

Training models of anatomic shape variability

Derek Merck,^{a)} Gregg Tracton, Rohit Saboo, Joshua Levy, Edward Chaney, Stephen Pizer, and Sarang Joshi
Medical Image Display & Analysis Group, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina 27599

(Received 4 June 2007; revised 27 March 2008; accepted for publication 15 May 2008; published 15 July 2008)

Learning probability distributions of the shape of anatomic structures requires fitting shape representations to human expert segmentations from training sets of medical images. The quality of statistical segmentation and registration methods is directly related to the quality of this initial shape fitting, yet the subject is largely overlooked or described in an *ad hoc* way. This article presents a set of general principles to guide such training. Our novel method is to jointly estimate both the best geometric model for any given image and the shape distribution for the entire population of training images by iteratively relaxing purely geometric constraints in favor of the converging shape probabilities as the fitted objects converge to their target segmentations. The geometric constraints are carefully crafted both to obtain legal, nonself-interpenetrating shapes and to impose the model-to-model correspondences required for useful statistical analysis. The paper closes with example applications of the method to synthetic and real patient CT image sets, including same patient male pelvis and head and neck images, and cross patient kidney and brain images. Finally, we outline how this shape training serves as the basis for our approach to IGRT/ART. © 2008 American Association of Physicists in Medicine. [DOI: 10.1118/1.2940188]

Key words: image segmentation, deformable shape models

I. INTRODUCTION

Our target application, adaptive and image-guided radiotherapy (ART,¹ IGRT²), requires a framework for accurately mapping anatomical objects from inter-fractional images into the same coordinate system as the planning or other reference image. Current research activity is aimed at developing practical and reliable methods for space-filling nonrigid mapping³⁻⁸ to overcome the shortcomings of current clinical methods based largely on rigid registration. This work focuses on the framework required for probabilistic image segmentation as a means for automatically computing nonrigid mappings. Probabilistic segmentation can also provide a basis for image guided surgery (e.g.,^{9,10}) and diagnosis.¹¹

Probabilistic segmentation is based on understanding the shape variability of the anatomic structures found in medical images. Robust probability density estimates of shape have been shown to be effective for object-based methods of probabilistic segmentation^{9,12,13} for two reasons: (1) their relatively low dimensionality allows us efficiently to compute optimal solutions, and (2) the optimal solutions yield anatomically credible objects. Geometrically legal space-filling nonrigid mappings can be derived from differential equations (e.g., optical flow,^{14,15} elastic flow,¹⁶ visco-elastic flow^{6,17,18} or finite-element models^{19,20}) and then reduced to lower dimension for statistical analysis by means of control points or decomposition via basis functions. However, as shown in Fig. 1, these methods cannot be restricted to credible shapes because one must trade off spatial resolution with the fact that the extremely high dimensions prevent learning the feature space of credible shapes from the limited number of training cases available.²¹ Furthermore, while these meth-

ods have been applied to images with similar topology such as brain or inter-fractional or 4D scans, they have not been shown to be applicable to cross-patient images. This article therefore restricts its attention to models of one or more anatomic objects, e.g., of their surfaces or skeletons.

The method presented here is based on statistical deformable shape models (SDSMs), which are object-based parametric shape models that characterize shape changes relative to a “typical” instance. The typical shape usually is taken to be the mean shape, and relative changes are encoded as a limited number of important eigenmodes of shape variability. Using a SDSM, any particular shape in the space defined by the training set can be completely and uniquely identified by a few coefficients to within some small truncation error. Once trained, SDSMs can be used as the basis of numerical methods for shape discrimination, comparison, and interpolation, for longitudinal shape studies,¹¹ and for deformable image segmentation and registration.²²⁻²⁴ (See Ref. 25 for a recent summary of clinical applications of both space filling and shape based image registrations.) Figure 2 diagrams the process of training and use of the estimated probability distributions on the object(s). This article is concerned primarily with defining a framework for fitting parametric models to training images while maintaining credible, legal shapes with locational correspondence across the training population. However, the method’s application to image guided radiotherapy is outlined in the discussion section.

We consider shape training as a special case of Bayesian image segmentation.²⁶ In SDSM-based segmentation, a starting shape is initialized to a target image and automatically deformed based on the eigenmodes of shape variation to op-

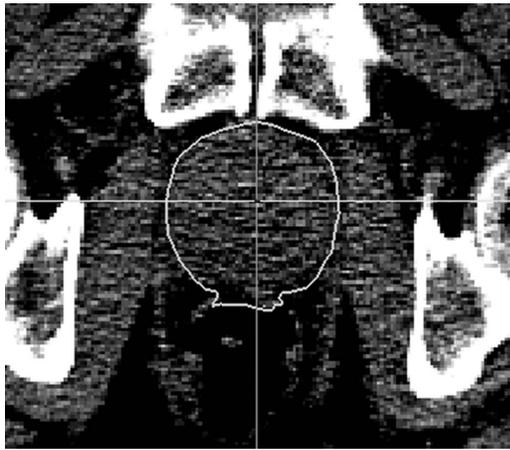


FIG. 1. Purely intensity driven registration can generate noncredible shapes when pulling daily images back to the atlas for IGRT/ART. Here the boundary of the prostate segmentation has drifted into the rectum during an intensity-driven registration. Such segmentations are not allowed with our shape based training (see Fig. 16).

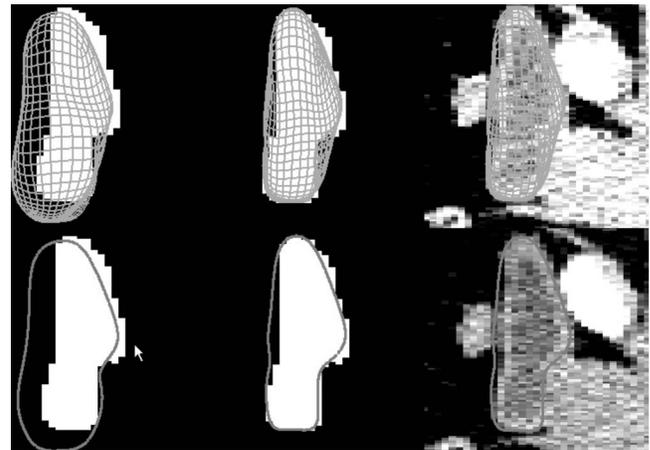


FIG. 3. Fitting a statistical deformable model to a target training image. (Top) 3D surface views and (bottom) single sagittal slice views of bladder template geometry (left) coarsely aligned to a target training image, (middle) deformably fit to that image, and (right) in the context of the actual gray-scale data.

timely match the image data, as illustrated in Fig. 3. The deformation is driven numerically through the SDSM parameter space by minimizing an objective function including a term reflecting the probability of the deformed shape and a term measuring how well the shape matches the image. However, by using a shape parameterization that supports reasonable physical deformations, it is possible to define purely geometric conditions which establish and maintain legality and the locational correspondence over the training population we need to compute correct statistics. After the training population has been fit, these first-round shape estimates are used to compute coarse statistics. The training images are then refit about a new mean, relaxing the purely geometric constraints as the shape statistics are iteratively refined. This recomputation of an atlas, unbiased to a particular subset of the training instances, by computing a succes-

sively better estimate of a mean as the fits improve, is comparable to the iterative computation of a Fréchet mean and fits to it used in the nonrigid registration framework.^{18,27} However, in our method not only the mean but also the variability is being recomputed at each iteration.

Our automatic method for shape training has been applied to a variety of anatomic structures from both synthetic and real 3D medical images, such as inter-fractional and cross-patient male pelvis, inter-fractional head and neck, and cross-patient brain and kidney images. Whereas inter-fractional SDSMs cover day-to-day anatomic shape change, such as the bladder filling and emptying, etc., cross-patient SDSMs must be able to account for much broader anatomic shape variability, such as how prostates differ in both shape and relationship to the neighboring bladder, rectum, and pubic bones from patient to patient.

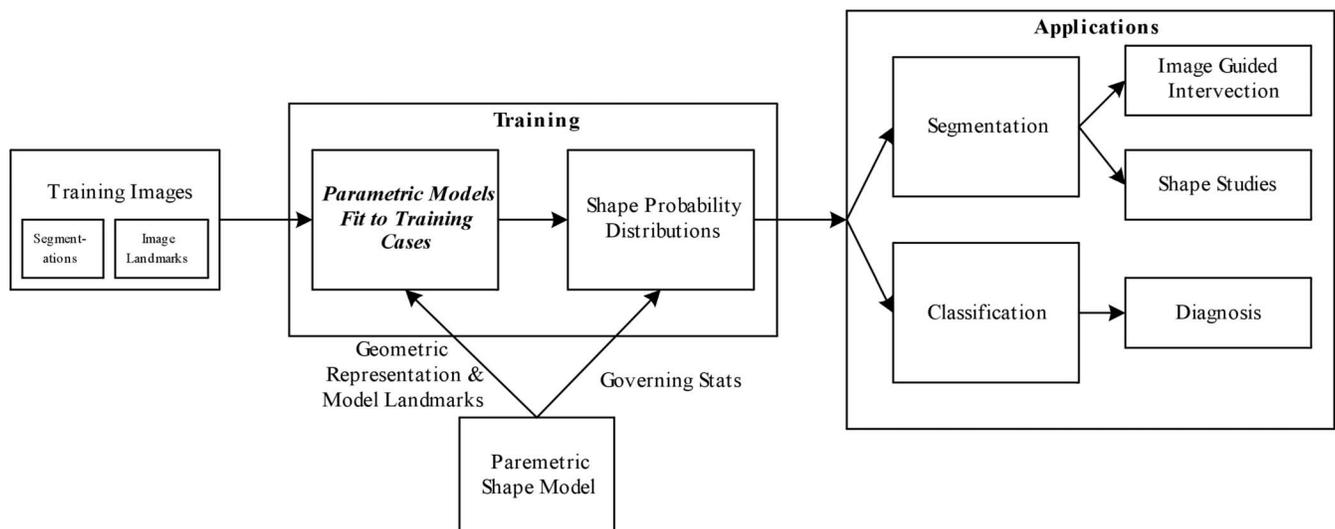


FIG. 2. The process of training and using probability distributions on object shapes in clinical applications. The step outlined in bold, fitting training images with parametric models as input to statistical analysis, is the main subject of this article.

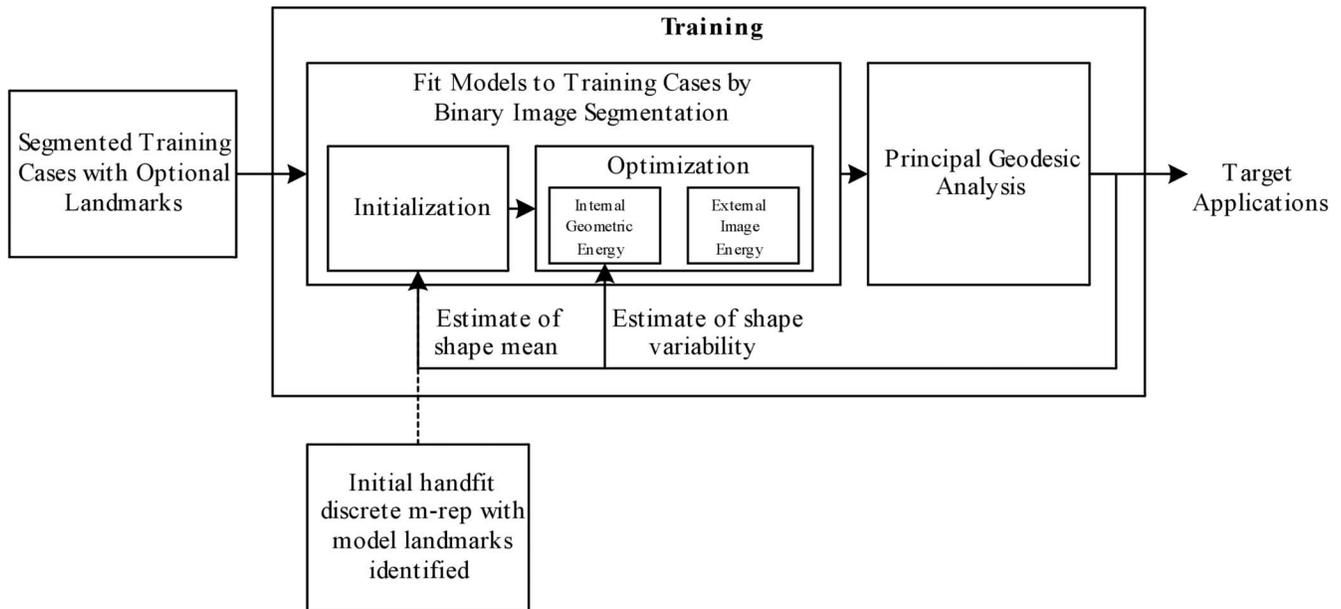


FIG. 4. A detailed view of our method for the training step shown in Fig. 2. Shapes are fit to the training images iteratively according to a binary image segmentation algorithm, with purely geometric terms relaxed in favor of converging group statistics. See the Appendix for a summary of the discrete medial representation and principal geodesic analysis.

II. METHOD

Extraction of shape representations of training cases has been accomplished in the following ways for the boundary-only shape models that are common in the statistical shape analysis literature. For these, the model is initialized by sampling the boundary voxels of each training sample to produce a nonfolding tiled surface. Then each case's surface is reparameterized to give correspondence according to geometric, mechanical, or local image properties. As examples, Ref. 28 uses no image match grounds, but requires a regular sampling with tight geometric distributions. Active appearance models^{12,29} require both a tight geometric distribution as well as similar local image intensities. The training method used in Refs. 30 and 31 reparameterizes the sampled surface according to an orthogonal decomposition based on spherical harmonics, then resamples the surface at locations corresponding to equally distributed points on the parameterizing sphere in that framework. Reference 32 uses a finite element model to track voxel correspondences for a retrospective dose study. Reference 33 describes parametric topological constraints for active contour models (snakes)^{24,34} or active shape models. Many of these approaches can obtain good results, but none directly constrain their shape statistics to the generation of only credible shapes or, for most, even of nonself-interpenetrating shapes.

To obtain correspondence while maintaining a space of credible objects, our process is designed as a special case of binary image segmentation within the probabilistic segmentation framework. The goal of our training process is therefore to jointly estimate both a parametric shape space that

covers the organ shape over all n images, and at the same time to obtain the best set of corresponding descriptors of parametric models for each organ in each image. As illustrated in Fig. 4 and overviewed in algorithm 1, in this framework an energy function composed of internal geometric forces and external image match forces is optimized over a shape space. Typically this space comes from the trained probability distributions, but in their absence, we rely on a shape parameterization that supports reasonable physical changes and a carefully written internal geometric term to constrain possible shape changes.

The parameterization that we use is a discrete medial representation called m-reps, although the ideas can be generalized to other SDSMs with a well defined volumetric coordinate system and an explicit encoding of volumetric legality. See the Appendix for a summary of m-reps¹² and its governing statistical model, principal geodesic analysis.³⁵ Since m-reps capture not only positional information (i.e., of the skeleton and the implied boundary) but also orientational information (of spokes between the skeleton and the boundary), the parameterization handles plausible physical changes such as local twisting and bending in a way that can be easily captured by the statistical model. Further, we can leverage differential geometry directly in our legality term to enforce nonself-interpenetration.

This method section of this article is organized to follow Fig. 4. After discussing the inputs to the process in Sec. II A, we describe the initial model-to-image alignment in Sec. II B. The energy function is overviewed in Sec. II C, the internal geometric term is explained in Sec. II D, and the external image term is explained in Sec. II E. The itera-

tive optimization framework is described in Sec. II F. Algorithm 1 gives an additional overview of the method. A major contribution of this article is the iterative fitting of models and shape probability estimation indicated in step 2 of algorithm 1. Additional novel contributions are describing the

principles and design of shape terms for the energy function used in steps 2.1.1–2.1.3 that maintain object correspondence and legality during the training fitting, and describing the iterative shape estimation where purely geometric terms are relaxed as shape statistics become available.

Algorithm 1: Iteratively Training Models of Shape Variability

Input: I , a collection of image segmentations i defining geometric truth

1. $R_0 = \text{HandFit}(i \text{ in } I \text{ s.t. } i \text{ is reasonably typical})$
2. Until difference $(R_k, R_{k+1}) < \text{convergence error}$:
 - 2.1. For each i in I :
 - 2.1.1. Find m_i s.t. $m_i = \text{Arg/Min}_{(\text{sim trans of } R_k)} \text{Energy}(R_k, i)$
 - 2.1.2. Find m_i s.t. $m_i = \text{Arg/Min}_{(\text{global deformations of } R_k)} \text{Energy}(R_k, i)$
 - 2.1.3. Find m_i s.t. $m_i = \text{Arg/Min}_{(\text{local deformations of } R_k)} \text{Energy}(R_k, i)$
 - 2.2. $(R_{k+1}, \text{Shape stats}) = \text{Principal Geodesic Analysis}(M = \{m_i\})$

II.A. Materials

For inter-fractional shape study, our method begins with a collection of CT scans from a single patient over the course of a multi-day treatment regimen. Typical images are $512 \times 512 \times 40$ voxels in dimension and have $1 \times 1 \times 3$ mm voxel resolution. The CT images include both treatment and nontreatment organs. Between 12 and 18 scans are collected for inter-fractional studies. Other studies have slightly different image characteristics, but the method remains the same; details are given in the results section for the various case studies presented.

Expert raters segment structures of interest in each image by drawing contours of target and nontarget organs on transaxial slices. Methods to generate ground truth from these segmentations comprise a separate area of research beyond the scope of this article.³⁶ Because the quality of the manual segmentations is critical to the quality the derived model, at the present time we train against individual raters whose segmentations have been reviewed and, if necessary, edited by other experts. Precautions are taken to control for intra-rater variability across images of the same patient. For example, segmentations from one image are transferred to the next, coarsely registered with the corresponding objects, and then edited to match the objects in that day's image. Moreover the prostate volume is not allowed to change from day to day. Based on anecdotal evidence these precautions are believed to control for random variations that would be observed when segmenting each image *de novo*. The segmentations are then scan converted to a label image per organ. These label images have the same extent and resolution as the source images, but all voxel values are set to either 1, if that voxel is inside or on the segmentation's boundary, or 0 otherwise. Throughout this article, we will refer to these per organ label training images as the set I_{l-n} , where n is the number of images in this population.

II.B. Initialization

Our training process starts by hand manipulating a model to be modestly representative of the shape we are attempting

to train. This typical shape is the presumptive geometric template for the shape space, R_0 . R_0 is our reference shape, from which all other shapes will initially be derived and relative distances measured. As long as the initial model has sufficient sample resolution and correct topology (e.g., tubes for vessels, a bent blob for kidney, a multi-blob figure for liver), our method converges for even arbitrarily poor starting models such as a generic slab. However, it is quite clear that the quality of the SDSMs we compute and the speed of convergence is directly related to the input quality at each step, making a good initial shape estimate particularly important.

Using R_0 as a coordinate system governing our region of interest, we can identify in R_0 a set of model-centric landmarks, LM , which we wish to keep explicitly in correspondence to within some tolerance throughout the training population (see Fig. 5). These LM correspond to a set of explicitly identified anatomic features, LI , landmarks noted by the raters on the segmentations they manually produce. For example, when training a prostate shape model, we use posterior and anterior poles and the urethral entrance and exit. Although we typically use only surface points as landmarks, our shape parameterization provides a volumetric model-centric coordinate system, which allows landmarking of positions both inside or on the object, such as the opening of the prostatic urethra, and outside the object, such as the nearest boney landmark to the prostate.

The initial step in fitting R_0 to a particular label image is then to align R_0 to the image according to the landmark

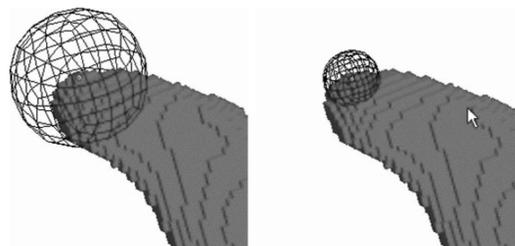


FIG. 5. An image landmark identified at the tip of a segmentation with (left) large tolerance and (right) a tighter tolerance.

pairings LM to LI. This alignment can remove global transformations such as translation, scale, or rotation from the shape space we are attempting to estimate. The more tightly aligned the training cases are to each other, the more tightly the geometric template will be able to fit the data, and the more representative the derived statistics of deformation will be. We have explored a variety of methods for designating this initial alignment: We frequently use a Procrustes similarity transform, although with some highly variable objects, we could use an affine transform. And one could imagine integrating even more complex nonlinear landmark based registrations such as Refs. 37 and 38, but then the probability distributions would not be describing these shape changes. In the absence of explicit landmarks, implicit landmarks from curvature or other derived properties may be computed. For simple synthetic objects, it is sometimes sufficient to align the centroid, volume, and orientation of R_0 's surface to by similarity transform to the image moments of the binary training image.

II.C. Defining the energy function

The best deformation of R_0 into each training image is computed numerically, by optimizing model parameters according to a metric that measures the goodness of a model, M , fit to a given image, I . This energy metric normally consists of two terms, an “external” image match term and an “internal” geometric typicality term. This division can be theoretically founded in probabilistic terms. We seek $M = \text{Arg Max}_M p(M|I)$, i.e., the model with the greatest probability density given the image. This formulation is frequently called the posterior density, so the method is called one of posterior optimization.²⁶ An application of Bayes rule together with an application of the logarithm to both sides of the equation and removing terms that are constant in M yields $M = \text{Arg Max}_M [\log p(M) + \log p(I|M)]$; $\log p(M)$ then measures the *geometric typicality*, and $\log p(I|M)$ measures the *image match*. Other choices for these two terms are possible; but if one assumes Gaussian distributions, it is natural to use squared distances as proxies for $\log p$ since the log of a normal distribution is linear in a squared distance. The remainder of this section describes how we go about designing the distance terms that are minimized in our energy function.

II.D. Internal geometric energies

Without statistics, estimating shape likelihood requires approximations specific to the characteristics of the underlying parameterization. The essential problem is that the underlying parameter space from which the models are drawn is much larger than the shape space that we are trying to cover. Therefore, we restrict the models generated during training to a compact and legal subspace where we are confident that features are in correspondence across the training population. To this end, we impose three purely geometric conditions on the posterior.

First, we impose a penalty against nonuniformity of samples, i.e., irregularity in the discrete medial grid. Second,

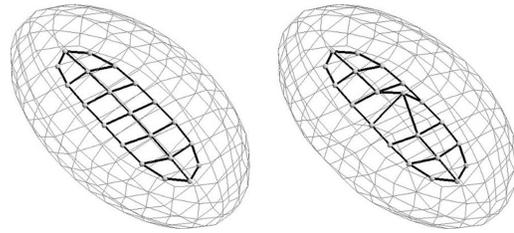


FIG. 6. A medial mesh (thick lines) and implied surface (thin lines) with a (left) high nonuniformity penalty and (right) low nonuniformity penalty. Meshes with high irregularity may imply similar surfaces as more regular meshes, but can result in qualitatively inferior results and break our volumetric correspondence assumptions.

we use a legality measurement based on differential geometry to identify and prevent local self-intersections. And third, we penalize shapes that are not near the presumptive mean. In addition, in recent trials a reward for model smoothness has been shown to be of use.

II.D.1. Sampling control

Establishing feature correspondence across the training population implies that there is a unique set of parameters that best describes each training case. However, it is possible to find two quite different models that have nearly the same image match in I , so we differentiate them by establishing an additional geometric criterion. As an example, there are many possible cubic approximations to a given function, but by preferring certain end conditions, we can identify a unique “best” approximating function.

In our case, we want a unique and legal best model to represent a given training image. That is, if shape models M and O both represent the same data, we need that $M \approx O$. Otherwise, the correspondence requirements will not be met. The geometric criterion that we use to distinguish medial sheets with similar implied boundaries is the uniformity of the medial grid. Medial sheets with samples bunched together form difficult to interpolate configurations that poorly capture the object. Following the Markov assumption that the likelihood of a sample conditioned on the model is the same as the likelihood of the sample conditioned on its neighbors, we define nonuniformity as a function of the agreement between any sample and the expectation of its neighbors. The total nonuniformity of the model is the sum of such sample agreements. Since we are producing a dissimilarity term, $\text{Nonuniform}(M)$ is actually defined as the distance $d^2(M, \bar{M})$, where $\bar{M} = M$ smoothed by a low-pass filter. As a side effect, this definition of smoothness also tends to act as an *ad hoc* enforcement for shape legality. Figure 6 shows an example of uniform and nonuniform organizations of medial samples.

In particular, in order to define regularity of a sampled medial representation, we use the Riemannian distances and means as described in the Appendix. We define \bar{M} as the model such that each medial sample \bar{m} is the distance-minimizing Fréchet mean of its neighboring samples. Practically, using our parameterization, \bar{M} is a curvature minimiz-

ing medial sheet with regularly spaced samples. $\text{Nonuniform}(M)$, the distance between M and its smoothed \bar{M} , can be reduced and rewritten as Eq. (1).

$$\text{Nonuniform}(M) \sim \sum_{m_i \in M} \sum_{m_j \in \text{Neighbrs}(m_i)} d^2(m_i, m_j) \quad (1)$$

This simple function, which essentially normalizes volume per respective medial sample point, is suitable to our parameterization because medial representations inherently imply desirable types of correlated surface deformation such as bending, twisting, and magnification.

II.D.2. Legality

Correspondence assumptions can also be undermined by illegal shapes. As we perturb the model parameters to find the best fit to a segmented image, the model's surface may fold or develop self-intersections. Medial representations have very precise analytic forms for computing shape legality on the medial manifold.³⁹ The largest eigenvalue of the radial shape operator yields a legality measure at every sample point with a well defined threshold for identifying folded surfaces.⁴⁰ As we narrow down the shape space during the fitting procedure, we tend to relax $\text{Nonuniform}()$ and rely more on the legality measure. Similar smoothness and legality functions can be mathematically derived for other shape representations specifying local orientations.

The penalty for illegality for any medial sample is set to zero when the largest eigenvalue of the radial shape operator is non-negligibly less than its threshold. Then it rises first slowly and then sharply as that eigenvalue passes its threshold. The overall penalty is the sum of the penalties of the samples.

II.D.3. Reference model

Given a space full of possible shapes, we desire to identify the mean shape and expand our shape model about it. Under Gaussian assumptions, given a mean and the variances of the principal directions of deformation, this method provides a shape-normalized Mahalanobis distance from any particular shape to the mean shape. Using this as the geometric prior in optimization tends to keep the candidate models clustered near the mean where our statistical shape model is most likely to be valid.

However, with limited *a priori* knowledge, we use our initializing model R_0 as a tentative reference point for the shape space. In the early iterations of model fitting we assume an isotropic shape probability distribution, i.e., that the difference between two shapes is measured only by a weighted Euclidean distance, we can estimate a geometric prior as the Riemannian model-to-model distance between a candidate model, M , and R_0 , with distance defined as in the previous section.

$$\text{Reference}(M, R) \sim \sum_{m_i \in M} d^2(m_i, r_i) \quad (2)$$

As both of the geometric penalties, $\text{Nonuniform}()$ and $\text{Reference}()$, are in the form of a distance, they can equally be

thought of as log probabilities on Gaussian distributions in the feature space. Taken together with the illegality penalty, they comprise the internal geometric energy in our optimization metric.

II.E. External image energies

The external image terms of the energy function measure how well the model's surface fits the label image. This is measured in two terms, one that computes distances between specific landmark points in the model to corresponding landmarks identified in the image, and another that computes a general surface-to-surface distance between the model surface points and the nearest boundary voxel in the label image.

II.E.1. Landmarks

The landmarks identified automatically or manually in Sec. II A require an explicit expression in the metric in order to prevent them from drifting out of alignment during optimization. We assume that the model coordinate landmarks should be normally distributed about the landmarks identified in the image, that is, the model landmarks have the distribution $N(LI, \rho^2)$, where ρ_i is the standard deviation, i.e., tolerance, for the i th landmark. Because the log probability of $N(LI, \rho^2)$ is simply $\rho^{-2} d^2(LI, LM)$ with d^2 the standard Euclidean square distance, we can express the error in landmark match as Eq. (3).

$$\text{Landmark}(LM, LI) \sim \sum_{lm_i \in LM} \frac{1}{\rho_i} d^2(lm_i, li_i). \quad (3)$$

II.E.2. Binary image match

The image match term is computed as the sum of squared distances between the boundary voxels, B , of the label image, found using a six-connected neighbor test, and the continuous boundary surface, Ω , implied by a candidate model M (see Appendix). This is equivalent to thinking of each surface point or boundary voxel as an equally weighted landmark corresponding to its nearest point in the other set.

Ideally we desire to measure two terms, the distance from each member of Ω to the closest voxel in B , denoted $d^2(\Omega, B)$, and the distance from each member of B to the closest surface point on Ω , denoted $d^2(B, \Omega)$, which are not equivalent. The distance $d^2(\Omega, B)$ can be computed fairly quickly across many candidate models because the label boundary remains static over all the trials. We generate a single space filling lookup table for distance from the label boundary by a modified anisotropic version of Danielsson's algorithm.⁴¹ Trilinear interpolation gives a very fast measure of the distance at any point in space to the closest boundary point on B . Discrete samples ω are taken arbitrarily densely from each candidate Ω and we then let $d(\omega, B)$ be the lookup of the position of ω in the distance map. Integrating over all surface samples gives $d^2(\Omega, B) = \sum_{\bar{\omega} \in \Omega} d^2(\bar{\omega}, B)$.

The inverse distance $d^2(B, \Omega)$, however, is computationally exorbitant given the finely sampled subdivision surfaces

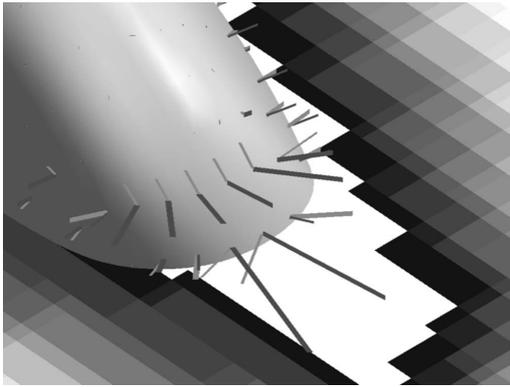


FIG. 7. A slice from a distance map and a suboptimally fit surface illustrating θ . The light gray lines show the distance map gradient direction; dark gray lines show the surface normal direction.

required for accurate matches and the large number of candidate surfaces generated for optimization. However, simply ignoring this part of the term can lead to undesirable results in areas of high curvature. In these areas, the $d^2(B, \Omega)$ distance-minimizing Ω tends to be more volume filling than the minimizer of the $d^2(\Omega, B)$ distance. Our solution is to compute the $d^2(B, \Omega)$ distance at only a minimal number of points where we would expect large distance asymmetries. We identify such points on Ω by computing the angle between the gradient of the distance map and the surface normal. If the angle is greater than a threshold, we compute a new distance along the surface normal at that point as can be seen in Fig. 7. With $\theta^2(x, y)$ as this modified minimum distance function, then our data likelihood term can be expressed as in Eq. (4).

$$\text{Image}(M, I) \sim \sum_{\omega_i \in \Omega} \theta^2(\omega_i, B). \quad (4)$$

This modified distance map method is dramatically faster and produces superior results to the standard approach of understanding binary images using local edge detectors, such as the derivative of Gaussian filter. While edge detectors are a logical extension of the model-to-image match as it would be computed in a grayscale image, they suffer from significant problems with capture range and orientation.

An additional advantage to using a medial parameterization for M is that M and B are strongly related according to morphological erosion. To fit models to images with structures that are only a few voxels in thickness, we can fit an initially dilated model to a dilation of the labeling, and then contract the model surface by the same amount by an inverse scaling of the thickness parameter. As seen in Fig. 8, this morphologically closed model approximates the thin object much better than is possible otherwise.

In the Bayesian framework, the image term, $P(I|M)$ is the probability that the surface of the model is in alignment with the boundary voxels of the segmentation. The landmark and binary image match energies we describe can be thought of as the log probabilities of a joint distribution and added together into the optimization metric.

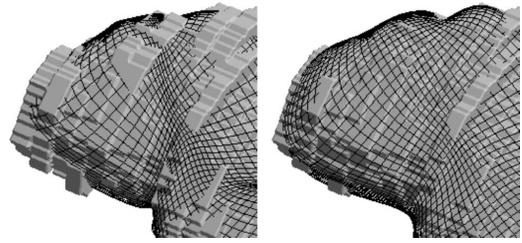


FIG. 8. Two candidate model meshes compared to tiled surface of a segmentation of the thin masseter muscle in the neck. The mesh on the right has been fit naively; the better fitting mesh on the left has been fit to a dilated image and then contracted.

II.F. Optimization framework

We now have an initial geometric template, R_0 , that has been coarsely aligned to the image and a well defined non-statistical objective function which has as its minimum a model configuration which should be both well fit to the target image and qualitatively similar to other models fit to other images of the population. Our process is then to deformably fit each image in the training population about R_0 . The mean of the population of the fit models, R_1 , is computed along with an estimate of the principal modes of deformation. The population is then refit iteratively about R until R converges to a suitable mean of the training population when $R_{i-1} \approx R_i$.

Each step of the iteration is a search for the error minimizing M , with error, E , computed via our complete dissimilarity metric, that is, the sum of (1–4) and the Legal() term described in Sec. II D 2, using α , β , δ , and γ as relative weighting factors.

$$E = \alpha \text{Landmark} + \beta \text{ImageMatch} + \delta \text{Nonuniform} + \text{Legal} + \gamma \text{Reference}. \quad (5)$$

Our optimization engine uses a conjugate gradient descent.⁴² Because conjugate gradient search performs best given a relatively isotropic global minimum, some experimentation is required to fine tune the weighting factors to new shape studies. These weighting factors are essentially scaling between the Euclidean voxel-space distances and the Riemannian sample-space distances.

Before we have modes of shape change, we assume that samples can move about freely in their feature space, restricted only implicitly by the distance metrics in the geometric constraints. After the first round of fitting, we can compute intermediate shape statistics by looking at correlations in feature relationships across the training samples. For boundary models with implicitly Gaussian distributions, Ref. 12 proposes using principal component analysis, which relies on the eigenstructure of the feature covariance to describe correlated shape change. Reference 43 gives a relevant example of applying PCA to anatomic shape variability. M-reps and other representations with explicit orientational components cannot use PCA because distances between features that encode orientation and scale are not Euclidean. These representations are governed by a generalized PCA known as

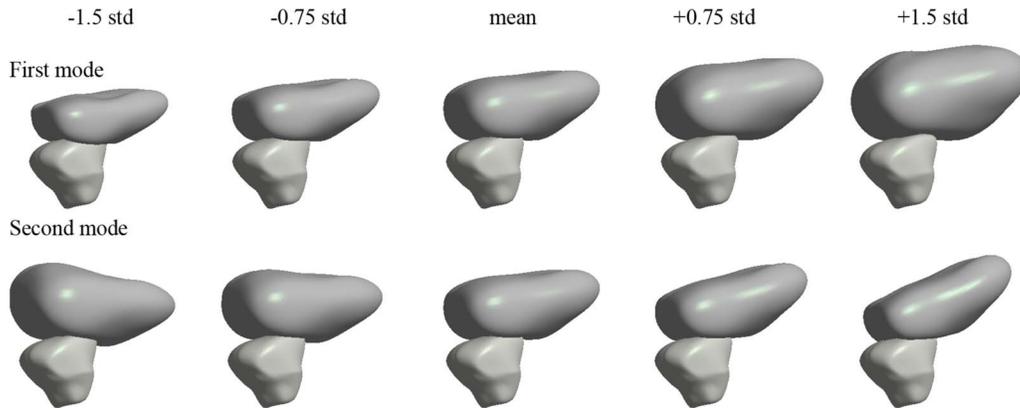


FIG. 9. A mean bladder (darker upper object) and its first two principal modes of deformation relative to the mean prostate (lighter lower object). These two modes of deformation together cover over 65% of the shape variability across the 18 images of this patient.

principal geodesic analysis (PGA).³⁵ As with PCA, an observation matrix relating measurements of corresponding features across the training population is formed and the covariance of the observation matrix is computed and diagonalized. However, the nonlinear orientation and size terms of our feature vectors are first mapped onto a tangent plane fit locally to the Fréchet mean of the shape space, R . The eigenvectors of the system orthogonal basis of correlated feature changes which are projected back into the nonlinear feature space and imply correlated surface deformations such as those shown in Fig. 9. This basis can be truncated according to eigenvalue thresholds, which describe the significance of each of direction of shape change. See the Appendix for details of how to decompose a set of m -reps according to PGA and how to apply the resulting statistics to individual instances.

As we iterate and gain confidence in our intermediate statistics, we replace our sampled medial shape parameterization with a short vector of coefficients of each principal mode of deformation. This restricts the shape changes we optimize over and both obviates our smoothness requirement and allows us to compute the geometric prior by the correct shape-to-mean Mahalanobis distance. This covariance weighted distance is computed directly as the eigenvalue scaled PGA coefficients. That is, the Mahalanobis distance squared $= \sum \alpha_i^2 / \lambda_i^2$, where α_i are the coefficients of M expressed in PGA eigenvectors and λ_i are the corresponding eigenvalues.

Following an intermediate step of statistical shape optimization, local refinements may still be required for individual

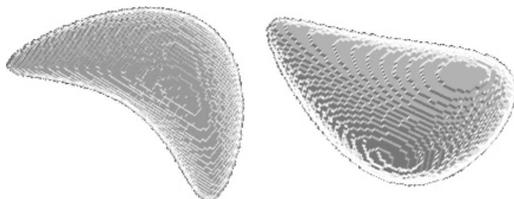


FIG. 10. Tiled surfaces from two procedurally generated warped ellipsoid test objects showing bending and tapering.

samples within the model. These refinements can be computed deterministically, by another optimization of our error metric (5) restricted to the parameters of a given sample, or they can be probabilistically estimated by PGA over the differences between the results of the statistical stage and the deterministic refinement. Tracking the residual optimization work on individual samples still required after an intermediate statistical fitting gives us a measurement of when the training is complete.

Extensions of the statistical model allow us to compute joint statistics of interdependent multiple figure shapes⁴⁴ and multiple object shape models, which leads to the multiscale methods for describing both global and local phenomena discussed in our image segmentation paper.¹²

III. RESULTS

Our method is routinely tested by computing a shape template and modes of deformation for a set of procedurally generated binary 3D images of bent, twisted, and tapered ellipsoids such as those in Fig. 10. A set of images is generated by randomly sampling values for each of the three parameters, analytically generating the corresponding ellipsoid and scan converting it into voxels. Vertex landmarks are computed for each image. The medial manifold of a standard ellipsoid is a simple primitive in our framework. This standard ellipsoid is designated R_0 . R_0 is initially aligned by landmarks to each of the training cases, and then the parameters of each medial sample are optimized according to the error function described in our method section. When the best model for each image has been found, we compute the final average and maximum voxel distances from each model to its corresponding image as shown in Fig. 11. Figure 12 is a histogram of these distances over the standard 20 test cases. Shape statistics as in Fig. 13 are computed and we refit the population by optimizing over the principal geodesics of shape change. Note that PGA has substantially reduced the complexity of the optimization. We input a three parameter distribution into the problem, our initial fitting optimizes on the order of 100 parameters, and using PGA we can reconstruct over 90% of the shape variability using only five



FIG. 11. Relationship between the boundary voxels of an ellipsoid binary training image (gray) and the fitted model's surface (black) through a transaxial slice.

model parameters. The distance histogram of the statistically fit models is also given in Fig. 12. Another encouraging result is that while the two training outliers we see in Fig. 12 were thrown out of the statistical analysis, they have been adequately covered by the statistical fitting.

The ultimate indication of our methodology's effectiveness is its application to real medical and scientific problems, such as modeling shape variability in inter-patient and cross-patient images. In Fig. 3, we show a mean bladder being initialized and deformably fit to a patient image, represented by a single sagittal slice of the label and corresponding gray image. In this case, we start with a model hand fit to day 1 of each patient series. Subsequent daily images are aligned according to the position and orientation of the urethral landmarks. We then proceed with two iterations of the training as outlined in our method, first a nonstatistical and then a coarse statistical stage. This procedure is repeated for the bladder and prostate. Figure 14 shows a histogram of average and maximum distances over models fit to bladders, prostates, and rectums from sets of inter-fractional images of 25 different patients. The average error across all organs is less than a voxel and the large maximum distance seen in the rectum training is artifactual of segmentation variability. The bladder is a particularly challenging object to model effectively because it exhibits large shape changes from day to day. The first two PGA modes of shape change for one patient's bladder and prostate are shown in Fig. 9. These two modes account for over 65% of the patient's daily shape variability in those two organs, and as we would hope, could be seen as

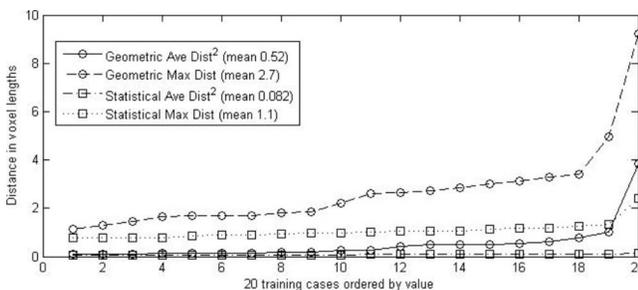


FIG. 12. Histogram of average and max distances for warped ellipsoid models over 20 training cases. The two outliers were excluded from the first round statistics, but were successfully fit in the next round using the recovered statistical modes of deformation.

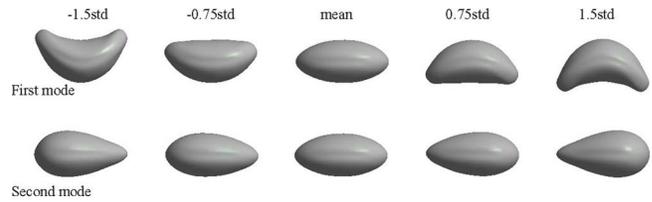


FIG. 13. As we would expect, the first two principal modes of deformation trained from 20 bent, twisted, and tapered ellipsoids reflect bending and tapering.

roughly accounting for the bladder filling and emptying. The statistical deformable shape models for each patient from such training, combined with a suitable CT image match term, allowed us to produce leave-one-out segmentations of actual images from four of the patients with mean volume overlaps (intersection/average) of 92.5%, 93.7%, 91.6%, and 94.7% respectively using only eight modes of deformation.⁴⁵

We also applied our method to a set of 50 cross-patient kidney images. In this case, the SDSM must account for actual anatomic variability across patients, a much broader kind of shape change than that usually seen in sets of within-patient images. The eigenvalues of the recovered principal modes of deformation show that we now need 15 modes of variance cover 95% of the anatomic variability. Our training resulted in 95.3% volume overlaps (int/ave) of the models to the segmentations. The trained 15 parameter SDSM applied to leave one out segmentation of the actual patient images resulted in models with volume overlaps only slightly lower, on average, 91.1%. For details of our kidney segmentation results, see Refs. 46–48.

Our method has also been applied to several other target areas. Head and neck models suitable for radiotherapy planning are shown in Fig. 15, left. The deep brain structures shown Fig. 15, right, were taken from an autism shape study where the researcher used the SDSM based shape-to-shape metric to discriminate autism by shape characteristics.¹¹ We have also modeled hippocampi for a large statistical shape study. Temporal extensions to the method have enabled studies of heart motion,⁴⁹ which has obvious extensions to lung motion and 4D-ART. Reference 50 uses a SDSM trained with our method to establish feature correspondences in dog hip joints for a longitudinal study characterizing canine hip dysplasia.

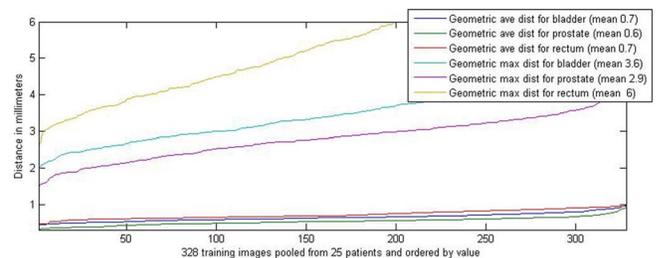


FIG. 14. Histogram of average and max distances for fit bladder, prostate, and rectum models over 328 training cases pooled from 25 sets of same patient inter-fractional images.

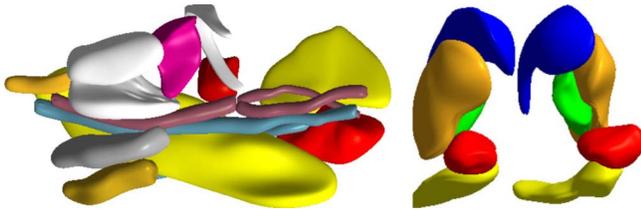


FIG. 15. Multi-object shape models. (Left) A 15 object complex of structures from the head and neck. (Right) Deep brain structures from an autism study, left and right hippocampus, amygdala, putamen, caudate, and globus pallidus.

All the results discussed were produced using a C++ implementation of the algorithm. Using our implementation on a modern 2 GHz desktop, a model can be fit to an image in less than 2 min. Batch fitting is trivially parallelizable onto any number of machines up to the number of images in the training population. Using such a parallel computing framework and given a set of segmented images and landmarks, a completely trained SDSM can be produced in less than 1 h.

IV. DISCUSSION

IV.A. Applications to IGRT/ART

The model-fitting methodology described in this article and the shape statistics computed from the resulting models are being routinely used in ongoing research involving segmentation of CT images acquired for IGRT and ART. An overview of this work in progress is presented here and greater detail will be provided in a future article. The main idea behind CT-guided radiotherapy is to extract anatomic information from images acquired immediately before or during treatment to guide patient positioning and perhaps beam shaping. In ART, treatment images are used to calculate delivered doses for comparison with the original treatment plan, followed when necessary by replanning.¹ Image segmentation is not required when the only goal is to position the isocenter within the target volume. Our research currently is focused on segmentation, via posterior optimization using *m*-reps as discussed, of target organs and organs at risk in the male pelvis as needed for IGRT beam shaping and ART dose calculation. So far we have investigated segmentation of the prostate, bladder, and rectum from conventional kV fanbeam CT images. This work makes use of the training methods and tools described in this article and more recently an image processing tool called ConStruct,⁵¹ which applies the resulting shape statistics to segment target images via posterior optimization.

We distinguish between two modes of segmentation based on the method for statistical training, i.e., cross-patient (*xpat*) versus within-patient (*inpat*) training. When a random target image is segmented, e.g., the planning CT image, the shape statistics for that patient are unknown. In this case only *xpat* statistics can be applied. These shape statistics describe geometric variation from one patient to another and comprise a mean shape with principal modes of variation computed

from collections of models fit as described in this article to expert human contours across many patients, one sample image per patient. Currently we train from 30–75 patients depending on experimental design. Segmentation quality is affected by a number of factors still being investigated for the male pelvis that are beyond the scope of this article to discuss in detail. Perhaps the most critical factors are the shape of the starting model and its initialization in the target image. If the model starts too far in shape and/or position from the target object, our gradient-descent optimization algorithm will get trapped in a local optimum. Moreover, when the target shape is far from the population mean, the geometric penalty is more likely to prevent the starting model from completely deforming to match the target shape. We note that these considerations apply to all posterior optimization methods. Our strategy for dealing with these issues is discussed later. The segmentation of a target object yields a patient-specific estimated mean shape that can be used as a starting mean for segmentation of subsequent target images, e.g., treatment CT images. In this case the principal modes applied to the estimated mean are computed from input shape statistics that capture day-to-day shape changes. These statistics are computed from collections of models fit to human-drawn contours across multiple sample images per patient, e.g., 10–20, and across many patients. Currently we are working with approximately 15 images each for 30 patients (~450 images total).

Our general approach for both IGRT and ART using ConStruct is to create a patient-specific estimated mean model for each target object on a reference image, preferably a planning image acquired without contrast medium to avoid artificial intensity patterns that will confound the image-match term of our objective function. If the planning image is available only with contrast medium, then the first day-treatment image serves as the reference. There are several ways to create estimated mean models. One method is to carefully hand segment the reference image, and then fit a model to the contours for each object via posterior optimization using population shape statistics for the geometric term and distance to contour points to compute the image match term. This can be implemented by performing only the final step of the shape training described here on a single image. Another method is to segment the gray-scale reference image directly with *m*-reps. One way to achieve good initialization of the starting *m*-rep for this approach is via posterior optimization as just described to deform the model to match points in the target image that are known to be on the surface of a target object, followed by fully automatic segmentation using gray-scale intensities to compute image match. We currently use several user-drawn contours per object, usually three, to define boundary points. This approach so far has yielded segmentations that are clinically acceptable in close to 90% of cases. When an automatic segmentation is unacceptable it can be edited with standard contour editing tools, and a model can be fit to the edited contours as just described. Although some level of user interaction likely will

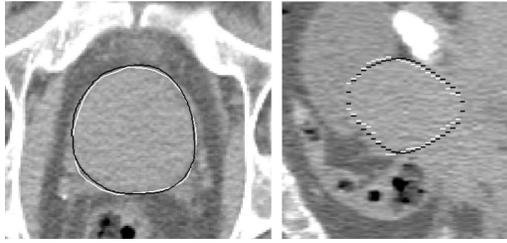


FIG. 16. Manual (white) and computed (black) segmentations of the prostate in a treatment image, trans-axial slice on the left, sagittal slice on the right. The segmentation computed based on shape training described in this article agrees with the manual segmentation for this day agree with an 89% int/ave volume overlap.

be required for creating reference models, interaction is undesirable and we are looking into ways to minimize it, especially for segmentation of treatment images.

When patient-specific models are created, ConStruct initializes the prostate in a treatment image by registration with the reference image, eliminating user interaction to define surface points. The registration is multiscale and concludes by focusing on a small region of interest including the prostate and its immediate surrounds. Fully automatic segmentation follows by applying inpat shape statistics to deform to model. Since the starting position and shape are close to the target, the final prostate segmentation compares well with expert human contours, as seen in Fig. 16. Because of the very large day-to-day variations of the shapes and intensity patterns associated with the bladder and rectum, initialization via registration alone often yields regions that do not agree well with human contours. These unacceptable regions are usually distant from the prostate, that is, out of the treatment beams. Agreement proximal to the prostate often is clinically acceptable and therefore useful for beam shaping and dose calculation. At the present time, user-identified surface points are needed to achieve good agreement away from the prostate.

IV.B. CONCLUSIONS

Shape training is a key step in the application of deformable shape models to such problems as image registration required for IGRT or ART, shape classification, and longitudinal shape studies. Our novel methodology for training statistical deformable models is a special case of image seg-

mentation, where both the best geometric model for any given image and the shape distribution for the entire population of training images is jointly estimated by iteratively relaxing purely geometric constraints in favor of the converging shape probabilities. The internal geometric terms that we use in the absence of *a priori* shape probabilities are crafted to guarantee training shapes that are regularly sampled, legal, and compact about a mean, which together cause credible shapes to be reflected in the converging statistics. The method described is fast and routinely gives good meshes for a variety of target shapes. The framework has been described in the context of the discrete medial representation, but should extend to b-reps with orientation or other representations where measures for sampling regularity and nonself-interpenetration can be analytically computed.

ACKNOWLEDGMENTS

The authors gratefully acknowledge our many collaborators for access to their training data, and in particular, Chi Villaruel for the synthetic ellipsoid data, Joshua Stough and Eli Broadhurst for the CT image segmentation results of the male pelvis and kidney, and Kevin Gorczowski for the brain structure models. This work was supported by NIH Grant Nos. P01 EB02779 and R01 RR0181615.

APPENDIX: THE DISCRETE MEDIAL REPRESENTATION

Medial geometry⁵² describes 3D objects in terms of a skeletal surface, a 2D curved sheet lying midway between opposing surfaces of the object, and a set of spokes extending to the object boundary from both sides of the skeletal surface. The medial manifold, \mathbf{M} , of a three-dimensional object has eight parameters at each point $(u,v): \mathbf{M}(u,v) = \{\text{position}(3), \text{spoke length}(1), \text{and two spoke directions}(2 \times 2)\}$. Some additional complexity is introduced along the crests at the edges.

The discrete medial representation, m-reps, samples the continuous manifold on a grid, yielding a set of eight-dimensional medial samples which taken together act as control points for the object's volume, as shown in Fig. 17. Additional medial points can be interpolated according to Ref. 53, which in turn imply a denser surface sampling. Alternatively, additional surface points can be approximated directly using a modified Catmull–Clark subdivision

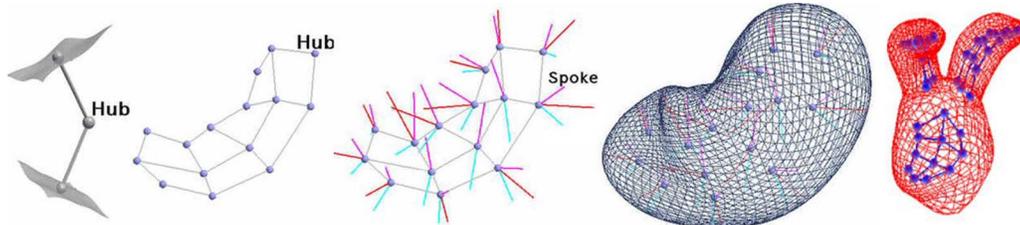


FIG. 17. Discrete medial representations. (Left) A medial sample with two equal length spokes that touch opposing surface patches. (Mid Left) A sampled skeletal sheet for a kidney with neighbor relations marked. (Middle) Spokes at each medial sample describe the orientation of the implied surface at that hub. (Mid Right) A densely sampled surface can be interpolated from the medial samples. (Right) A prostate model with subfigures defined for the left and right seminal vesicles.

algorithm⁵⁴ with additional normal constraints. \mathbf{M} also implies a volume filling hexahedral mesh useful for computing mechanical deformations according to finite element methods.⁵⁵ An object made from a single grid of medial samples is called a *figure*. A single column grid implies a tube figure; a multi-column grid implies a slab figure. Indentations and protrusions are handled as attached subfigures.⁴⁴ A figure along with any associated subfigures is called a *model*, shown in Fig. 17 (right).

The discrete medial representation is well suited to Bayesian segmentation because it provides (1) a volumetric coordinate system in which to gather image intensity statistics, (2) a low dimensional parametric shape space amenable to numerical optimization, and (3) an analytic description of legality. Each figure can be examined sample by sample where we need only eight parameters to represent complex object-based deformations such as bending, twisting, and magnification. Parameter changes can be constrained to imply only legal models by computing a radial shape operator similar to the surface shape operator described in Ref. 56, but dependent both on local shape and on local orientation. Measuring the largest eigenvector of the spoke length radial shape operator at each grid position^{39,40} gives an immediate indication of local self-interpenetration in the volume.

1. Distances and means of samples and models

Every medial sample can be understood as a translation, magnification, and rotation of any other. Thus, we can define Riemannian distances between them for computing means and variabilities. The spoke directions describe a rotation, which along with the radius, are multiplicative terms which must be logarithmically mapped so that sample-to-sample distances can be computed in a Euclidean space.³⁵ The Fréchet mean, \bar{m} of a set of samples $\{m_i\}$ can be computed as the point in the parameter space that minimizes its distance to every member of $\{m_i\}$. Samples on the grid are also given neighbor relationships, which along allows for Markov fashion predictions (see Sec. II D 1). Sample-to-sample distances can be extended to a distance between two models, M and O with samples $\{m_1, \dots, m_n\}$ and $\{o_1, \dots, o_n\}$, respectively, by taking the sum of the Riemannian distances between corresponding samples as $d^2(M, O) = \sum_{m_i \in M} d^2(m_i, o_i)$. Distances taken from object to object and then refined sample by sample provide the basis for multi-scale shape analysis.

Using these definitions for distance, a typical shape and descriptions of shape variability can be computed from a set of models according to a generalizing of principal component analysis called principal geodesic analysis.³⁵ The mean for each atom over a training population is computed by the Fréchet approach of minimizing the sum of squared geodesic distances to the result point. The atoms in each training m-rep are projected (mathematically called the Log map) to a tangent space at the computed mean, where PCA is done on the collection of atoms. Mapping the resulting principal direction vectors back into an m-rep increment from the mean creates modes of global variability that encode not only position but orientation and width. The relative weight of the

associated eigenvalues give a reciprocal standard deviation weighting to this space of deformations and provide the basis for a shape-to-shape Mahalanobis distance. The Mahalanobis distance from the mean gives the log likelihood of the shape, except for a linear multiplier and an additive constant and thus it serves as the basis for methods for Bayesian segmentation or statistical shape analysis.

See Refs. 12 and 48 for additional details of the discrete medial shape representation and its applications to medical image analysis.

^{a)} Author to whom correspondence should be addressed. Electronic mail: derek@cs.unc.edu

¹D. Yan, D. Lockman, D. Brabbins, L. Tyburski, and A. Martinez, "An off-line strategy for constructing a patient-specific planning target volume in adaptive treatment process for prostate cancer," *Int. J. Radiat. Oncol., Biol., Phys.* **48**(1), 289–302 (2000).

²J. R. Wong, L. Grimm, M. Uematsu, R. Oren, C. W. Cheng, S. Merrick, and P. Schiff, "Image-guided radiotherapy for prostate cancer by CT-linear accelerator combination prostate movements and dosimetric considerations," *Int. J. Radiat. Oncol., Biol., Phys.* **61**, 561–569 (2005).

³W. Lu, M.-L. Chen, G. Olivera, K. Ruchala, and T. Mackie, "East free-form deformable registration via calculus of variations," *Phys. Med. Biol.* **49**, 3067–3087 (2004).

⁴H. Wang, I. Dong, M. Lii, A. Lee, R. De Crevoisier, R. Mohan, J. Cox, D. Kuban, and R. Cheung, "Implementation and validation of a three-dimensional deformable registration algorithm for targeted prostate cancer radiotherapy," *Int. J. Radiat. Oncol., Biol., Phys.* **61**, 725–735 (2005).

⁵R. Mohan, X. Zhang, H. Wang, Y. Kang, X. Wang, H. Liu, K. Ang, D. Kuban, and L. Dong, "Use of deformed intensity distributions for on-line modification of image-guided IMRT to account for interfractional anatomic changes," *Int. J. Radiat. Oncol., Biol., Phys.* **61**(4), 1258–1266 (2005).

⁶M. Foskey, B. Davis, L. Goyal, S. Chang, E. Chaney, N. Strehl, S. Tome, J. Rosenman, and S. Joshi, "Large deformation 3D image registration in image-guided radiation therapy," *Phys. Med. Biol.* **50**, 5869–5892 (2005).

⁷D. Freedman, R. Radke, T. Zhang, Y. Jeong, M. Lovelock, and G. Chen, "Model-based segmentation of medical imagery by matching distributions," *IEEE Trans. Med. Imaging* **24**(3), 281–292 (2005).

⁸V. Pekar, T. McNutt, and M. Kaus, "Automated model-based organ delineation for radiotherapy planning in prostatic region," *Int. J. Radiat. Oncol., Biol., Phys.* **60**, 973–980 (2004).

⁹T. Klinder, C. Lorenz, J. von Berg, S. P. M. Dries, T. Bülow, and J. Ostermann, "Automated model-based rib cage segmentation and labeling in CT images," *Medical Image Computing and Computer-Assisted Intervention (MICCAI) Conference*, Brisbane, Australia, 2007, pp. 195–202.

¹⁰W. E. Grimson, R. Kikinis, F. A. Jolesz, and P. M. Black, "Image-guided surgery," *Sci. Am.* **280**(6), 62–69 (1999).

¹¹M. Styner, K. Gorczowski, T. Fletcher, J. Y. Jeong, S. M. Pizer, and G. Gerig, "Statistics of pose and shape in multiobject complexes using principal geodesic analysis," *Medical Imaging and Augmented Reality Conference*, Shanghai, China, 2006, LNCS 4091, pp. 1–8.

¹²S. M. Pizer, P. T. Fletcher, S. Joshi, G. Gash, J. Stough, A. Thall, G. Tracton, and E. L. Chaney, "A method and software for segmentation of anatomic object ensembles by deformable M-Reps," *Med. Phys.* **32**(5), 1335–1345 (2005).

¹³T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam, "The use of active shape models for locating structures in medical images," *Information Processing in Medical Imaging (IPMI) Conference*, Flagstaff, Arizona, 1993, pp. 33–47.

¹⁴J.-P. Thirion, "Image matching as a diffusion process: An analogy with Maxwell's demons," *Med. Image Anal.* **2**(3), 243–260 (1995).

¹⁵B. K. Horn and B. G. Schunck, "Determining optical flow," Technical Report No. AIM-572, Massachusetts Institute of Technology (1980).

¹⁶J. C. Gee, M. Reivich, and R. Bajcsy, "Elastically deforming 3D atlas to match anatomical brain images," *J. Comput. Assist. Tomogr.* **17**(2), 225–236 (1993).

¹⁷G. E. Christensen, M. I. Miller, and M. Vannier, "3D Brain mapping using a deformable neuroanatomy," *Phys. Med. Biol.* **39**, 609–618 (1994).

- ¹⁸S. Joshi, B. Davis, M. Jomier, and G. Gerig, "Unbiased diffeomorphic atlas construction for computational anatomy," *Neuroimage* **23**(Suppl 1), 151–160 (2004).
- ¹⁹K. K. Brock, D. L. McShan, R. K. Ten Haken, S. J. Hollister, L. A. Dawson, and J. M. Balter, "Inclusion of organ deformation in dose calculations," *Med. Phys.* **30**(3), 290–295 (2003).
- ²⁰K. K. Brock, S. J. Hollister, L. A. Dawson, and I. M. Balter, "Technical note: Creating a four-dimensional model of the liver using finite element analysis," *Med. Phys.* **29**(7), 1403–1405 (2002).
- ²¹K. E. Muller, Y.-Y. Chi, J. Ahn, and J. S. Marron, "Limitations of high dimension, low sample size principal components for Gaussian data," under revision for resubmission.
- ²²IEEE Trans. Med. Imaging Imaging, Special Issue on Model-Based Analysis of Medical Images, **18**(10), (1999).
- ²³K. Haris, S. N. Efstratiadis, N. Maglaveras, C. Pappas, J. Gourassas, and G. Louridas, "Model-based morphological segmentation and labeling of coronary angiograms," *IEEE Trans. Med. Imaging*, Special Issue on Model-Based Analysis of Medical Images, **18**(10), 1003–1015 (1999).
- ²⁴T. McInerney and D. Terzopoulos, "Topology deformable surfaces for medical image volume segmentations," *IEEE Trans. Med. Imaging*, Special Issue on Model-Based Analysis of Medical Images, **18**(10), 840–850 (1999).
- ²⁵D. J. Hawkes, D. Barratt, J. M. Blackall, C. Chan, P. J. Edwards, K. Rhode, G. P. Penney, J. McClelland, and D. L. Hill, "Tissue deformation and shape models in image-guided interventions: A discussion paper," *Med. Image Anal.* **9**(2), 163–175 (2005).
- ²⁶R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. (Wiley-Interscience, New York, 2000).
- ²⁷S. Marsland and C. J. Twining, "Constructing diffeomorphic representations for the groupwise analysis of nonrigid registrations of medical images," *IEEE Trans. Med. Imaging* **23**(8), 1006–2010 (2004).
- ²⁸J. Cates, P. T. Fletcher, M. Styner, M. Shenton, and R. Whitaker, "Shape modeling and analysis with entropy-based particle systems," *Inf. Process Med. Imaging* **20**, 333–345.
- ²⁹T. Cootes, C. Taylor, D. Cooper, and J. Graham, "Active shape models – their training and application," *Comput. Vis. Image Underst.* **1**, 38–59 (1995).
- ³⁰C. Brechbuhler, G. Gerig, and O. Kubler, "Parameterization of closed surfaces for 3-D shape description," *Comput. Vis. Image Underst.* **61**, 195–170 (1995).
- ³¹A. Kelemen, G. Szekely, and G. Gerig, "Elastic model-based segmentation of 3-D neuroradiological data sets," *IEEE Trans. Med. Imaging* **18**(10), 828–839 (1999).
- ³²D. Yan, D. A. Jaffray, and J. W. Wong, "A model to accumulate fractionated dose in a deforming organ," *Int. J. Radiat. Oncol., Biol., Phys.* **44**(3), 665–675 (1999).
- ³³H. Delingette and J. Montagnat, "Shape and topology constraints on parametric active contours," *Comput. Vis. Image Underst.* **83**(2), 140–171 (2001).
- ³⁴M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *Int. J. Comput. Vis.* **1**, 321–331 (1987).
- ³⁵P. T. Fletcher, L. Conglin, S. M. Pizer, and S. Joshi, "Principal geodesic analysis for the study of nonlinear statistics of shape," *IEEE Trans. Med. Imaging* **23**, 995–1005 (2004).
- ³⁶S. K. Warfield, K. H. Zou, and W. M. Wells, "Simultaneous truth and performance level estimation (STAPLE): An algorithm for the validation of image segmentation," *IEEE Trans. Med. Imaging* **23**(7), 903–921 (2004).
- ³⁷F. L. Bookstein, "Principal warps: Thin-plate splines and the decomposition of deformations," *IEEE Trans. Pattern Anal. Mach. Intell.* **1**(6), 567–585 (1989).
- ³⁸F. L. Bookstein, *Morphometric Tools for Landmark Data* (Cambridge University Press, Cambridge, 1991).
- ³⁹J. Damon, "Geometry and Medial Structure," in *Medial Representations: Mathematics, Algorithms and Applications*, edited by K. Siddiqi and S. Pizer (Springer, Heidelberg, Germany, in press, 2007).
- ⁴⁰Q. Han, D. Merck, J. Levy, C. Villarruel, E. Chaney, and S. M. Pizer, "Geometrically proper models in statistical training, Presented at *Information Processing in Medical Imaging (IPMI) Conference*, edited by Nico Karssemeijer and Boudewijn Lelieveldt, Kerkrade, The Netherlands, 2007.
- ⁴¹P. E. Danielsson, "Euclidean distance mapping," *Comput. Graph. Image Process.* **14**(3), 227–248 (1980).
- ⁴²W. H. Press, W. T. Vetterling, S. A. Teukolsky, and B. P. Flannery, *Numerical recipes in c++: The art of scientific computing* (Cambridge U.P., 2002).
- ⁴³M. Söhn, M. Birkner, D. Yan, and M. Alber, "Modeling individual geometric variation based on dominant eigenmodes of organ deformation: Implementation and evaluation," *Phys. Med. Biol.* **50**, 5893–5908 (2005).
- ⁴⁴Q. Han, S. M. Pizer, D. Merck, S. Joshi, and J.-Y. Jeong, "Multi-figure anatomical objects for shape statistics," Presented at *Information Processing in Medical Imaging (IPMI) Conference*, Glenwood Springs, Colorado, 2005.
- ⁴⁵E. Chaney, S. Pizer, S. Joshi, R. Broadhurst, T. Fletcher, G. Gash, Q. Han, J. Jeong, C. Lu, D. Merck, J. Stough, G. J. Tracton, M. Bechtel, J. Rosenman, Y. Chi, and K. Muller, "Automatic male pelvis segmentation from CT images via statistically trained multi-object deformable m-rep models," Presented at *American Society for Therapeutic Radiology and Oncology (ASTRO) Conference*, 2004.
- ⁴⁶M. Rao, J. Stough, Y. Chi, K. Muller, G. Tracton, S. Pizer, and E. Chaney, "Comparison of human and automatic segmentations of kidneys from CT images," *Int. J. Radiat. Oncol., Biol., Phys.* **61**(3), 954–960 (2005).
- ⁴⁷R. Broadhurst, J. Stough, S. Pizer, and E. Chaney, "A statistical appearance model based on intensity quantile histograms," *Proceedings of the 2006 IEEE International Symposium on Biomedical Imaging*, Washington D.C.
- ⁴⁸S. Pizer, Q. Han, S. Joshi, P. T. Fletcher, P. Yushkevish, and A. Thall, "Synthesis, deformation, and statistics of 3D objects via m-reps," in *Medial Representations: Mathematics, Algorithms and Applications*, edited by K. Siddiqi and S. Pizer (Springer, Heidelberg, Germany, in press, 2007).
- ⁴⁹R. Pilgram, P. T. Fletcher, S. M. Pizer, O. Pachinger, and R. Schubert, "Common shape model and inter-individual variations of the heart using medial representation: A pilot study," Technical report, Institute for Medical Knowledge Representation and Visualization, University for Health Informatics and Technology, Tyrol, Austria, 2003.
- ⁵⁰W. S. Vanden Berg-Foels, S. J. Schwager, R. J. Todhunter, and A. P. Reeves, "Subluxation during development alters femoral head shape in a canine model of hip dysplasia," (Abstract #1847) *Orthopaedic Research Society Transactions* Vol. 33, San Francisco, CA, 2008, <http://www.ors.org/web/Transactions/54/1847.PDF>.
- ⁵¹M. Foskey, A. Gash, Q. Han, G. Tracton, S. Joshi, S. Pizer, and E. Chaney, "A software toolkit for multi-image registration and segmentation in IGRT and ART," *Med. Phys.* **34**, 2351 (2007).
- ⁵²H. Blum, "A transformation for extracting new descriptors of shape," in *Models for the Perception of Speech and Visual Form*, edited by W. Wathen-Dunn (MIT Press, Cambridge, 1967), pp. 363–380.
- ⁵³Q. Han, S. M. Pizer, and J. Damon, "Interpolation in discrete single figure medial objects," in *Proceedings of the 2006 Conference on Computer Vision and Pattern Recognition Workshop (June 17–22, 2006)*, CVPRW, IEEE Computer Society, Washington, D.C., 85.
- ⁵⁴E. Catmull and J. Clark, "Recursively generated b-spline surfaces on arbitrary topological meshes," *Comput.-Aided Des.* **10**, 183–188 (1978).
- ⁵⁵J. R. Crouch, S. M. Pizer, E. I. Chaney, and M. Zaider, *Medially Based Meshing with Finite Element Analysis of Prostate Deformation: Medical Image Computing and Computer-Assisted Intervention Proceedings*, Part I, edited by Randy E. Ellis, Terry M. Peters, 108–115 (Springer, New York, 2003), pp. 108–115.
- ⁵⁶J. J. Koenderink, *Solid Shape* (MIT Press, Cambridge, 1990).