Submitted for
Conference Review

A Coarse-to-fine Shape Prior for Probabilistic Segmentations Using A Deformable M-rep

Xiaoxiao Liu, Ja-Yeon Jeong , Joshua H. Levy,
Rohit R. Saboo, Edward L. Chaney , Stephen M. Pizer

Abstract

Training a shape prior has been potent scheme for anatomical object segmentations, especially for images with noisy or weak intensity patterns. When the shape representation lives in a high dimensional space, Principal Component Analysis (PCA) is often used to calculate a low dimensional variation subspace from frequently limited number of training samples. However, the eigenmodes of the subspace tend to keep the coarse variation of the shape only, losing the detailed localized variability which is crucial to accurate segmentations. In this paper, we propose a coarse-to-fine shape prior for probabilistic segmentation to enable local refinement, using a deformable medial representation, called the m-rep.

Tests on the goodness of the shape prior are carried out on large simulated data sets of a) 1000 deformed ellipsoids with mixed global deformations and local perturbation; b) 500 simulated hippocampus models. The predictability of the shape priors are evaluated and compared by a squared correlations metric and the volume overlap measurement against different training sample sizes. The improved robustness achieved by the coarse-to-fine strategy is demonstrated, especially for low sample size applications. Finally, posterior segmentations of bladder in 3D CT images from multiple patients in day-to-day adaptive radiation therapy validate the local residual statistics that are introduced by this method satisfactorily improves the segmentation accuracy.

1. Introduction

Automatic 3D medical image segmentation via deformable models is challenging; adoption occurs only if its accuracy is competitive with manual segmentations. Deformable models with various representations have been proven effective in capturing variation of a population of geometric entities [8]. It is important to collect the statistics about the shape space via proper training in order to constrain the deformation field, especially for processing noisy image patterns. Given a sizable training samples, an off-line learning process can build up the probability distribution of the shape variation [2, 11], namely the shape prior.

When having high dimensional data representations, the Principal Components Analysis (PCA) based methods are often used to extract a robust subspace with a fairly small number of dimensions as the shape prior. However, as also pointed out in [10], the trained shape variation space often can not reflect higher order dependencies of the object, which tends to produce oversmooth segmentations. Further, with often a small number of training samples, the extracted shape space is not robust in terms of limited predictability.

Considering the various degrees of localities for geometric features in any shape representations, an efficient way of representing the shape space is to decompose them into different scales. Multiscale spherical wavelets is a successful example[10] to address this issue via wavelet basis functions with shape localities. Here we propose a more general approach to "decompose" the variation space by rediscovering the residual statistics hierarchically, which tends to be applicable to various representations or parametrization of the deformable model with notions of localities.

A recent study shows medial representation has superior efficiency over boundary representation[7].We use a discrete medial representation method, called the *m-rep* (Fig.1) [16]. M-reps has been successfully applied to various medical image segmentation tasks[15] and statistical shape analysis[18]. And in terms of average closest point distance and Dice similarity coefficient, the m-rep segmentation is among one of the best in literature of CT segmentation for male pelvic organs [3, 13, 5, 12, 17]. However, the global m-rep segmentation results are not accurate enough due to the oversmooth boundaries, which often are several voxels off the manually segmented boundaries in local regions. Joshi [6] et al. developed a multiscale m-rep segmentation where the global scale stage optimize the posterior probability of the shape with trained statistics, and the local scale stage is a deterministic optimization on the full space of atom deformations. However, without shape prior information, the local scale deformation can go wild with low-contrast intensity patterns and the ad hoc weight-

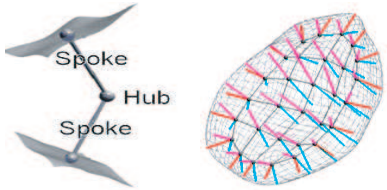

Figure 1. An m-rep model: An interior medial atom (left) with two spokes and an object (right). The object is composed of multiple medial atoms. Interior atoms have two spokes each. Exterior atoms (on the crest region) have three spokes each.

ing penalties are tedious to optimize and lack of validation.

The m-rep model of objects with simple geometry usually contains a sparse grid of atoms, which naturally provide the local scale defined by each of the atoms. A two-scale (object scale and atom scale) probabilistic segmentation scheme via m-rep is developed in this paper as an important step towards the ideal coarse-to-fine strategy, in which more careful studies on the number of scales and the size of each scale are needed.

The rest of the paper is organized as follows: Section 2 introduces the coarse-to-fine shape prior via m-rep. And the robustness of the statistics against training sample size is tested on 1000 simulated warped ellipsoids with both global and local deformations and 500 hippocampus m-rep models sampled from 51 trained variation modes. The posterior segmentation methodology for applying shape statistics is presented in Section 3. In Section 4, segmentation results on 79 clinical CT image data of bladder are discussed in detail. Section 5 summarizes and discusses the method and results.

2. Coarse-to-fine Shape Space Analysis

2.1. Localities in M-rep

In single-figure m-rep models, an object is a sheet of medial atoms represented by a quadrilateral mesh (Fig.1), which carries geometric properties such as the widening, bending and tapering, with the locality scale given according to the grid spacing. Each atom controls a local boundary region implied by its spokes' ending points. The object as a whole gives these properties in a way reflecting global interrelations of atoms, whereas each atom, and its relation to its immediate neighbors, provides more local features.

For simple biological objects represented by a sparse grid of medial atoms, typically less than 50 atoms, object and atom scales are sufficient to describe the shape variation globally and locally. For rich geometrical structure that characterized by hundreds of atoms, it might be necessary to define more locality scales by grouping n-nearest neighboring atoms, where n controls the size of each scale. For a model with N atoms, we focus on the simple situation of $n = \{N, 1\}$ in this paper.

2.2. Coarse-to-fine shape prior

A m-rep model $\mathbf{m}$ lives in a curved space $\mathbf{G} = \mathbb{R}^3 \times \mathbb{R}^+ \times \mathbb{S}^2 \times \mathbb{S}^2$. The ith atom $\mathbf{A}_i$ in m is a vector of 8 dimensions:$(\mathbf{p}_i, \mathbf{r}_i, \mathbf{u}_i, \mathbf{v}_i)$, with the hub position $\mathbf{p}_i \in \mathbf{R}^3$, the two spokes have the same radius $\mathbf{r}_i \in \mathbb{R}^+$, and two unit spoke directions $\mathbf{u}, \mathbf{v} \in \mathbf{S}^2$. For the external atoms in the crest region, there is an extra spoke with a different length than the other two spokes, thus makes some of the atom vectors 9 dimensional.

The generalized version of PCA for nonlinear data, Principal Geodesic Analysis (PGA) [4], is used for training the shape priors. Given the training samples, we can compute a Fréchet mean and the principal geodesic eigenmodes for the curved shape space.

The coarse-to-fine shape prior is built by PGA on object and atom scales separately. But the atom prior is built upon the residual space of the object prior.

2.2.1 Object prior

For the object scale, the whole grid of atoms is taken as a big vector in the high dimensional shape space. For example, a model with 21 atoms in a 3 by 7 grid has 184 dimensions. The first several eigenmodes with big eigenvalues produced by PGA usually cover most of the variation from the training samples. By only using the subspace composed of the few eigenmodes, we successfully reduce the dimension of the shape space from several hundred to a number often less than 10. Moreover, with the often limited number of training cases, the eigenmodes with little variance are not robust against the large dimensions of the object representation, which necessitates limiting the number of eigenmodes used in the object scale and training the residual statistics in a local scale.

2.2.2 Atom residual prior

If we regard the object prior as the *main* shape space, we can think of the difference between *main* shape space and the actual space as the *residual* shape space. The *residual* shape space contains high frequency signals of the shape details that are left from the coarse object scale. We again use PGA to analyze this *residual* shape space, but locally on each atom to match the "signal" with the "resolution".

An atom in m-rep lives in the curved space $\mathbf{G}$, as does an atom shape residual. Given two medial atoms $\mathbf{A}_i$, $\mathbf{A}'_i \in \mathbf{G}$, the residual $\Delta\mathbf{A}_i$, from $\mathbf{A}'_i$ to $\mathbf{A}_i$, is calculated by the subtraction operator:

$$\ominus : \mathbf{G} \times \mathbf{G} \to \mathbf{G}, \tag{1}$$
$$\begin{aligned}\Delta\mathbf{A}_i &\doteq \mathbf{A}_i \ominus \mathbf{A}'_i \\ &\doteq (\Delta\mathbf{p}_i, \Delta\mathbf{r}_i, \Delta\mathbf{u}_i, \Delta\mathbf{v}_i). \\ &\doteq (\mathbf{p}_i - \mathbf{p}'_i, (\mathbf{r}_i - \mathbf{r}'_i)/\mathbf{r}'_i, \mathbf{R}_{\mathbf{u}'_i}(\mathbf{u}_i), \mathbf{R}_{\mathbf{v}'_i}(\mathbf{v}_i)),\end{aligned}$$

where $\mathbf{A}_i = (\mathbf{p}_i, \mathbf{r}_i, \mathbf{u}_i, \mathbf{v}_i)$ with the hub position $\mathbf{p}_i \in \mathbf{R}^3$. $\mathbf{R}_\mathbf{w}$ represents the rotation along the geodesics in $\mathbb{S}^2$ that moves a point $\mathbf{w} \in \mathbb{S}^2$ to the North Pole (0, 0, 1) $\in \mathbb{S}^2$. The addition operator is correspondingly defined as

$$\oplus : \mathbf{G} \times \mathbf{G} \to \mathbf{G}, \tag{2}$$
$$\begin{aligned}\mathbf{A}_i &\doteq \mathbf{A}'_i \oplus \Delta\mathbf{A}_i \\ &\doteq (\mathbf{p}'_i + \Delta\mathbf{p}_i, \mathbf{r}'_i \cdot (\Delta\mathbf{r}_i + 1), \mathbf{R}^{-1}_{\mathbf{u}'_i}(\Delta\mathbf{u}_i), \mathbf{R}^{-1}_{\mathbf{v}'_i}(\Delta\mathbf{v}_i)).\end{aligned}$$

The shape residual can be calculated as

$$\Delta\mathbf{A}_i = \mathbf{A}_i^{train} \ominus \hat{\mathbf{A}_i^{obj}}, \tag{3}$$

where $\mathbf{A}_i^{train}$ is the training m-rep and $\mathbf{A}_i^{obj}$ is the projection of the training model onto the object prior. The dot product of the model deviation vector with the N eigenvectors produces the coefficients for the projection model:

$$\begin{aligned}x_j &=< \mathbf{A}_i^{train} \ominus \mu_i^{train}, v_j^{obj} > \\ \hat{\mathbf{A}}_i^{obj} &= \mathbf{A}_i^{(proj,obj)} \\ &= \mu_i^{train} \oplus \sum_{j=1}^{N} x_j v_j^{obj},\end{aligned} \tag{4}$$

where μ_i^{train} is the Fréchet mean of the training sample.

Based on the projections onto the object prior, the training samples for computing the atom prior therefore have to be the same with the object prior. We apply PGA to get the principal modes v_j^{atom} for each atom. The atom parametrization is 8 or 9 dimensional, depending on whether the atom is interior atom or exterior atom(Fig.1). The PGA results on a single atom often yields on average 4 modes that are bigger than zero. Similar to the object scale projection, we project the residual into its trained residual shape space. Together with the object scale projection, we calculate the multiscale projection as follows:

$$\begin{aligned}\mathbf{x}_j &=< \mathbf{A}_i^{train} \ominus \mathbf{A}_i^{(proj,obj)}, v_j^{atom} > \\ \hat{\mathbf{A}}_i^{MultiStage} &= \mathbf{A}_i^{(proj,obj+atom)} \\ &= \hat{\mathbf{A}_i^{(proj,obj)}} \oplus \sum_{j=1}^{K} x_j v_j^{atom}, K \leq 8.\end{aligned} \tag{5}$$

To date we assumed that the probabilistic distribution of local residual $\Delta\mathbf{A}_i$ on atom i is independent of its neighbor atoms. Namely, we assume that the large range correlation among the atoms are captured by the object scale PGA so that the residual shape variation space of each atom can be think of isolated with each other. In situations with more than two scales, the entities in the same scale are considered independent with each other.

2.3. Robustness measurement of the shape prior

According to Muller [9] the robustness of the probability distribution estimation can be measured by the fitness of a test population to the estimated probability distribution, over a variety of training sample sizes. The fitness is evaluated by the squared correlation ρ^2 on the sum squared distances between test cases and their projections in the space estimated from training cases, which reduces to the following formula:

$$\rho^2 = \frac{\sum_{i=1}^{N} d(\hat{m}_{i,test}, \overline{m}_{train})^2}{\sum_{i=1}^{N} d(m_{i,test}, \overline{m}_{train})^2}, \tag{6}$$

where $\hat{m}_{i,test}$ is the projection of the test model $m_{i,test}$ on the shape space and $\overline{m}_{train}$ is the Fréchet mean of the training sample. The distance is the Euclidian distance for Cartesian space, and the geodesic distance for the manifold in our m-rep case. The projected model indicates the closest model to the test model in the trained shape space, according to geodesic distances d in the feature space. The closer ρ^2 is to 1, the better predictability provided by the shape prior.

2.3.1 Toy experiment on deformed ellipsoids

We evaluate the robustness of the statistics first on a simulated data set of 1000 warped ellipsoids. M-rep models of warped ellipsoids were produced by successively applying independent globally bending, tapering, and twisting to a base ellipsoid, followed by a relatively small amount of local perturbation on the hub position of a fixed atom, which visually create a protrusion or indentation on that local region of the surface (Fig.2). Each transformation was sampled from a zero-mean Gaussian with its respective variance. The accumulated percentages of variation of object eigenmodes are demonstrated in Fig.2, which shows the distribution of the shape variation among the trained eigenmodes from 20 randomly selected m-rep models. We found that the first three eigenmodes obtained by PGA separately capture each of the three motions, namely bending, twisting and tapering. And the small amount of perturbation on that fixed region can be found in the 4th and 5th eigenmodes , but mixed with small variations on other random regions due to the limited number of training samples.

Shape priors using a) three object eigenmodes, followed by atom eigenmodes, and b) three object eigenmodes only

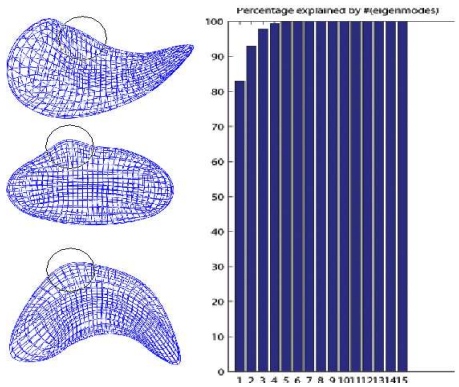

Figure 2. Left: 3 typical training m-reps that are generated by a random mix of three global variation and one local variation (circled areas). Right: the accumulative variation percentage plot for the principal geodesic eigenmodes on randomly selected 20 m-reps from 1000 training m-reps.

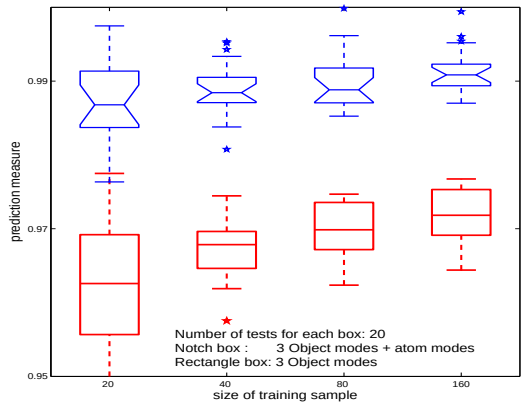

Figure 3. Prediction measurement against training sample sizes are shown for the shape prior with 3 object modes and the shape prior containing extra atom modes.

were compared. The projections introduced in section 2.2.2 were used to estimate the closest model in the training space. We can clearly see the improvement of the predictability of the shape space from Fig. 3. The test group used for each ρ^2 calculation has no overlapping over the corresponding training samples. For each training sample size, 20 tests were and the stars in the figures are the outliers.

From this toy experiment with carefully designed artificial variation, predictability of the two-scale shape prior is shown. It suggests that the atom scale residual modes are able to capture the fairly small amount of local region variation ignored by the object scale modes .

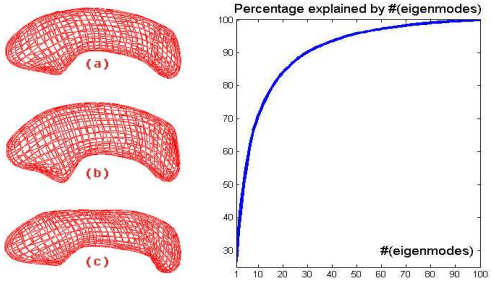

Figure 4. Left: (a) The Fréchet mean of the 290 m-reps. (b)and(c)are randomly selected from the 500 samplings out of the 51 modes. Right: The accumulated percentages of the variances explained by the eigenmodes trained from the 290 m-reps.

2.3.2 Experiments on simulated hippocampus models

To further investigate the advantages of the coarse-to-fine prior on real anatomical structures, we simulated a large group of hippocampus m-rep models from hand-segmented 290 CT images. First, the m-rep models well fitted to images were obtained by a semi-automated process. Those models were then used to form the sampling shape space by global PGA. As can be seen from the right of the Fig.4, 51, 33, 17 eigenmodes respectively capture 95%, 90% and 80% total variation of the 290 models. By a Gaussian random sampling on the 51 modes, 500 hippocampus models were generated. Examples of the simulated models are shown in Fig.4.

Shape priors with different number of eigenmodes were evaluated. We first compared the 33 object-modes prior with the prior composed of 10 object modes and atom modes, shown in Fig.5 (top). Note that for training sample size N, there are only maximum $(N-1)$ modes actually available. So at sample size of 20, 19 object modes are actually used. We also calculated the volume overlap measurement in Fig.5 (bottom) for the same comparison. The Dice Similarity Coefficents were computed on the volumes of the test models and the corresponding projected models. The two measurements consistently demonstrate the improvements of the robustness.

In general CT image segmentations, 33 dimensional space is already fairly large to search for the optimum shape, robustness is challenged by local minimums during the optimization. In this example, using only 10 object modes followed by atom modes shows much better predicability. More importantly, much better performance were seen in tests with small training sample size, which makes

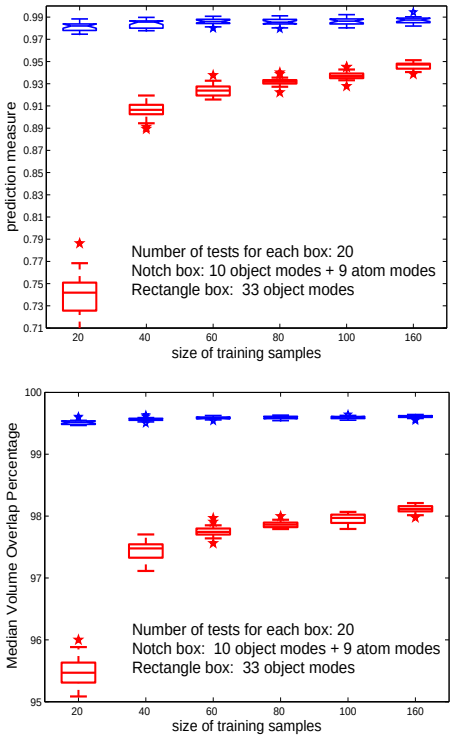

Figure 5. Comparisons on the robustness of the shape priors between using 33 object modes with using 10 object modes followed by atom modes. Top: ρ^2 measurement; Bottom: Volume Overlap between the test models and the projected models.

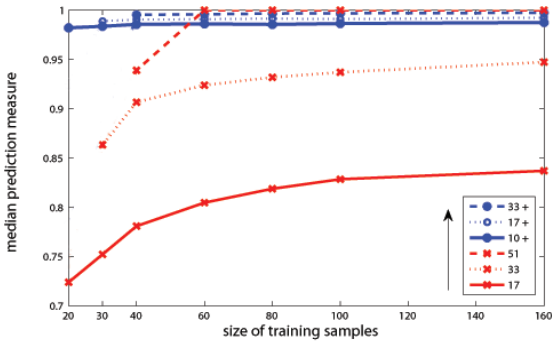

Figure 6. Six ρ^2 tests with different combinations of the object modes and atom modes: each marker shows the median ρ^2 number of 20 tests. Red curves are the results of using only object modes of 17, 33,51. Blue curves are the results of using 10, 17, 33 object modes followed by atom modes (tagged by"+").

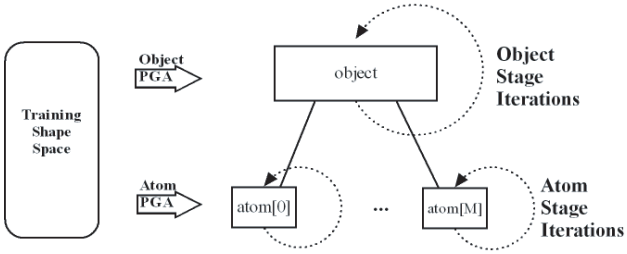

Figure 7. The probabilistic multiscale segmentation scheme for single object m-rep model: The object and atom stage shape statistics are collected from training samples and then used during the optimization process in the corresponding stages respectively.

this coarse-to-fine strategy superior in High Dimensional Low Sample Size problems.

More comparisons are shown in Fig.6. Six shape priors with different combination of modes were tested. To make the comparison easier, only the median ρ^2 number for each group of 20 tests are shown. The two-scale priors are much more robust than the single object priors overall. As the training sample size increases, each the curve converges in certain level. For the object priors (in red), predictability is getting better with a higher number of modes used. Two-scale priors (in blue) converge much faster as the training sample increases and show their strength with training sample sizes less than 60. No measurements are calcuated at training sample size 20 and 30 for the priors with bigger number of eigenmodes. The improvements from the object prior are most noticeable at smaller training sample sizes.

3. Coarse-to-fine Posterior Optimization

A coarse deformation (object stage) followed by a fine deformation (atom stage) can optimize the model hierarchically, as shown in Fig.7. In each stage, a Bayesian posterior optimization scheme is applied using the off-line trained statistics in the corresponding scale.

In the object stage the entire sheet of medial atoms deform together, restricted by the variation learned from the multiple sheets of atoms of the training samples. While in the atom stage, each atom are constrained in its own variation space trained from corresponding atoms of the training samples. The object scale segmentation results provides a good starting point for the local refinement stage, where the model deforms locally towards more "correct" configurations.

Besides the shape prior training, the image intensity patterns are also trained by PCA on Regional Intensity Quantile Functions (RIQFs) of local regions corresponding to the spoke ends of the atoms, which probabilistically represent the appearance of the local region in the image. The details for the image appearance model can be found in [1]. With the training statistics, at each of the two scales the model deforms by optimizing the sum of two penalty terms: the log geometric prior and the image match measurement computed by the log probability distribution on the RIQFs

of the implied boundary of the m-rep, since by Bayesian rule $\arg\max_{\mathbf{m}}(\log p(\mathbf{m}|\mathbf{I})) = \arg\max_{\mathbf{m}}[\log p(\mathbf{I}|\mathbf{m}) + \log p(\mathbf{m})]$, where $\mathbf{m}$ is the whole grid of medial atoms, and the I is the image intensity. The optimization is over the principal geodesic coefficients of the shape $\mathbf{m}$, thus restricting the result to the shape space spanned by the principal geodesic directions. The local scale posterior optimization is described by the formula:

$$\arg\max_{\Delta\mathbf{A}_i}(\log p((\mathbf{A}_i \oplus \Delta\mathbf{A}_i)|\mathbf{I}_i)$$
$$= \arg\max_{\Delta\mathbf{A}_i}[\log p(\mathbf{I}_i|\mathbf{A}_i \oplus \Delta\mathbf{A}_i) + \log p(\Delta\mathbf{A}_i)] \quad (7)$$

where $\Delta\mathbf{A}_i = \sum(x_j v_j^{atom})$ is the atom scale variation. The optimization is over the coefficient vector x for atom i. Starting from the object stage resulting model, we randomly loop over all the atoms and update each atom by adding the residual deformation that gives the best log posterior probability value.

4. CT Image Segmentation Results

We tested the methodology on a data set of CT images from 5 patients, and on-average 16 daily CT scans of the male pelvic area taken for each patient during the image-guided radiation therapy. The image has an in-plane resolution of 512×512 with pixel dimension of 0.98 mm × 0.98 mm and an inter-slice distance of 3 mm. The manual segmentation by an expert was provided.

For each patient, we first carried out a leave-one-day-out bladder segmentation. Training was done on all other days when segmenting the target day image. The results were compared with the manual segmentations and measured in terms of average closest surface point distance. Fig.8, depicts the segmentation results in box plot. The training models are the m-rep models fitted to the binary images from the human segmentation, which can be regarded as the desired segmentation results. The initialization are done by optimizing the model over the object shape space by evaluating three slices of contours only, which can be regarded as a part of the object stage. After initialization, the model is pretty much correctly located in the image.The object stage then further optimize into a coarse shape as a good starting point for the atom stage. The observed improvements in median average surface distance, on the order of 0.1 millimeter, were often the result of significant improvements, i.e. of several millimeters, localized to just a few regions of the model. The decreased interquartile range indicated improved robustness. The atom scale refinement can be seen across the population, which on average brings the accuracy into the subvoxel level.

We also found that the improvements in the bladder segmentation are most noticeable at parts of the boundary where the contrast is high (Fig.10). An example of

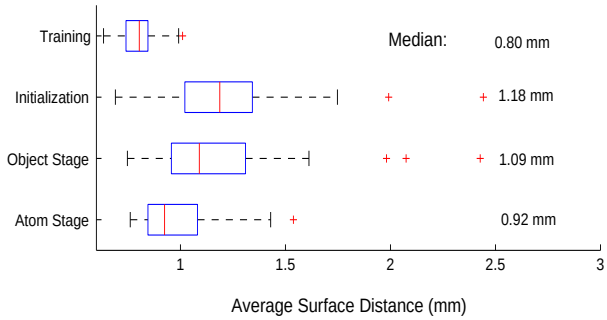

Figure 8. Leave-one-day-out bladder segmentation results for 79 images from 5 patients.

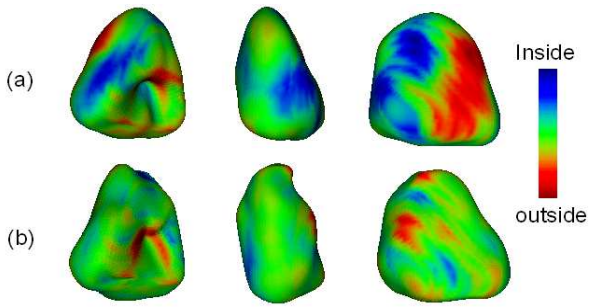

Figure 9. Surface distance errors shown in color from blue (5.5 mm inside) to red (5.8 mm outside) (a) Object stage result with average surface distance of 1.15 mm. (b)Atom stage result with average surface distance of 0.92 mm.

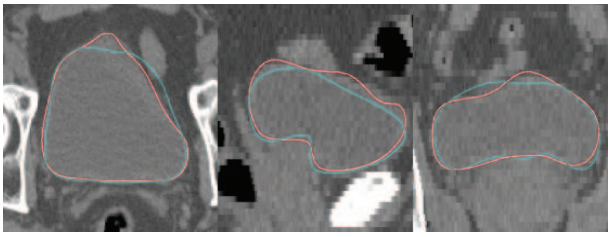

Figure 10. 2D contours comparison: object stage contours (in cyan) and the atom stage contours (in red) in 2D slices from axial (left), coronal (middle) and sagittal (right) views.

typical (median surface distance) bladder segmentation results from the leave-one-day-out study is shown to demonstrate the difference made by the atom stage visually on the boundaries. Fig.9 shows surface distance in color map from different views. And Fig.10 shows the three orthogonal views through the models with the corresponding slices of gray scale CT image shown as the background.The localized refinement are clearly shown in the three slices.

In addition, we applied the statistical atom stage to a leave-one-patient-out [14] study. Among the 5 patients, one was chosen as the target patient and the other 4 became the

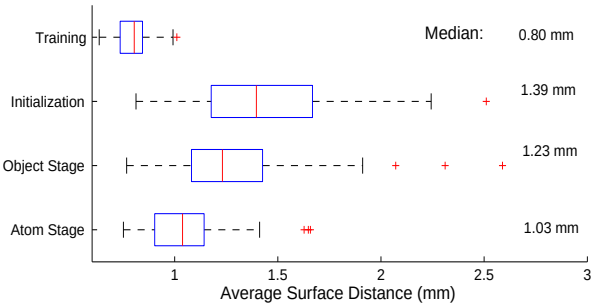

Figure 11. Leave-one-patient-out bladder segmentation results for 79 images from 5 patients.

training sample. We pooled the trained variation from different patients together after a pelvic bone based alignment [14]. During the segmentation time, we took the previous days' mean model as the initialization for the next day's image segmentation, with the first day's segmentation being provided beforehand. Compared to the leave-one-day-out study, although the training sample size is about 3 times larger, this experiment was more challenge in terms of the more varied shape space across different patients and the less strong correspondences for the local regions across patients. Nevertheless, the refinement results can be seen from Fig. 11, where we can see the encouraging improvement.

For CT image applications, we also applied the strategy to caudate segmentation and got satisfying improvement. We conclude that the local residual statistics effectively enable the refinement from a coarse shape. For several low contrast regions where even human have difficulties to locate the boundary, the atom stage tends to avoid further deformations.

5. Conclusion

In this paper, we presented a novel method to build up a coarse-to-fine shape prior by discovering the fine residual statistics from the coarse scale hierarchically via m-rep. We showed improved robustness on two large simulated data sets, evaluated in ρ^2 and volume overlap measurements against different training sample sizes. We applied the coarse-to-fine shape prior into a multiscale posterior optimization segmentation framework to improve the segmentation accuracy. Our experiment suggests that the residual shape statistics can be properly used to restrict results in the trained shape space and to penalize the local deformation towards better posterior estimates, combined with the image appearance likelihood. Segmentation experiments on 79 CT bladder images demonstrated improved segmentation accuracy by the local scale refinement.

A condition for building up the localized shape priors is the correspondence of the local regions cross training samples. Better correspondence condition generates the tighter statistics. Here in m-reps, atoms with the same index correspond to each other geometrically. In other presentations without natural correspondence, resampling techniques can be used to set the correspondence in different scales.

For High Dimensional Low Sample Size problems, the strategy of multiscale is shown effective to improve the estimations of the shape. Computationally, optimizing in a high dimensional space is much more likely to yield local optima and less efficient than the successive optimization over far fewer dimensions at each stage provided by the coarse-to-fine approach.

Besides m-reps, the usage of the local residual statistics is applicable to various representations with notions of localities. More studies need to be carried out on decisions on the number of scales and the size of each scale.

References

[1] R. E. Broadhurst, S. M. Stough, Joshua V.and Pizer, and E. Chaney. A statistical appearance model based on intensity quantiles. *IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 422–425, 2006. 5

[2] T. Cootes, D. H. Taylor, C. J.and Cooper, and J. Graham. Active shape models-their training and application. *CVGIP: Image Understand*, 61(1):38–59, 1994. 1

[3] M. Costa, H. Delingette, and N. Ayache. Automatic segmentation of the bladder using deformable models. *Proc. of IEEE Int.Symposium on Boimedical Imaging(ISBI)*, pages 904–907, 2007. 1

[4] P. Fletcher, C. Lu, S. Pizer, and S. Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *Medical Imaging, IEEE Transactions on*, 23:995–1005, 2004. 2

[5] D. Gibou, F nd levy. Partial differential equationscbased segmentation for radiotherapy treatment planning. *IMathematical biosciences and engineering*, 2(2):209–226, 2005. 1

[6] S. Joshi, S. M. Pizer, P. T. Fletcher, P. Yushkevich, A. Thall, and J. S. Marron. Multiscale deformable model segmentation and statistical shape analysis using medial descriptions. *IEEE TRANSACTIONS ON MEDICAL IMAGING*, 21(5), 2002. 1

[7] K. Leonard. Efficient shape modeling: -entropy, adaptive coding, and blum's medial axis versus the boundary curve. *International Journal of Computer Vision*, 72(2):183–199, 2007. 1

[8] T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: A survey. *Medical Image Analysis*, 1(2):91–108, 1996. 1

[9] K. E. Muller, S. Ray, and Y.-Y. Chi. Goodness of prediction for principal components, including high dimension, low sample size. *In prepertation*, 2008. 3

[10] D. Nain, S. Haker, A. Bobick, and A. Tannenbaum. Multiscale 3d shape analysis using spherical wavelets, 2005. 1

[11] C. Nikou, F. Heitz, J. Armspach, G. Bueno, and D. Vernon. A physically-based statistical deformable model for brain im-

age analysis. *Proceedings of the 6th European Conference on Computer Vision-Part II*, 1843:528–542, 2000. 1

[12] T. Pasquier, D.and Lacornerie, R. J. Vermandel, M., E. Lartigau, and N. Betrouni. Automatic segmentation of pelvic structures from magnetic resonance images for prostate cancer radiotherapy. *Int. J. Rad Onc. Biol. Phys.*, 68(2):592–600, 2007. 1

[13] V. Pekar, T. McNutt, and M. Kaus. Automated model-based organ delineation for radiotherapy planning in prostatic region. *Int. J. Rad Onc. Biol. Phys.*, 60(3):973–980, 2004. 1

[14] S. M. Pizer, R. E. Broadhurst, J.-Y. Jeong, Q. Han, R. Saboo, J. V. Stough, G. Tracton, and E. L. Chaney. Intra-patient anatomic statistical models for adaptive radiotherapy. *MICCAI Conference: Workshop on From Statistical Atlases to Personalized Models: Understanding Complex Diseases in Populations and Individuals*, pages 43–46, 2006. 6, 7

[15] S. M. Pizer, E. L. Chaney, and R. E. e. a. Broadhurst. Segmentation of kidneys and pelvic organs from ct by posterior optimization of m-reps. *techreport*, 2006. 1

[16] S. M. Pizer, P. T. Fletcher, and S. C. e. a. Joshi. Deformable m-reps for 3d medical image segmentation. *International Journal of Computer Vision*, 55(2-3):85–106, 2003. 1

[17] M. Rousson, A. Khamene, and M. e. a. Diallo. Constrained surface evolutions for prostate and bladder segmentation in ct images. *CVBIA*, pages 251–260, 2005. 1

[18] M. Styner and G. Gerig. Three-dimensional medial shape representation incorporating object variability. *Computer Vision and Pattern Recognition*, pages 651–656, 2001. 1