text{Exp}(\phi) &= \left( u_1 \frac{\sin|u|}{|u|}, u_2 \frac{\sin|u|}{|u|}, \cos|u| \right), \\
 \text{where } |u| &= \sqrt{u_1^2 + u_2^2}, \\
 \text{Exp}(\theta) &= \tan^{-1}(\theta) + \frac{\pi}{2}, \text{ and} \\
 \text{Exp}(r) &= \exp(r).
 \end{aligned} \tag{8}$$

Here,  $\langle y_1, y_2, y_3 \rangle$  represents a point on the unit sphere

and  $\langle u_1, u_2 \rangle$  is its projection on the tangent plane. Note that the exponential and logarithmic maps for  $\phi$  preserves distances from the point  $(1, 0, 0)$  on  $S^2$ . To make the units of all the components commensurate in the logarithmic 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.

Interpolation between two atoms is performed by interpolating in the tangent space and then projecting back to the manifold space. In this case, we must be careful that the logarithmic and exponential maps for one of the atoms is taken centered on the first atom.

Once we have a distance metric and these maps, we can define the Fréchet mean  $\mu$  of a set of atoms as the one that minimizes the sum of squared geodesic distances from all the corresponding atoms from the population.

$$\mu = \underset{\mu}{\operatorname{argmin}} \sum_i |\mathbf{m}_i - \mu|^2. \tag{9}$$

We then take the logarithmic 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 describe more than 90% of the variation. Fig. 3 shows the shape variation along the first mode of variation of the rectum from our study. It also shows the percentage of the variation explained by each of the eigenmodes. In this case we needed the first 8 eigenmodes.

## 2.4. End Atom Normalization

Organs in the shape of tubes in the human body are usually quite long in length. However, in applications such as radiotherapy, surgery and treatment planning, one is interested in a small section. During the step of manual segmentation for generating training images for the model, the cut-off may be somewhat arbitrary. If we compute shape statistics on models trained on these images, then the variable position of the ends may be a dominant mode of variation after PGA.

We compute the mean of the length of the medial curve between the first and the second atoms across all the training models. We then interpolate or extrapolate the first atom of each model so that it is at the mean distance from the second atom along the medial curve. This is also repeated for the other end. If  $\mathbf{m}_{0,i}$  and  $\mathbf{m}_{1,i}$  are the first and second atoms from the  $i^{\text{th}}$  model, then the new first atom,  $\mathbf{m}'_{0,i}$  is given

by the following relation.

$$\begin{aligned}
 l_i &= \int_0^1 \gamma_i(u) dt, \\
 \text{mean length } \bar{l} &= \sum_i \frac{l_i}{s} \\
 \mathbf{m}'_{0,i} &= \text{Exp} \left( \frac{\bar{l}}{l_i} \text{Log}(\mathbf{m}_{0,i}) \right), \quad (10)
 \end{aligned}$$

where  $s$  is the number of samples.

## 2.5. Regional Intensity Quantile Functions

Broadhurst [1] and Stough [7] developed the method of regional image intensity quantile functions. In this method, the surface is divided into small regions centered around each spoke. Image intensity quantile functions  $\mathbf{Q}_{x,k}$  are recorded for each region  $x$  in each training image  $k$ . A quantile function is the inverse of the cumulative density function on intensity. Because the space of such quantile functions typically forms a low-dimensional linear space, PCA yields a mean  $\mu_x$  and a few modes of variation  $\{\lambda_{x,i}, \mathbf{v}_{x,i}\}$ . The local image match function measures statistical distance between the observed quantile function  $\hat{\mathbf{Q}}_x$  in a region of the target image and the  $\mu_x$  computed. The local image match is thus

$$f_x(m, I) = \sum_i^n \frac{((\mathbf{Q}_x - \mu_x) \cdot \mathbf{v}_{x,i})^2}{\lambda_{x,i}} + \frac{|\mathbf{r}|^2}{\lambda_r}. \quad (11)$$

The first term is the Mahalanobis distance between the mean and the observed quantiles with only the first  $n$  modes of variation used. The second term accounts for the residue  $\mathbf{r}$  outside this PCA space.

## 3. Applications and Results

Adaptive radiotherapy is one of the common solutions to treat prostate cancer, which is a growing medical problem. Radiotherapy requires the precise delineation, also known as segmentation, of the organ that is being irradiated. The neighboring organs also need to be segmented so that the radiation beams can be configured to avoid any sensitive surrounding tissue and at the same time deliver the required dose to the organ being treated.

The prostate is abutted by the bladder on one side and the rectum on the other. Its shape and position are influenced by both these organs. There has been success reported in segmenting the bladder-prostate complex in Pizer et. al. [6]. The rectum is an object that is like a tube in shape, so it is not well segmented via an m-rep slab.

We trained some models on real rectum data and obtained a mean model with shape statistics from these trained

models. These statistics were then sampled to produce 99 synthetic rectum models, which were truly tubular. Grayscale images for these models with additive Gaussian noise were generated to train the intensity statistics on. We first trained tubular medial models on each of the 99 cases and computed shape statistics on these in a leave-one-out manner. Corresponding image statistics were also computed in a leave-one-out fashion. Landmarks corresponding to the two ends of the rectum and the point closest to the prostate were also automatically generated and used for the initialization during the segmentation.

Some of these results are shown in Fig. 4. In Fig. 6, we graph the performance of our method over all the samples. The green line shows how far away the trained models are from the ground truth using the average distance measure. The red line shows the result of the segmentation. Note that the size of a voxel in the simulated data is  $2mm^3$ . Thus on the average, all the results are less than half a voxel away from the ground truth.

Having performed very well on synthetic data, we then tested our method on real data that is only approximately tubular. The first data chosen for the test comprises of 17 images of the male pelvis of the same patient over different days. We try to model and segment a small section of the rectum and perform a leave-one-day-out analysis on this data set. Four of the results are shown in Fig. 5. A series of 2D slices from the segmentation shown in the left of this figure are shown in Fig. 8. Initialization was done with the help of four landmarks. On the left are two good segmentations. On the right are two bad segmentations. We note that our segmentation method fails in capturing the ends properly, but then it is designed to ignore the ends as they are arbitrary. Also the segmentation fails on an image of one of the days where the patient has a bloated rectum. This patient does not have such a bloated rectum on any other day; thus, we can consider this case as an outlier. From Fig. 7, we can see that on the average our segmentation method performs reasonably well.

Next, we tried our method on a small set of six common carotid arteries from different patients. Common carotid arteries are a challenging in that they are very small and sit right next to the internal jugular vein that has about the same contrast. Further, head and neck CT images are usually plagued with metal streak artifacts. In our experiment, we only model the section of the common carotid artery that sits near the sternocleidomastoid muscle. Initialization during segmentation was done with the help of three manually added landmarks. Some of these results are shown in Fig. 9. The method performs reasonably well on three of these images. In one of the cases where the method does not work, the carotid artery gets attracted into the nearby internal jugular vein. In another case, the artery is much bent compared with others and is thus a statistical outlier. These






Figure 4. Set of segmentation results in decreasing order of performance on synthetic rectum images: The segmented rectum is shown in a gray color and the ground truth is shown by a translucent blue color. The average distances between the two in clockwise order from the top left in units of mm are: 0.40, 0.47, 0.55, 0.64, 1.52, 0.98, 0.87 and 0.84.




Figure 5. Set of segmentation results in decreasing order of performance on real rectum images: Again, the segmented rectum is shown in a gray color and the ground truth is shown by a translucent blue color. The average distances between the two from left to right in units of mm are: 0.9, 1.2, 1.4 and 4.0.




Figure 6. The red line shows the average distance our segmentation results were from the ground truth on synthetic rectum data. The green line shows the same for the trained models. Note that the size of a voxel is 2mm.

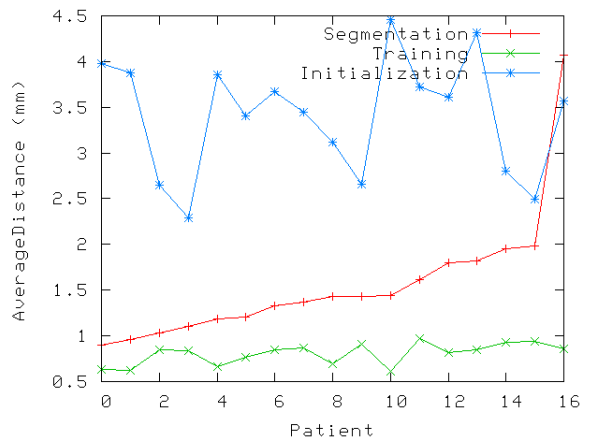


Figure 7. The red line shows the average distance our segmentation results were from the ground truth on synthetic rectum data, the green line shows the same for the trained models, and the blue line represents the initialization used for segmentation.

#### 4. Conclusion

The method performs really well on synthetic rectum data. We are just beginning to apply this method to real world data, and the results on real world data are satisfac-

are shown in Fig. 10. We hope to resolve these problems with more training and better initialization.




Figure 8. Set of 2D slices from a segmentation of a rectum: The red line shows our segmentation. Notice the really poor contrast in the last few slices.




Figure 9. The top three common carotid segmentations look reasonable, whereas the one at the bottom is not so good.

tory. Real world data poses more challenges in the form of poor contrast and the fact that organs in the body are not exactly tubular.

We have developed a new medial method for segmentation of tubular structures from 3D images and tested it on the data set of rectums for one patient. The method per-




Figure 10. While segmenting the common carotid in the top image, it gets attracted to the nearby jugular vein. In the bottom case, the artery is too bent compared with others and the method fails to get any reasonable segmentation.

forms quite well on most of the data failing badly only in one case that represents a bloated rectum. Most importantly, the portion of the rectum that is close to the prostate is identified quite accurately. These segmentations can be further improved with the help of a smaller scale stage that would model the deviations of the rectum from a tube. In the case of head and neck images, we have just started our analysis and the results look satisfactory.

From the point of view of the medical problem, in the future, we might incorporate the knowledge gained from these segmentations to locate and segment the prostate. Also, we intend to test this method for the segmentation of other tubular structures in the body. Methodologically, we intend to support branching tube models.

## 5. Acknowledgements

This work was supported under NIH grant number P01 EB02779. We would also like to thank Eli Broadhurst, Joshua H. Levy, Ja-Yeon Jeong and Qiong Han for some valuable discussion and help with different stages of the segmentation process.

## References

- [1] R. Broadhurst, J. Stough, S. Pizer, and E. Chaney. A statistical appearance model based on intensity quantiles histograms. *ISBI*, 2006. [2](#), [5](#)
- [2] J. Damon. Determining the geometry of boundaries of objects from medial data. *International Journal of Computer Vision*, 63:45–64, June 2005. [3](#)
- [3] Y. Fridman, S. M. Pizer, S. Aylward, and E. Bullitt. Segmenting 3d branching tubular structures using cores. In *MICCAI*, pages 570–577, November 2003. [1](#)
- [4] Q. Han, S. M. Pizer, and J. N. Damon. Interpolation in discrete single figure medial objects. In *Proceedings of the 2006 Conference on Computer Vision and Pattern Recognition Workshop*, page 85, Washington, DC, USA, 2006. IEEE Computer Society. [3](#)
- [5] S. Pizer, T. Fletcher, Y. Fridman, D. Fritsch, A. Gash, J. Glotzer, S. Joshi, A. Thall, G. Tracton, P. Yushkevich, and E. Chaney. Deformable m-reps for 3d medical image segmentation. *International Journal of Computer Vision - Special UNC-MIDAG issue*, 55(2):85–106, November-December 2003. [1](#)
- [6] S. M. Pizer, R. E. Broadhurst, J.-Y. Jeong, Q. Han, R. R. Saboo, J. Stough, G. Tracton, and E. L. Chaney. Intra-patient anatomic statistical models for adaptive radiotherapy. *MICCAI Workshop From Statistical Atlases to Personalized Models: Understanding Complex Diseases in Populations and Individuals*, pages 43–46, October 2006. [5](#)
- [7] J. Stough, R. E. Broadhurst, S. M. Pizer, and E. L. Chaney. Regional appearance in deformable model segmentation. In *Proceedings of Information Processing in Medical Imaging*, 2007. [2](#), [5](#)