

Statistical Modeling of Shape and Appearance Using the Continuous Medial Representation

Paul A. Yushkevich, Hui Zhang, and James C. Gee

Department of Radiology, University of Pennsylvania

Abstract. We describe a novel approach to combining shape and appearance features in the statistical analysis of structures in medical images. The continuous medial representation is used to relate these two types of features meaningfully. The representation imposes a shape-based coordinate system on structure interiors, in a way that uses the boundary normal as one of the coordinate axes, while providing an onto and nearly one-to-one parametrization. This coordinate system is used to sample image intensities in the context of shape. The approach is illustrated by the principal components analysis of the shape and appearance of the hippocampus in T1-weighted MRI from a schizophrenia study.

1 Introduction

In medical image analysis, it is typical to describe anatomical structures in terms of either shape or appearance. Combining these two classes of descriptors is challenging because appearance features are local measurements of certain tissue properties (e.g., T1 relaxation) that are sampled on a lattice, while shape features are derived from geometric loci (e.g., boundaries) that are modeled by meshes or parametric manifolds. The importance of combining shape and appearance features in the statistical analysis of anatomical structures is underscored by the wide use of Active Shape and Active Appearance models (ASM/AAM) [4,5].

In this paper, we demonstrate how the continuous medial representation (*cm-rep*) [13] can be used to co-analyze shape and appearance. Specifically, we focus on the unique way in which *cm-reps* impose a common shape-based coordinate system on anatomical structures in a population. This coordinate system extends the boundary parametrization $\mathbf{y}(u^1, u^2)$ onto the entire volumetric region enclosed by the boundary, in a way that (1) is depth-based, i.e., it preserves coordinates u^1 and u^2 along vectors normal to the boundary, and (2) is onto and, except at a codimension 1 set of points, one-to-one. The first property is akin to ASMs, which use fixed-length intensity profiles normal to the boundary to associate shape features with appearance features. The second property is analogous to AAMs. Hence, our method combines two attractive, but mutually exclusive, properties of ASMs and AAMs. However, it does so at the cost of representational accuracy, as *cm-rep* models are restricted to a class of shapes with non-branching skeletons. In [13], we estimated the accuracy with which *cm-reps* can describe the hippocampus. In this paper, we illustrate how *cm-reps* can be used to study the variability in hippocampal shape and appearance.

2 Methods and Experimental Results

2.1 Modeling Anatomical Structures with *CM-Reps*

In the first approximation, the *cm-rep* approach is the continuous analog of the Pizer et al. *m-rep* method [11]. Anatomical structures are modeled by inverting the process of skeletonization: the skeleton of a structure is defined explicitly and the boundary is derived from the skeleton analytically. In our method, the skeleton is described parametrically as a combination of a *medial manifold* and a positive-valued *radial scalar field* given at each point on this manifold (the modeling of branching skeletons as a set of connected medial manifolds is the subject of ongoing research). The radial scalar field is in turn derived from a *radial conductance scalar field* by solving a Poisson PDE. This step is necessary to conform to the equality constraints imposed on the radial scalar field by the medial geometry. The medial manifold and the radial conductance scalar field are defined using basis functions whose coefficients can be varied in order to apply deformations to the model. Actual anatomical structures are represented by fitting a template to characteristic images in a Bayesian estimation framework. In the following paragraphs, we describe our approach in greater detail. However, for a complete treatment of the subject, we refer the reader to [13].

A *cm-rep* model is defined uniquely by two sequences of coefficients: a vector-valued sequence $\mathbf{w}_1, \dots, \mathbf{w}_N$ in \mathbb{R}^3 and a real-valued sequence $\omega_1, \dots, \omega_M$. These coefficients are used together with a sequence of orthogonal twice-differentiable basis functions $f_i(u^1, u^2)$ on a regular domain $\Omega \in \mathbb{R}^2$ to define the medial manifold \mathbf{x} and the radial conductance scalar field ρ as

$$\mathbf{x}(u^1, u^2) = \sum_{i=0}^N \mathbf{w}_i f_i(u^1, u^2); \quad \rho(u^1, u^2) = \sum_{i=0}^M \omega_i f_i(u^1, u^2). \quad (1)$$

Currently, we use the real components of the Fourier basis to define functions f_i , but we expect that in the future a wavelet basis will prove to be better suited for Bayesian estimation of the coefficients.

The radial scalar field R is derived from the medial manifold \mathbf{x} and the radial conductance field ρ by solving the following variant of the Poisson PDE:

$$\Delta_{\mathbf{x}} R^2 = \rho; \quad \|\text{grad}_{\mathbf{x}} R\| = 1 \text{ on } \partial\Omega, \quad (2)$$

where $\Delta_{\mathbf{x}}$ and $\text{grad}_{\mathbf{x}}$ denote, respectively, the Laplace-Beltrami operator and the Riemannian gradient on the manifold \mathbf{x} . These operators are intrinsic manifold generalizations of the Laplacian and gradient operators in \mathbb{R}^n . Despite the unusual non-linear boundary condition, the PDE carries many desirable properties, such as uniqueness (which we can prove formally), existence and stability (which are supported by empirical evidence), invariance under similarity transforms applied to \mathbf{x} , etc. We solve this PDE numerically using Newton's method in the Finite Differences framework.

The boundary surface \mathbf{y} associated with a *cm-rep* model is generated analytically from the medial manifold and the radial field. This surface is naturally partitioned into two halves, \mathbf{y}^+ and \mathbf{y}^- , one on each side of the medial manifold:




Fig. 1. The local geometry of a point on the medial manifold

$$\mathbf{y}^\pm = \mathbf{x} + R \mathbf{U}^\pm, \quad \text{where} \quad \mathbf{U}^\pm = -\text{grad}_{\mathbf{x}} R \pm \sqrt{1 - \|\text{grad}_{\mathbf{x}} R\|^2} \mathbf{N}_{\mathbf{x}}, \quad (3)$$

and $\mathbf{N}_{\mathbf{x}}$ is the unit normal to the medial manifold. This expression is essentially the inverse of Blum's Medial Axis Transform [1] in 3-D. It describes two points of tangency between the sphere of radius R centered at \mathbf{x} and the boundary surface. Unit vectors \mathbf{U}^\pm are orthogonal to the boundary surface at \mathbf{y}^\pm . Fig. 1 illustrates the medial-boundary relationship described by (3).

It is easy to verify from the boundary condition in (2) that boundary halves \mathbf{y}^+ and \mathbf{y}^- coincide along $\partial\Omega$. This coincidence is a form of an equality constraint imposed by the medial geometry on functions \mathbf{x} and R . If, instead of deriving R using the PDE (2), we had modeled R explicitly as a weighted sum of basis functions, we would be presented with a severely overconstrained problem, as the number of places where the constraint holds would be infinite (all of $\partial\Omega$), while the number of coefficients defining \mathbf{x} and R would be finite. In addition to this equality constraint, there are certain inequality constraints that the coefficients \mathbf{w}_i and ω_i must satisfy in order to ensure that \mathbf{y}^+ and \mathbf{y}^- form a smooth closed surface. These are handled in the course of Bayesian estimation.

The three stages of *cm-rep* construction are illustrated in Fig. 2: the first panel shows the medial manifold and the radial conductance field, the second panel plots the radial field, and the third panel shows the boundary surface. Note that the medial manifold has corners; this is an undesirable side effect of using the unit square as the domain Ω . We are working to extend our method to arbitrary domains in \mathbb{R}^2 .




Fig. 2. The three steps of constructing a *cm-rep*. **a.** A medial manifold \mathbf{x} with the radial conductance function ρ . **b.** The radial function R computed by solving the Poisson equation (2) on the manifold. **c.** Boundary surface \mathbf{y} computed using (3).

The basic ideas behind *cm-rep* deformable modeling are derived from the discrete m-rep methodology [11], which is based on pattern theory [10]. First, a template *cm-rep* is generated. Currently that is done in an *ad hoc* manner, but in the future we plan to study optimal template selection, perhaps adapting techniques developed for m-reps [12]. The deformable template is fitted to instances of the anatomical structure by minimizing the posterior probability of the coefficients $\{\mathbf{w}_i\}$ and $\{\omega_i\}$, given data in the form of a binary characteristic image of the structure. In a classical Bayesian estimation framework, the posterior probability is factored into a likelihood term, which measures the match between the *cm-rep* boundary/interior and the binary image, and a prior term. We use volumetric overlap (which can be computed efficiently using the *cm-rep* interior parametrization, Sec. 2.2) and boundary-based match metrics to compute the likelihood. Penalty prior terms are used to enforce the inequality constraints on $\{\mathbf{w}_i\}$ and $\{\omega_i\}$ ‘softly’. In addition, a regularization prior is used to minimize the distortion in area element on the medial manifold, thus providing a rudimentary correspondence between instances.

The deforming *cm-rep* template can only assume shapes for which the skeleton is a single manifold. This clearly limits the representational ability of the model. However, in practice, it appears that many anatomical structures can be modeled with *cm-reps* fairly accurately. Indeed, Styner [12] proved that certain subcortical structures can be represented using single-figure discrete m-reps with sub-voxel error. To evaluate the representational ability of *cm-reps* in a similar manner, we fitted the hippocampus template to 174 (87 right, 87 left) segmentations of the hippocampus from a MRI schizophrenia study [3]. The segmentation was computed using the Joshi et al. [9] algorithm for large deformation diffeomorphic registration with manually placed anatomic landmarks. This approach is used extensively in brain morphometry [6] and was shown to be more accurate and reliable than manual segmentation [8]. The data in the form of boundary meshes was graciously provided by Profs. Guido Gerig (UNC Depts. of Comp. Sci. and Psychiatry) and Sarang Joshi (UNC Dept. of Rad. Onc.).

The results of the fitting are illustrated in Fig. 3. The fit was computed in a multi-resolution procedure where the number of basis functions and coefficients was gradually increased. At the highest resolution, $8 \times 12 \times (3 + 1)$ coefficients were used. After the fitting, we computed the average of several goodness-of-fit scores over the entire data set. These include mean squared distance from the



Fig. 3. Examples of a template fitted to instances of the hippocampus. The solid blue surface is the boundary of the subject hippocampus, and the white mesh is the boundary of the fitted *cm-rep* template.

cm-rep boundary to the hippocampus boundary (0.201 mm), maximum distance from *cm-rep* to the hippocampus (1.576 mm), maximum distance from the hippocampus to the *cm-rep* (2.558 mm) and volume overlap between the *cm-rep* and the hippocampus (91.3%). These error scores are slightly worse than the *m-rep* results (avg. model-to-hippo. distance of 0.17 mm.) [12] that were computed on the same images, but using a different manual segmentation. The *m-rep* model does not enforce strict conformance to medial geometry, so we expect it to fit somewhat better than *cm-reps*.

2.2 CM-Rep Interior Parametrization

One advantage of *cm-reps* over discrete *m-reps* and boundary shape shape representations is the ability to impose, with ease, a shape-based coordinate system on the *cm-rep* interior, i.e. the region of space enclosed by the *cm-rep* boundary. For ‘valid’ *cm-reps*, the interior is homeomorphic to a unit ball. The vectors $\mathbf{U}^\pm(u^1, u^2)$ with tails at $\mathbf{x}(u^1, u^2)$ span the *cm-rep* interior. We can use these properties to define a shape-based coordinate system that consists of the coordinates $(u^1, u^2) \in \Omega$ and a scalar $\xi \in [-1, 1]$ that describes a point’s location with respect to the medial axis and the boundary. Formally, we define a mapping from $\Omega \times [-1, 1]$ onto the *cm-rep* interior:

$$\mathbf{z}(u^1, u^2, \xi) = \mathbf{x}(u^1, u^2) + |\xi|R(u^1, u^2) \mathbf{U}^{\text{sign}\xi}(u^1, u^2) \quad (4)$$

Under this parametrization, points on the medial manifold have $\xi = 0$ and points on the boundary have $\xi = \pm 1$. If Q is some point on the boundary, then all points along the inward boundary normal vector with tail at Q have the same first two coordinate values as Q , due to the fact that the vectors \mathbf{U}^\pm are orthogonal to the *cm-rep* boundary. The distance from a point on the *cm-rep* interior to the nearest point on the *cm-rep* boundary is equal to $(1 - |\xi|)R(u^1, u^2)$.

Thus, we have parameterized the entire object interior, in a way that, roughly speaking, associates each interior point with the nearest point on the boundary. This type of interior parametrization is consistent with the way that the ASM [4] samples image intensities inside objects using profiles that extend in the normal direction from the boundary. Unlike ASM, but like AAM [5], our parametrization is onto. It is also one-to-one, with the exception of a codimension 1 set of points whose coordinates $(u^1, u^2) \in \partial\Omega$. For these points, $z(u^1, u^2, \xi) = z(u^1, u^2, -\xi)$. Thus, we may say that our coordinate system combines the best of ASM and AAM: the preservation of boundary normal direction and (nearly) one-to-one and onto parametrization.

When *cm-reps* are fitted to anatomical objects in medical images, we are able to use this coordinate system to sample image intensities. This ability is illustrated in Figs. 4 and 5a. Here, a *cm-rep* has been fitted to the hippocampus in a T1 image (Fig. 4a.). The values of u^1, u^2 and ξ coordinates of on the *cm-rep* interior are shown using color maps in Figs. 4b-d. Fig. 5a shows the image intensities mapped back to the u^1, u^2 and ξ space. This mapping of image intensities associates shape features with intensity features, and is a unique way of establishing across-subject correspondences between intensities on the basis of geometry.




Fig. 4. The shape-based coordinate system induced by a *cm-rep*. **a.** A slice through the hippocampus in a T1 weighted MRI, to which a *cm-rep* template has been fitted. **b, c.** The values of the coordinates u^1, u^2 , which span the medial manifold, plotted at each point in the hippocampus using a color map. **d.** The values of the ξ coordinate, which goes from the medial manifold to the boundary.




Fig. 5. a. MRI intensities of the hippocampus in Fig 4a, sampled on a lattice in the shape-based coordinate system. The slices are taken as ξ goes from -1 to 1, and the axes are along the u^1 and u^2 coordinates. Sampling used cubic interpolation. **b.** The PCA mean of the hippocampal image intensities in the shape-based coordinate system. The CSF adjacent to the hippocampus can be seen near the edges; this illustrates the error in the initial segmentation combined with the error in *cm-rep* fitting.

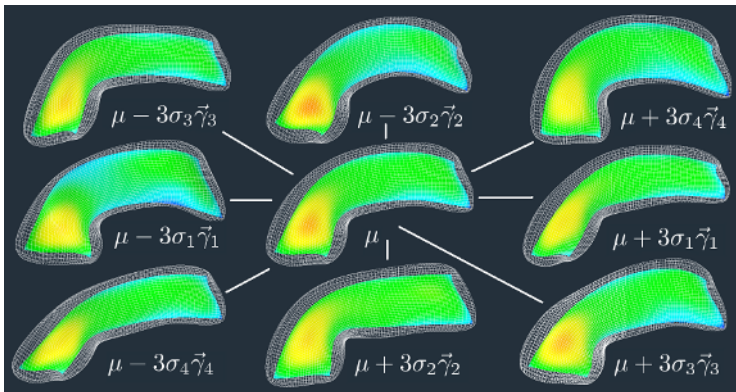


Fig. 6. Four principal modes of variability in the left hippocampus shape. The eigenvalues corresponding to these modes are 2.57, 2.28, 1.81 and 0.96, the total spectrum is 11.38, so the modes shown here represent 66.8% of total variability. The color map represents the radius function.

2.3 Statistical Modeling

To show how *cm-reps* would be used for shape characterization, we performed principal component analysis (PCA) on the shape and appearance features of *cm-reps* fitted to the hippocampus data. The shape features were computed by taking the values of the coefficients $\{\mathbf{w}_i\}, \{\omega_i\}$ after aligning the *cm-reps* using the Generalized Procrustes method [7]. The appearance features were sampled on a regular lattice in the u^1, u^2, ξ space. Fig. 6 shows the mean shape and principal modes of shape variability, and Fig. 5b shows the mean intensity pattern. Our PCA did not take into account the fact that some linear combinations of the *cm-rep* parameters that PCA generates can violate one of the inequality constraints mentioned above. However, no invalid *cm-reps* were generated by staying within 2.8 standard deviations from the mean in the first 10 principal modes. This indicates that valid and invalid *cm-reps* are well separated in feature space.

3 Discussion and Conclusions

We have presented *cm-reps*: a shape representation that models the continuous geometric relationship between the boundaries and skeletons of objects. The ability to impose a coordinate system on the interior of structures in a way that preserves two of the three coordinates along boundary normals and is nearly one-to-one and onto is the strength of the representation. The lossy nature of the representation, which is a potential weakness, was evaluated for the hippocampus, and the average representational error was found to be relatively small. The utility of *cm-reps* for the statistical analysis of shape and appearance was demonstrated by applying PCA to the *cm-reps* of the hippocampus. In the future, we plan to address the problem of *cm-rep* correspondence more directly and to use the *cm-rep* PCA as a component in an algorithm for the segmentation the hippocampus in structural MRI. We also intend to use *cm-reps* for structure-oriented fMRI analysis.

Other approaches to depth-based parametrization of interiors of anatomical structures have appeared in the recent literature. These methods typically employ distance transforms or skeletonization algorithms to assign a depth-based coordinate system to objects. For instance, Bouix et al. [2] find the dominant medial surface in the hippocampus and flatten it to form a reference space. The advantage of our method is that it is model-based, so a consistent shape-based coordinate system is given by construction, while methods such as [2] require pruning of skeletal branches and registration to find a common coordinate frame.

Acknowledgement

This work was supported by the NIH grant NS045839. We thank Prof. Charles L. Epstein, Prof. Jean Gallier and Marcelo Siqueira (Penn) for insightful discussions. We are indebted to Profs. Guido Gerig, Stephen M. Pizer, Sarang Joshi, and Martin Styner (UNC) for providing data and inspiration for this work.

References

1. H. Blum and R.N. Nagel. Shape description using weighted symmetric axis features. *Pattern Recognition*, 10(3):167–180, 1978.
2. Sylvain Bouix, Jens C Pruessner, D. Louis Collins, and Kaleem Siddiqi. Hippocampal shape analysis using medial surfaces. *Neuroimage*, 25(4):1077–89, May 2005.
3. M. H. Chakos, S. A. Schobel, H. Gu, G. Gerig, D. Bradford, C. Charles, and J. A. Lieberman. Duration of illness and treatment effects on hippocampal volume in male patients with schizophrenia. *Br J Psychiatry*, 186(1):26–31, 2005.
4. T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models – their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.
5. T.F. Cootes, G.J. Edwards, and C.J. Taylor. Active appearance models. In *European Conference on Computer Vision*, volume 2, pages 484–498, Freiburg, Germany, June 1998.
6. J. Csernansky, S. Joshi, L. Wang, J. Haller, M. Gado, J. Miller, U. Grenander, and M. Miller. Hippocampal morphometry in schizophrenia via high dimensional brain mapping. In *Proc. National Academy of Sciences*, volume 95, pages 11406–11411, 1998.
7. J.C. Gower. Generalized procrustes analysis. *Psychometrika*, 40:33–51, 1975.
8. J.W. Haller, A. Banerjee, G.E. Christensen, M. Gado, S. Joshi, M.I. Miller, Y.I. Sheline, M.W. Vannier, and J.G. Csernansky. Three-dimensional hippocampal MR morphometry by high-dimensional transformation of a neuroanatomic atlas. *Radiology*, 202:504–510, 1997.
9. S. Joshi, U. Grenander, and M. Miller. On the geometry and shape of brain sub-manifolds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:1317–1343, 1997.
10. D. Mumford. Pattern theory: a unifying perspective. In D. C. Knill and W. Richards, editors, *Perception as Bayesian Inference*, pages 25–62. Cambridge University Press, 1996.
11. S. M. Pizer, P. T. Fletcher, S. Joshi, A. Thall, J. Z. Chen, Y. Fridman, D. S. Fritsch, A. G. Gash, J. M. Glotzer, M. R. Jiroutek, C. Lu, K. E. Muller, G. Tracton, P. Yushkevich, and E. L. Chaney. Deformable m-reps for 3D medical image segmentation. *International Journal of Computer Vision*, 55(2):85–106, Nov 2003.
12. M. Styner, G. Gerig, S. Joshi, and S.M. Pizer. Automatic and robust computation of 3D medial models incorporating object variability. *International Journal of Computer Vision*, 55(2):107–122, November 2003.
13. P.A. Yushkevich, H. Zhang, and J.C. Gee. Parametric medial shape representation in 3-D via the Poisson partial differential equation with non-linear boundary conditions. In G.E. Christensen and M. Sonka, editors, *Information Processing in Medical Imaging*, pages 162–173, 2005.