Conditional-Mean Initialization Using Neighboring Objects in Deformable Model Segmentation

Ja-Yeon Jeong[a], Joshua V. Stough[a], J. Stephen Marron[b], and Stephen M. Pizer[a]

Department of Computer Science and Statistics[a] and Operation Research[b]
University of North Carolina, Chapel Hill, NC, USA

ABSTRACT

Most model-based segmentation methods find a target object in a new image by constructing an objective function and optimizing it using a standard minimization algorithm. In general, the objective function has two penalty terms: 1) for deforming a template model and 2) for mismatch between the trained image intensities relative to the template model and the observed image intensities relative to the deformed template model in the target image. While it is difficult to establish an objective function with a global minimum at the desired segmentation result, even such an objective function is typically non-convex due to the complexity of the intensity patterns and the many structures surrounding the target object. Thus, it is critical that the optimization starts at a point close to the global minimum of the objective function in deformable model-based segmentation framework.

For a segmentation method in maximum a posteriori framework a good objective function can be obtained by learning the probability distributions of the population shape deformations and their associated image intensities because each penalty term can be simplified to a squared function of some distance metric defined in the shape space. The mean shape and intensities of the learned probability distributions also provide a good initialization for segmentation. However, a major concern in estimating the shape prior is the stability of the estimated shape distributions from given training samples because the feature space of a shape model is usually very high dimensional while the number of training samples is limited. A lot of effort in that regard have been made to attain a stable estimation of shape probability distribution.

In this paper, we describe our approach to stably estimate a shape probability distribution when good segmentations of objects adjacent to the target object are available. Our approach is to use a conditional shape probability distribution (CSPD) to take into account in the shape distribution the relation of the target object to neighboring objects. In particular, we propose a new method based on principal component regression (PCR) in reflecting in the conditional term of the CSPD the effect of neighboring objects on the target object. The resulting approach is able to give a better and robust initialization with training samples of a few cases. To demonstrate the potential of our approach, we apply it first to training of a simulated data of known deformations and second to male pelvic organs, using the CSPD in m-rep segmentations of the prostate in CT images. Our results show a clear improvement in initializing the prostate by its conditional mean given the bladder and the rectum as neighboring objects, as measured both by volume overlap and average surface distance.

Keywords: Shape, Statistical methods, Segmentation, Initialization

1. INTRODUCTION

In medical image analysis, segmentation of human organs is one of central problems. Anatomical structure is complex: some organs are clumped together and touch each other, and the low physical contrast in the image between such structures can make significant parts of the boundary of the target organ ambiguous. Segmenting such an object slice-by-slice manually can be particulary a hard and laborious task even to clinicians, and certainly also to automatic segmentation methods. However, if image intensities at the boundaries of some objects near the target object are easier to identify than the target object, then positions, orientations or material properties of these objects relative to those of the target object can be used as guides to pick out the target object.

Send correspondence to Ja-Yeon Jeong: jeong@cs.unc.edu

Based on this observation, we propose to take a roundabout approach to tackle the challenging problem of automatic segmentation of an object of ambiguous boundary when its nearby objects are easier to segment. The idea is to incorporate the relation of the target object to its neighboring objects into shape priors so that the shape variation of the target object is constrained not only by its own shape variations but also by its positional, orientational information, and even shape deformations relative to neighbor objects over a population. A natural and intuitive way to include this supplementary information into the shape prior is by means of conditional probability distribution. Thus we consider the conditional shape probability distribution (CSPD) of the target shape model given neighboring shape models for segmentation of the target object. The CSPD provides a shape prior for finding a most probable shape of the target object given the configuration of its neighboring objects in relation to the target object.

Estimating both the conditional mean and the shape distribution involves computing an inverse of a sample covariance matrix. The limited sample size and high dimensional feature space of shape descriptors provides a challenge in computing an appropriate pseudo-inverse of the covariance matrix. In estimating the conditional mean, that is, finding a relation of the target mean shape to its neighboring shapes, we employ the principal components regression (PCR) method.[1] PCR corresponds to performing principal component analysis (PCA) and regressing in the feature space given by a few first principal directions that capture major variations explained in a given sample. Using PCR to estimate the CSPD has the following advantages. First, doing PCA on the neighboring shape features reduces the dimension of the feature shape space of neighbor objects. As explained in section 2.3, in the reduced features shape space, we can compute a non-singular sample covariance matrix and get a stable estimation of its inverse. Second, it can help to tighten the distribution of the estimated CSPD since major variation of the neighboring objects is explained in the reduced feature shape space.

Our segmentation method is based on a maximum a posteriori (Bayesian) framework. Deformable shape descriptors together with statistical shape analysis offers an efficient and direct way to include in the image segmentation the prior knowledge of shapes of target objects observed in training samples. As the shape representation, we use a single-figure m-rep[2] as the geometric model for each object. The benefit of using m-reps lies in that they are designed to describe relations within a figure or an object or among objects, the essential property for the objectives of this work.

M-reps consist of sheets of medial atoms that represent object interiors as well as object surfaces. By medial atom transformations, m-reps explicitly capture local shape information such as position, orientation, and magnification, which allows the neighboring geometric entities to be characterized as medial atom transformations of each other. M-reps also supply a local coordinate system for object interiors, thereby giving correspondences across instances of an object. These main properties enable m-reps to describe inter-relations among objects.

A medial atom is defined as a tuple $\underline{\mathbf{m}} = (\mathbf{x},\ r,\ \mathbf{U}^{+1}, \mathbf{U}^{-1}) \in R^3 \times R^+ \times S^2 \times S^2$, where the hub position $\mathbf{x} \in R^3$, the spoke length $r \in R^+$, and the two unit spoke directions $\mathbf{U}^{+1}, \mathbf{U}^{-1} \in$ the unit sphere S^2. Thus $\mathcal{M} = R^3 \times R^+ \times S^2 \times S^2$ can be considered as the space of all possible medial atoms. An m-rep model with n medial atoms is in $\mathcal{M}^n$, i.e., the space of the direct product of n copies of $\mathcal{M}$.

In constrast to the linear point distribution model (PDM) introduced by Cootes et al.,[3] which represents the object by a subset of surface points, the m-rep is a nonlinear shape representation due especially to its rotational components. Fletcher et al.[4] showed that $\mathcal{M}^n$ is actually an Riemannian nonlinear symmetric space. In shape analysis of the PDM, PCA is the most commonly used dimension-reduction technique to limit the object deformation from the whole space to a reasonable linear subspace learned from training samples. PCA has been proven to be effective to characterize anatomical shape variability by mean point positions and their modes of variation and to ease the high dimensional low sample size problem in shape analysis. However, PCA cannot be directly applied to m-reps because of the nonlinearity of m-reps. Our shape analysis of m-reps is based on the work of Fletcher et al.,[4] which lays out a version of PCA using geodesic distances on symmetric spaces, and thus called principal geodesic analysis (PGA).

2. METHODOLOGY

As indicated above, our method aims to get a reliable estimation of the CSPD of shape deformations of a target object given its neighboring objects by employing PCR with shape features of neighboring objects as input data

and shape feature of the target object as the images of input data. We make use of the augmentation described in Pizer et al.[5] to deal with a target object's inter-relation with other objects. Based on the evidence that the atoms T in a target object are highly correlated with neighboring atoms N in the neighboring objects, we choose the atom set N. Typically T is comprised of all of the atoms of the target object, whereas N is a subset of the atoms describing the neighbor objects. These atoms are chosen manually right now, but the choice can be done automatically by setting some threshold of correlation between medial atoms between two neighboring objects. We then use T and N to estimate the conditional probability distribution.

Sections 2.1 and 2.2 describe principal component regression and the conditional probability distribution. Section 2.3 describes the use of PCR in the estimation of conditional probability distribution. In these sections plain capital letters stand for feature vectors and bold face capital letters stand for data matrices of the corresponding feature vector.

2.1 Principal Component Regression

Principal component regression is a multivariate linear regression method using a new set of coordinates (features) that has been extracted from the data by PCA.[1] Let $\mathbf{X}$ be a data matrix in which each row is an input data sample feature vector (independent variables). The desired outputs (dependent variables) are stored in a matrix $\mathbf{Y}$. General multivariate regression can be considered as the optimization of least squares equation

$$\min_{\mathbf{W}} \parallel \mathbf{XW} - \mathbf{Y} \parallel^2 \tag{1}$$

where $\mathbf{W}$ is a regression matrix and the norm is taken as the Frobenius matrix norm, i.e., the sum of the squared norms of the individual errors.

By singular value decomposition (SVD), $\mathbf{X}'$ can be factored into a diagonal matrix $\mathbf{\Lambda}$ between two orthogonal matrices $\mathbf{U}$ and $\mathbf{V}$, i.e., $\mathbf{X}' = \mathbf{U\Lambda V}'$. The columns of $\mathbf{U}$ are the principal directions. Let $\mathbf{U}_k$ be the first k columns of the matrix $\mathbf{U}$ and $\mathbf{\Lambda}_k$ be the diagonal matrix of the first k rows of $\mathbf{\Lambda}$. PCR takes the first k eigenvectors of $\mathbf{X}'\mathbf{X}$ as the features while leaving $\mathbf{Y}$ unchanged. The data matrix now becomes $\mathbf{XU}_k$ and then the least squares equation (1) for PCR becomes

$$\min_{\mathbf{W}} \parallel \mathbf{XU}_k\mathbf{W} - \mathbf{Y} \parallel^2 . \tag{2}$$

Since the norm does not change by pre-multiplication of an orthogonal matrix, (2) can be further simplified using $\mathbf{I} = \mathbf{V}'\mathbf{V} = \mathbf{U}'\mathbf{U}$ and $\mathbf{\Lambda}_k = \mathbf{\Lambda}'\mathbf{U}'\mathbf{U}_k$ as follows:

$$\min_{\mathbf{W}} \parallel \mathbf{XU}_k\mathbf{W} - \mathbf{Y} \parallel^2 = \min_{\mathbf{W}} \parallel \mathbf{V}'(\mathbf{V\Lambda}'\mathbf{U}'\mathbf{U}_k\mathbf{W} - \mathbf{Y}) \parallel^2 = \min_{\mathbf{W}} \parallel \mathbf{\Lambda}'_k\mathbf{W} - \mathbf{V}'\mathbf{Y} \parallel^2 .$$

The solution of $\mathbf{W}$ with minimal Frobenius norm is now simply

$$\widehat{\mathbf{W}} = \mathbf{\Lambda}_k^{-1}\mathbf{V}'\mathbf{Y} = \mathbf{\Lambda}_k^{-2}\mathbf{U}'_k\mathbf{X}'\mathbf{Y}. \tag{3}$$

2.2 Conditional Probability Distribution

We assume that the combined feature vector $X = (T, N)$ follows the multivariate normal distribution. Then, the conditional probability distribution of T given N, $P(T|N)$, also follows the multivariate normal distribution[6] whose mean $E(T|N)$ and covariance matrix $Cov(T|N)$ are

$$E(T|N) = \Sigma_{TN}\Sigma_{NN}^{-}N \tag{4}$$

$$Cov(T|N) = \Sigma_{TT} - \Sigma_{TN}\Sigma_{NN}^{-}\Sigma_{NT} \tag{5}$$

where Σ_{NN}^{-} is a generalized inverse of the covariance matrix of N, i.e., $\Sigma_{NN}\Sigma_{NN}^{-}\Sigma_{NN} = \Sigma_{NN}$ and the covariance matrix of the combined tuple X is

$$\Sigma = \begin{bmatrix} \Sigma_{TN} & \Sigma_{NN} \\ \Sigma_{NN} & \Sigma_{NT} \end{bmatrix}. \tag{6}$$

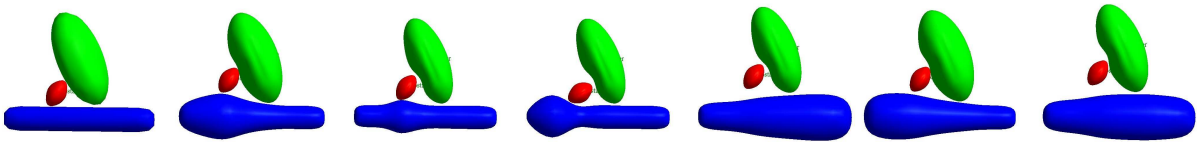

Figure 1. Samples from simulation of a quasi-bladder, quasi-prostate, quasi-rectum complex. The base case is on the left.

Since the target atoms T and the neighbor atoms N are elements in nonlinear symmetric space, we need to take Riemannian log map of X at the Fréchet mean μ of X to map N and A into linear tangent space at μ.* Strictly, the values T and N used above are actually maps of the feature tuples T and N in the nonlinear feature space to the tangent space after mapping. Also, in the equation of conditional mean (4), both T and N must be centered (mean zero) in the tangent space.

2.3 Estimation of Conditional Shape Distribution using PCR

Let $\mathbf{N}$ and $\mathbf{T}$ be data matrix of N and T respectively. By PCA or SVD, we can get principal axes of $\mathbf{N}$. Let $\mathbf{N}' = \mathbf{U}\mathbf{\Lambda}\mathbf{V}'$ by SVD; $\mathbf{U}_k$ has columns of the first k principal axes. Let $\widetilde{\mathbf{N}}$ be the principal components of $\mathbf{N}$, i.e., $\widetilde{\mathbf{N}} = \mathbf{N}\mathbf{U}_k$. When the sample size is n, the covariance matrix of $\widetilde{\mathbf{N}}$ is

$$\Sigma_{\tilde{N}\tilde{N}} \;=\; \frac{1}{n-1}\widetilde{\mathbf{N}}'\widetilde{\mathbf{N}} \;=\; \frac{1}{n-1}\mathbf{U}_k'\mathbf{N}'\mathbf{N}\mathbf{U}_k \;=\; \frac{1}{n-1}\mathbf{U}_k'(\mathbf{U}\mathbf{\Lambda}\mathbf{V}')(\mathbf{V}\mathbf{\Lambda}\mathbf{U}')\mathbf{U}_k \;=\; \frac{1}{n-1}\mathbf{\Lambda}_k^2, \tag{7}$$

where $\mathbf{\Lambda}_k$ is the diagonal matrix of the first k diagonal components of $\mathbf{\Lambda}$. It is clear that estimating the inverse of $\Sigma_{\tilde{N}\tilde{N}}$ is more stable compared to ridge regression approach[1,7] that tries to regularize Σ_{TT} by adding a typically positive small constant and computing the inverse.

Now we can show that the estimated conditional mean $E(T|\tilde{N})$ is equivalent to the estimator $\widehat{\mathbf{T}}$ by PCR, that is, $\widehat{\mathbf{T}} = \mathbf{N}\mathbf{U}_k\widehat{\mathbf{W}}$. With (7) and $\Sigma_{T\tilde{N}} = 1/(n-1)\mathbf{T}'\widetilde{\mathbf{N}}$, the conditional mean is

$$E(T|\tilde{N}) \;=\; \Sigma_{T\tilde{N}}\Sigma_{\tilde{N}\tilde{N}}^{-1}\widetilde{\mathbf{N}}' \;=\; \mathbf{T}'\widetilde{\mathbf{N}}\mathbf{\Lambda}_k^{-2}\widetilde{\mathbf{N}}' \;=\; \mathbf{T}'\mathbf{N}\mathbf{U}_k\mathbf{\Lambda}_k^{-2}\mathbf{U}_k'\widetilde{\mathbf{N}}'. \tag{8}$$

We can easily see that (8) is the same as $\widehat{\mathbf{T}}'$ when $\widehat{\mathbf{W}}$ is replaced with the identity (3).

Equation (5) allows the computation of the CSPD's covariance matrix $Cov(T|\tilde{N})$ in the tangent space. Once we have this estimate of this conditional covariance matrix, we perform PCA on it and choose the subset of the principal directions that is sufficient to describe the variability of the target m-rep shape space.

3. APPLICATIONS AND RESULTS

We applied the new conditional shape model to two data sets. Section 3.1 describes and illustrates synthetic data of multi-object models to which simple transformations are applied and the estimated principal geodesic modes of variation of the CSPD estimated from the synthetic data. Section 3.2 describes segmentation of the prostate given neighbor atoms from bladder and rectum objects of real patients with estimated conditional means as initializations of the segmentation.

3.1 Simulated Multi-objects

The first data set we used to apply our method is a randomly generated multi-object complex whose shape changes and motions simulate the bladder, prostate, and rectum in the male pelvis. The complex is made up of the pelvic arrangement of two slab ellipsoid (quasi-bladder and quasi-prostate) and one tube (quasi-rectum) m-reps. Fig. 1 shows the base model and sample deformations of the base model. The quasi-bladder slab ellipsoid is a 3×7 grid of medial atoms, the quasi-prostate slab ellipsoid is a 3×5 grid of medial atoms, and the

*Refer to Fletcher et al.[4] for detailed explanation of the log map and the exponential map.

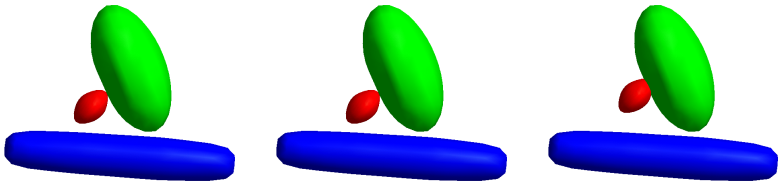

(a) (b) (c)

Figure 2. Shape variability of simulated multi-object complex (a) mean+1σ along the 1st principal geodesics of conditional distribution of quasi-prostate given augmented atoms from quasi-prostate and quasi-bladder (b) mean of simulated multi-object complex (c) mean+1σ along the 1st principal geodesics of non-conditional distribution of quasi-prostate

quasi-rectum tube is a column of 8 tube atoms. The transformations applied to each object are to imitate shape and motion changes that are observed in real organs and interaction between objects. The transformations that simulate the effect of neighboring objects depend on one random parameter α that indicates the stiffness of a quasi-prostate. We model $\alpha \sim N(0, \sigma_\alpha)$ where $0 < \sigma_\alpha < 1$ although we do not use any value of α greater than 1.

Quasi-bladder Ellipsoid M-rep We set a point x_{bp} on the base ellipsoid quasi-bladder surface to simulate the location through which the urethra goes from quasi-bladder through the quasi-prostate. One random transformation is applied on the base quasi-bladder. This transformation imitates filling of the bladder with urine. The transform scales up the lengths of the major axes of the ellipsoid to produce a uniform thickening of the ellipsoid. If a, b, c are lengths of the major axes, then their ratio $a : b : c$ follows a log-normal distribution, i.e., $a : b : c = e^l$ where $l \sim N(0, \sigma_l)$. l is 1 independent parameter of the transformation. x_{bp} is relocated to x'_{bp} by this transformation.

Local Effect from Quasi-bladder to Quasi-prostate As the quasi-bladder expands, the quasi-prostate shifts and in turn dents the quasi-bladder when the expansion of the quasi-bladder is large. If x_{bp} would have moved to x'_{bp} in the absence of a quasi-prostate, x_{bp} is restricted to move only by the fraction α of the distance between x_{bp} and x'_{bp}, and the quasi-prostate is shifted by the same amount along the same direction. The amount of dent in the quasi-bladder at x_{bp} by the quasi-prostate is determined deterministically to prevent interpenetration.

Quasi-Rectum Tube M-rep One random transformation is applied to the quasi-rectum tube to imitate the bulging of the rectum due to gas. The simulated bulge is cylindrical along the axis of tube, so the cross-section of the bulge is a circle. The transformation has 3 random parameters: the center position of the bulge along the tube axis, the width of the bulge, and the height of the bulge. The last two parameters determine the shape of the bulge.

Local Effect from Quasi-rectum to Quasi-prostate Bulging of the quasi-rectum causes the quasi-prostate to rotate and to shift. Consider a coronal plane that goes through the center of the quasi-prostate. If bulging happens in one side of the coronal plane, the quasi-prostate tilts toward the other side of the plane. If bulging happens right below the quasi-prostate, it pushes up the quasi-prostate. The direction and the displacement of the quasi-prostate is constrained so that these two objects do not interpenetrate each other. The displacement direction of quasi-prostate is from x_t to m_p where m_p is the center of the quasi-prostate and x_t is a top point of bulging on quasi-rectum. The displacement amount is also determined by α and the distance from the top point x_t to a point where the quasi-prostate is closest to the quasi-rectum.

Training Results We trained the CSPD of the quasi-prostate given selected neighbor atoms from the quasi-bladder and quasi-rectum on 188 synthetic multi-object m-reps. As neighbor atoms for the quasi-prostate, 4 atoms were selected from the quasi-rectum located right below the quasi-prostate, and 9 atoms were selected from the quasi-bladder near the end of the quasi-prostate meeting the quasi-bladder. Since these are simulated objects, we did not apply any alignments before training. The dimensions of the tube atom and slab atom in tangent space are 8 and 9 respectively.[4,8] So the dimension of the combined T and N atoms in tangent space is now $4\times8+9\times9 = 113$. In computing the conditional mean and shape distribution of the quasi-prostate, we chose

the first 4 principal components of the N atoms to reflect the 4 random transformations of the quasi-bladder and the quasi-rectum; this is a large reduction in dimension from 113 to 4. This reduction allowed us to have a stable estimation of conditional covariance matrix.

Fig. 2 shows the comparison between the first principal geodesics of conditional and non-conditional distributions of the quasi-prostate. To do this comparison, the first mode of conditional shape variation was applied to the non-conditional mean of the quasi-prostate. As we can see in fig. 2, the first non-conditional shape mode is not constrained by the neighboring two objects and at 1 standard deviation from the mean the deformed quasi-prostate penetrates the quasi-bladder. The other modes of variations from the CSPD show that 1) they capture the tilting and the shift toward the quasi-bladder caused by the bulge of the quasi-rectum and 2) they are within the confines of neighboring objects. Although non-conditional modes capture the tilting and shift, the range of their variation, especially for the shift of the quasi-prostate, is too large. The eigenvalues of non-conditional modes are all at least an order of magnitude larger than those of conditional modes.

3.2 Objects in Pelvic Region

Our data consist of four patients' image sets each of which is a series of approximately 16 CT scans taken during radiotherapy treatments. They are fractionated over a series of dates, and a new image is acquired before each fraction. Each image has an in-plane resolution of 512×512 with voxel dimension of 0.98mm×0.98mm and 3mm distance between slices. The images contains the bladder, prostate, and rectum. We have expert manual segmentations of these three objects in each image. M-reps were fit to the manual segmentations: 5×6 grids of medial atoms for the bladder, 7×4 for the prostate, and 17 columns of quasi-tube[8] atoms for the rectum. Our fitting procedure prevents penetration among models and folding of the interior of the object and maintains regularity of grid for correspondence of medial atoms.[9] For the purpose of this study the bladder and rectum models used during prostate segmentation are these training models.

Segmentation Structure The prostate in each image was segmented by Bayesian optimization over m-rep principal geodesic coefficients. The objective function for optimization has two terms: geometric penalty and image match. For the image match term, we use Broadhurst's regional intensity quantile functions (RIQF).[10] The m-rep defines a coordinate system that provides an explicit correspondence between deformation of the mean m-rep model and the 3D volume in the object boundary region. The fitted m-reps are placed into the associated grayscale images, and image intensity quantile functions (inverse cumulative histograms) in regions interior to and exterior to the target object boundary are recorded in each training image. The image match function measures statistical distance between observed quantile functions in a target image and trained quantile functions over the population for the corresponding region of the training set. For the geometric penalty term, we use Mahalanobis distance function from a mean model to a deformed model during the segmentation. Segmentation starts with the mean model placed in a new image of a patient by a similarity transformation computed by aligning the pubic bone of the patient. Segmentation then proceeds by a conjugate gradient optimization of the posterior of the geometric parameter given the new image data.

Table 1. Number of fractions (sample size) per patient

Patient Id	3101	3106	3108	3109
Number of fractions	14	16	17	17

Results We did intra-patient successive segmentations. Table 1 shows the sample size per patient. At any fraction i, the CSPD of each patient was trained using training fits from fraction 1 up to fraction $i-1$. To allow comparison to leave-one-day-out experiments that we had previously done, the RIQF per fraction i was trained by all days' grayscale images of the patient except for fraction i. Fig. 3 shows the comparison of mean volume overlaps of all 4 patients at each fraction among conditional means to date, non-conditional means to date, and leave-one-out means from fraction 6 to 13. It also shows the same comparison by mean average surface distances. We did not include the first 5 fractions in the comparison because 5 samples are not enough to train covariance matrices to compute the conditional mean. For neighbor atoms to the prostate from the bladder and rectum, we chose around 6 to 9 different set of atoms of bladder per patient and 1 fixed atom of rectum across patient.

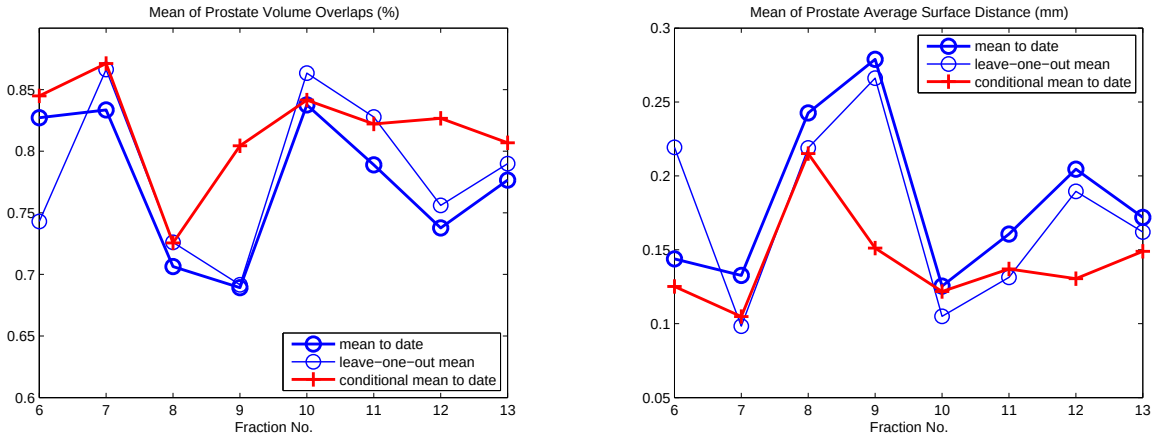

(a) (b)

Figure 3. (a) Mean *Jaccard similarity coefficients (JSC)*[11] of mean prostate volume overlaps with training prostate binary images; (b) Mean of average surface distance in mm between mean prostate and training prostate binary images. Mean of these measures are taken over 4 patients at each fraction, from fraction 6 to fraction 13. The thin blue line with the round marker indicates leave-one-out mean prostate, the thick blue line with the round marker indicates mean prostate to date, and the red line with the cross marker indicates conditional mean given neighbor atoms from the bladder and prostate.

We took the coefficients of the first three principal axes of the neighbor atoms of the bladder and rectum as new features for neighbor atoms in estimating the CSPD and its mean.

Result 1: Initialization The comparison clearly indicates that the conditional mean prostate to date gives much better initialization than the mean to date. Compared with the mean to date, the conditional mean prostate to date improved the volume overlap (intersection over union) in 32 out of 46 cases. Even compared with leave-one-out mean, the conditional mean prostate to date improved the volume overlap in 27 out of 46 cases. (46 is the total fractions of all patients from fraction 6.)

Table 2. Mean *JSC* of figure 3

Fraction No.	6	7	8	9	10	11	12	13
Mean to date	0.8272	0.8336	0.7064	0.6893	0.8376	0.7890	0.7378	0.7766
Leave-one-out mean	0.7429	0.8664	0.7261	0.6915	0.8634	0.8278	0.7560	0.7900
Conditional mean to date	0.8449	0.8713	0.7255	0.8044	0.8415	0.8221	0.8267	0.8068

Result 2: Segmentation Since it was clear that the conditional mean outperforms the non-conditional mean, we initialized segmentations with the conditional mean to date. We then did another comparison between conditional and non-conditional principal geodesics by performing the segmentation with both principal geodesics. Fig. 4 shows the comparison results both in volume overlap and average surface distance.

In average, the conditional distribution gives better results from around fraction 9 than the best leave-one-out segmentations that we have now. However, it provides roughly equivalent results to the non-conditional principal geodesics, and it uses fewer eigenmodes. From fraction 6 to 8, the conditional distribution does worst, but for later fractions non-conditional and conditional distributions are superior to the leave-one-out non-conditional distributions.

The training sample size determines the number of possible principal geodesics (eigenmodes) used in estimating a probability distribution. During the training we set a threshold of 10^{-10} for eigenvalues and keep ones of higher eigenvalue. In this experiment the maximum number of eigenmodes was set to 6. Table 3 shows the

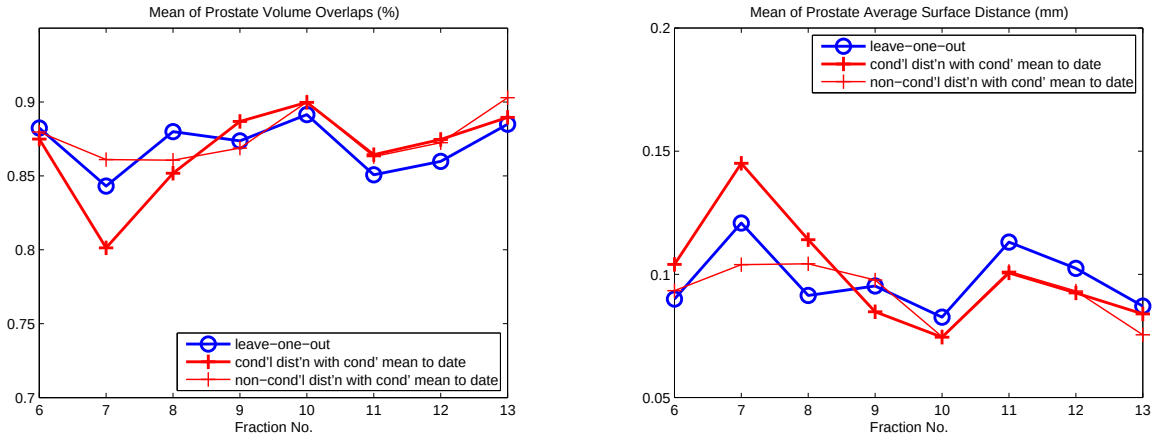

(a) (b)

Figure 4. (a) Mean *Jaccard similarity coefficients (JSC)* of segmented prostate volume overlaps with training prostate binary images; (b) Mean of average surface distance in mm between segmented prostate and training prostate binary images. Mean of these measures are taken over 4 patients at each fraction, from fraction 6 to fraction 13. The thin blue line with the round marker indicates leave-one-out segmentation initialized with leave-one-out mean, the thick red line with the cross marker indicates successive segmentation starting with conditional mean to date and the conditional mean model deformed by conditional principal geodesics also trained up to a target fraction, and the thin red line with the round marker shows successive segmentation starting with the conditional mean to date but the conditional mean deformed by non-conditional principal geodesics of the prostate trained up to a target fraction.

number of eigenvalues kept for patient 3106. The same pattern appears for the other patients. From table 3, we can see that for first a few fractions there are not enough conditional principal geodesics to adequately deform the initial model. The reason appears when we look at equation (5) for the covariance matrix of the CSPD. The second term in this equation is designed to correct the non-conditional covariance matrix using the reduced features of neighbor atoms. When there are too few training samples, the correction used to form the conditional probability distribution inappropriately removes shape variability from the non-conditional probability distribution.

With the segmentation results, we confirmed the importance of the initialization since the segmentation initialized with the conditional mean to date and proceeded with the non-conditional principal geodesics trained cumulatively outperforms the segmentation with the leave-one-out mean and the non-conditional principal geodesics.

Table 3. Number of principal geodesics for patient 3106

Fraction No.	6	7	8	9	10	11	12
Conditional	0	1	2	3	4	5	6
Non-conditional	3	4	5	6	6	6	6

4. CONCLUSION AND DISCUSSION

We outlined a new method to estimate the shape probability distribution of an object conditioned on its neighboring objects. The method is based on principal component regression analysis that provides a way to get a stable and robust estimation of the conditional mean and shape distribution. Since this approach uses additional information extracted from the relation between the target object and its neighboring objects, it is capable of segmenting an object of insufficient boundary intensity information when good segmentations of neighboring objects are available.

Initialization by conditional means outperforms non-conditional means, and segmentation starting from conditional means also looks promising. However, as mentioned before the sample size is still a critical factor in estimating conditional principal geodesics, and an issue on how to determine the number of principal directions from neighbor atoms in estimating the conditional mean and probability distribution still remains.

Another open issue is the choice of neighbor atoms. The choice affects the estimate of the conditional mean and conditional principal directions. We can choose these atoms by looking at correlations at every pair of atoms. However, the computational cost of this approach can be very high.

The most challenging issue is correspondence. Although the correspondence derived from our fitting mechanism works reasonably well to represent shape of the object itself, it may not be the right correspondence to describe the interrelation among the object.

Finally, the initial models for bladder and rectum used in segmentation are from the training bladder and rectum models. We need to verify that our current initialization and segmentation results still hold when training bladder and rectum models are replaced by good segmentations of the bladder and rectum.

ACKNOWLEDGMENTS

The authors would like to thank all members in MIDAG group for their contributions to data preparations and to code for segmentation, especially Edward Chaney for providing us valuable insights about anatomy , Joshua Levy and Rohit Saboo for their constant help in getting the latest segmentation results, and Qiong Han for working with us to build the program that generate the simulated data. The work described in this paper was done under the partial support of NIH grant P01 EB02779

REFERENCES

1. J. Shawe-Taylor and N. Cristianini, *Kernel Methods for Pattern Analysis*, Cambridge University Press, 2004.
2. 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, (O. Faugeras, K. Ikeuchi, and J. Ponce, eds.)* **55**(2), pp. 85–106, 2003.
3. T. Cootes, C. Taylor, D. Cooper, and J. Graham, "Active shape models – their training and application," *Computer Vision and Image Understanding* **61**, pp. 38–59, 1 1995.
4. P. T. Fletcher, C. Lu, S. M. Pizer, and S. Joshi, "Principal geodesic analysis for the study of nonlinear statistics of shape," *IEEE Transactions on Medical Imaging* **23**, pp. 995–1005, 8 2004.
5. S. Pizer, J.-Y. Jeong, C. Lu, K. Muller, and S. Joshi, "Estimating the statistics of multi-object anatomic geometry using inter-object relationships," in *in International Workshop on Deep Structure, Singularities and Computer Vision (DSSCV), (O.F. Olsen, L. Florack, and A. Kuijper, eds.)*, *LNCS*, pp. 60–71, 6 2005.
6. R. J. Muirhead, *Aspects of Multivariate Statistical Theory*, Wiley, 1982.
7. M. de Bruijne, M. T. Lund, L. B. Tanko, P. P. Pettersen, and M. Nielsen, "Quantitative vertebral morphometry using neighbor-conditional shape models," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 1–8, 2006.
8. R. R. Saboo, J. H. Levy, E. M. Chaney, and S. M. Pizer, "Modeling populations of nearly tubular objects through medial models." Submitted to CVPR 2008.
9. D. Merck, G. Tracton, R. Saboo, E. Chaney, S. Pizer, and S. Joshi, "A methodology and implementation for constructing geometric priors for deformable shape models," tech. rep., 2006.
10. R. Broadhurst, J. Stough, S. Pizer, and E. Chaney, "A statistical appearance model based on intensity quantiles histograms," in *IEEE International Symposium on Biomedical Imaging*, pp. 422–425, 4 2006.
11. W. R. Crum, O. Camara, and D. L. G. Hill, "Generalized overlap measures for evaluation and validation in medical image analysis," *IEEE Transactions on Medical Imaging* **25**, pp. 1451–1461, 11 2006.