

An Object-Based Volumetric Deformable Atlas for the Improved Localization of Neuroanatomy in MR Images

Tim McInerney^{1,2} and Ron Kikinis³

¹University of Toronto, Toronto, ON, Canada M5S 3H5

²Massachusetts Institute of Technology, Cambridge, MA 02139, USA

³Harvard Medical School, Boston MA 02115, USA

Abstract. We present a hierarchical object-based deformable atlas, a promising new approach for the automatic localization and quantitative analysis of neuroanatomy in MR images. The 3D finite element-based elastic atlas combines the advantages of both volumetric- and surface-based deformable atlases in one single unifying framework. This multi-resolution framework is not only capable of deforming entire volumes or subvolumes but can deform individual atlas objects, allowing greater and more effective use of object shape and local image feature information. Object surface representations are embedded in the volumetric deformable atlas and image-feature-derived forces acting on these surfaces are automatically transferred to the containing 3D finite element lattice. Consequently, spatial relationship constraints of the atlas objects are maintained via the elastic lattice while an object is deformed to match a target boundary. Atlas objects are deformed in a hierarchical fashion, beginning with objects exhibiting well-defined image features in the target scan and proceeding to objects with slightly less well-defined features. Experiments involving several subcortical atlas objects are presented.

1 Introduction

The automatic localization of neuroanatomy in MR images and the subsequent quantitative analysis using 3D elastically deformable atlases is gaining increased attention in medical imaging research [10,11,6,3,12,5,8,13,7]. These model-based techniques can dramatically decrease the time required for the localization task over interactive methods as well as improve the objectivity, reproducibility, and, potentially, the accuracy of the localization. A fitted anatomical atlas can then be used as a fundamental component for the assessment of structural brain abnormalities, for mapping functional activation of the brain onto the corresponding anatomy, and for computer-assisted neurosurgery.

There are essentially two approaches to deformable atlas matching: volumetric-based and surface-based. While both approaches offer a powerful strategy for efficient localization and analysis, they suffer from several well-known deficiencies affecting the accuracy of the localization. Volumetric approaches maintain the spatial relationships of the atlas objects implicitly via the elastic medium in

which they are embedded. However, these methods are sensitive to their initial placement – if the initial rigid alignment is off by too much, parts of the atlas may incorrectly warp onto the boundaries of neighboring features. This problem is exacerbated by the fact that volumetric methods discard the shape information of each atlas object and use only local intensity variations between the atlas and the target scan to drive the matching process.

Surface-based methods typically manually initialize several parametrically defined deformable surfaces – each representing a different neuroanatomical structure – and subsequently elastically deform the surfaces to extract the shapes of the target object boundaries. Each atlas object surface and its corresponding deformed surface in the target scan are then matched to produce surface warping functions. Using this surface warping information, a volumetric warp can be calculated via interpolation to deform the atlas material between the surfaces and register it with the target scan. One problem with this approach is that the surfaces are warped independently – the spatial relationship constraints provided by the elastic medium in the volumetric approach are initially discarded. Furthermore, if generic deformable surfaces are used – that do not make use of the known atlas object shape – then pieces of the object in the target scan could be missed by the deformable surface due to poor resolution or noise.

In this paper, we describe a new hierarchical object-based deformable atlas that combines the advantages of both the volumetric and surface based models in one unifying framework. The multi-resolution framework is not only capable of deforming entire volumes or subvolumes but can deform individual atlas objects and make use of image feature knowledge of an object. By embedding smooth surface representations of atlas objects into the finite element-based volumetric deformable atlas, image-feature-derived forces acting on these surfaces are automatically transferred to the surrounding volumetric finite elements. The subsequent deformation of the elements automatically deforms the embedded surfaces. Consequently, we maintain the spatial relationship constraints of the atlas objects via the elastic “medium” while also making use of shape and intensity information of each object. The atlas is deformed in a hierarchical fashion, beginning with an initial rigid alignment and elastic match over the entire volume. We then warp individual atlas objects and the surrounding volume in a specified neighborhood, starting with objects exhibiting well-defined image features in the target scan (such as the lateral ventricles), and then proceeding to objects with slightly less well-defined image features, and so on.

The motivation behind our approach is that while the standard volumetric deformable atlas approach can provide an automatic, efficient and good “overall” match of the atlas, there are still mismatches in individual atlas objects that can only be corrected by adjusting the object itself. At the same time, this “fine-tuning” should maintain and use the correct spatial relationships of neighboring objects. In the remainder of the paper, we will first describe an initial implementation of our model and then present some preliminary results using several subcortical structures to demonstrate the potential of this promising new approach.

2 Deformable Atlases

The idea behind a deformable anatomical atlas is to take the information contained in the atlas (typically a set of labeled voxels where the labels correspond to anatomical structures) and transfer this information onto the target dataset via a nonlinear warping function. To perform the warp, the atlas is typically modeled as a physical object and is given elastic properties. After an initial global alignment, the atlas deforms and matches itself onto corresponding regions in the brain image volume in response to forces derived from image features such as voxel similarity or image edges. To maintain the atlas topology and connectivity, the elastic properties of the atlas give rise to internal forces which regularize the deformation. The assumption underlying this approach is that at some representational level, normal brains have the same topological structure and differ only in shape details.

The idea of modeling the atlas as an elastic object was originated by Broit [4]. Bajcsy and Kovacic [1] subsequently implemented a multiresolution version of Broit's system where the deformation of the atlas proceeds step-by-step in a coarse to fine strategy, improving the robustness of the technique. The elastically deformable volume approach has become a very active area of research [10,11,6,3,7], and has recently been extended to a viscous fluid deformable volume [6] in an attempt to overcome the small deformation assumptions inherent in the linear elastic formulations. Although surface based deformable models have been widely used to segment medical images, surface-based deformable brain atlases are a more recent development [12,5,8,13].

3 Hierarchical Object-Based Deformable Atlas

Our model combines the advantages of the volumetric and surface based methods by integrating both approaches into one framework. This hybrid model embeds the surface of atlas objects into the solid deformable finite element mesh or lattice representing the elastic atlas. The result is that the deformation of atlas objects can be controlled individually or in combination while automatically maintaining the spatial relationships of each object via the elastic lattice. The model also has other distinct advantages:

- Multiple object-based image features can be used to attract an atlas object towards its boundary in the image. For example, known image intensity statistics of individual objects and *neighboring objects* can be used to weight a pressure force, driving the model towards salient boundary edges.
- Computing forces on the object surface and distributing these forces to the volumetric finite elements results in an accurate deformation of the object without discarding the contextual knowledge of the neighboring objects. The displacements of the deformed object are automatically passed on to the neighbors through the surrounding elastic lattice, resulting in improved localization of these objects. In addition, several neighboring objects can be

deformed concurrently in which case each object “competes” for the ownership of image features.

- The atlas is initially deformed over its entire volume to provide a good initialization of each atlas object. Atlas objects can be then deformed in a hierarchical fashion - from objects exhibiting very well-defined image features to objects with weaker features. After an object has been deformed we compute forces on it to maintain its equilibrium position. This acts to constrain the deformation of neighboring objects with adjacent boundaries (and less well-defined image features).
- The material properties of the 3D finite elements can be controlled on an object basis. Each object can use predefined elastic properties based on knowledge of the amount of deformation that typically will take place. If very large deformations are required an object can be deformed in phases, using coarse and rigid finite element meshes in initial phases and using finer, more flexible meshes in subsequent phases.

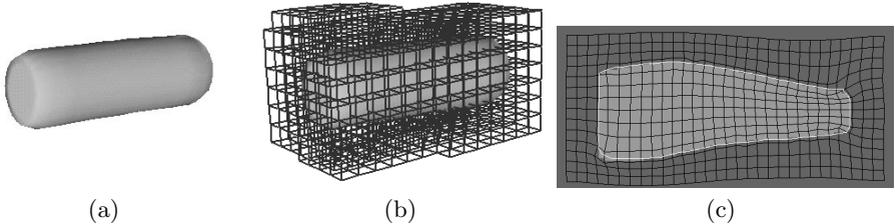


Fig. 1. (a) Synthetic object surface, (b) surface embedded in elastic 3D lattice, (c) cross-sectional view of deformed lattice and surface (surface: white, lattice: black, target data: gray).

3.1 Model Structure

We construct our object-based elastic atlas using a deformable 3D lattice or grid consisting of cubical finite elements. Each finite element corresponds to a labeled atlas voxel (or group of voxels). To deform a specified atlas object, we reconstruct a smooth triangulated surface of the object from the atlas using a modified marching cubes algorithm (Figure 1(a)). For each surface vertex, we compute its containing cubical element (Figure 1(b)) and the relative position of the vertex within the element. Image feature forces are calculated for each vertex and are then distributed to the nodes of the containing finite elements using the element basis (interpolation) functions. The elastic elements are then deformed and the nodal displacements calculated (Figure 1(c)). The new positions of the surface vertices are then computed based on their relative positions within their containing cubical element and the displacement of the element nodes. This process is repeated for a specified number of iterations.

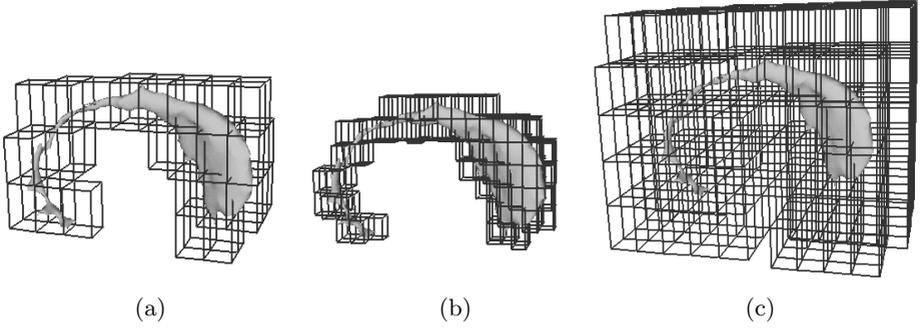


Fig. 2. Surface of caudate nucleus embedded in finite element mesh, (a) element size: 8 voxels, element layers: 1, (b) es: 4, el: 1, (c) es: 8, el: 2.

The lattice of cubical elements surrounding an atlas object can be controlled in a number of ways (Figure 2) by varying the element size, the number of element layers surrounding the object, and the element material properties.

3.2 Finite Element Implementation

The elastically deformable lattice model is implemented using the displacement-based finite element method. A brief description of this method for linear elastic bodies is provided below. Readers are referred to [2] for complete details.

In the displacement-based finite element method, a three dimensional body is located in a fixed coordinate system $\mathbf{X} = [X, Y, Z]^T$. The body is subjected to externally applied forces:

$$\mathbf{f}^B = [f_x^B \ f_y^B \ f_z^B]^T, \quad \mathbf{f}^S = [f_x^S \ f_y^S \ f_z^S]^T, \quad \mathbf{F}^i = [F_{x_i} \ F_{y_i} \ F_{z_i}]^T \quad (1)$$

where \mathbf{f}^B are the body forces (force per unit volume), \mathbf{f}^S are the surface traction forces (force per unit area), and \mathbf{F}^i are the concentrated forces (i denotes the point of force application) due to attachments. The displacements of the body from the unloaded configuration are measured in the coordinate system \mathbf{X} and are denoted by $\mathbf{U}(X, Y, Z) = [UVW]^T$. The strains corresponding to \mathbf{U} are:

$$\boldsymbol{\epsilon}^\top = \left[\frac{\partial U}{\partial X}, \frac{\partial V}{\partial Y}, \frac{\partial W}{\partial Z}, \frac{\partial U}{\partial Y} + \frac{\partial V}{\partial X}, \frac{\partial V}{\partial Z} + \frac{\partial W}{\partial Y}, \frac{\partial W}{\partial X} + \frac{\partial U}{\partial Z} \right]^\top, \quad (2)$$

the corresponding stresses are denoted $\boldsymbol{\tau}^\top = [\tau_{XX} \ \tau_{YY} \ \tau_{ZZ} \ \tau_{XY} \ \tau_{YZ} \ \tau_{ZX}]$, and finally the stress-strain relationship for a linear elastic body is given by $\boldsymbol{\tau} = \mathbf{D}\boldsymbol{\epsilon}$, where \mathbf{D} is the stress-strain material matrix (the elastic coefficients).

Assuming a linear elastic continuum with zero initial stresses, the total potential energy of the body can be written as

$$\Pi(\mathbf{U}) = \frac{1}{2} \iiint_V \boldsymbol{\epsilon}^\top \mathbf{D} \boldsymbol{\epsilon} dV - \iiint_V \mathbf{U}^\top \mathbf{f}^B dV - \iint_S \mathbf{U}^{S^\top} \mathbf{f}^S dS - \sum_i \mathbf{U}^{i^\top} \mathbf{F}^i. \quad (3)$$

Assuming, furthermore, a continuous displacement field \mathbf{U} that satisfies the displacement boundary conditions, the solution \mathbf{U} is at the minimum of this energy and is found at the vanishing of the first variation of Π (i.e. $\delta\Pi = 0$).

We approximate the body as an assemblage of discrete finite elements interconnected at nodal points on the element boundaries. The displacements, measured in a local coordinate system x, y, z within each element, are assumed to be a function of the displacements at the N finite element nodal points: $\mathbf{u}^j(x, y, z) = \mathbf{H}^j(x, y, z)\hat{\mathbf{U}}$, where \mathbf{H}^j is the displacement interpolation matrix of element j , and $\hat{\mathbf{U}}$ is a vector of the three global displacement components U, V, W at all nodal points. Currently we use an 8-node hexahedral element in our implementation (higher-order elements can also be used) and the element interpolation functions are specified in [2].

One typically calculates all element matrices using only element nodal displacements and the corresponding nonzero components of \mathbf{H}^j for element j so that: $\mathbf{u}^j = \mathbf{N}^j\hat{\mathbf{u}}^j$, where $\hat{\mathbf{u}}^j$ and \mathbf{N}^j are the element nodal displacements and basis functions, respectively. Using equations (2) and the stress-strain relationship, we can then evaluate the element strains as, $\boldsymbol{\epsilon}^j = \mathbf{B}^j\hat{\mathbf{u}}^j$, where \mathbf{B}^j is obtained by differentiating the components of \mathbf{N}^j . Using these two equations, the potential energy (3) can be rewritten in terms of its elemental contributions and subsequently minimized on an element-by-element basis:

$$\begin{aligned} \Pi = \sum_j \Pi^j(\mathbf{u}) = \sum_j \left\{ \frac{1}{2} \int_{V^j} \hat{\mathbf{u}}^{j\top} \mathbf{B}^{j\top} \mathbf{D}^j \mathbf{B}^j \hat{\mathbf{u}}^j dV \right. \\ \left. - \int_{V^j} \hat{\mathbf{u}}^{j\top} \mathbf{N}^j \mathbf{f}^{B(j)} dV - \int_{S^j} \hat{\mathbf{u}}^{j\top} \mathbf{N}^j \mathbf{f}^{S(j)} dS - \sum_i \hat{\mathbf{u}}_i^{j\top} \mathbf{N}^j \mathbf{F}_i^j \right\}. \end{aligned} \tag{4}$$

Solving for the first variation of Π leads to the equilibrium equation $\frac{\partial \Pi^j}{\partial \hat{\mathbf{u}}^j} = \mathbf{K}^j\hat{\mathbf{u}}^j - \mathbf{f}^j$, where \mathbf{K}^j and \mathbf{f}^j are the element stiffness and load matrices, respectively. We introduce a simple velocity-proportional damping force and rewrite the equilibrium equations as $\mathbf{C}^j \frac{d\hat{\mathbf{u}}^j}{dt} + \mathbf{K}^j\hat{\mathbf{u}}^j = \mathbf{f}^j$, where \mathbf{C}^j is a diagonalized damping matrix with velocity damping coefficients γ along the main diagonal. We currently integrate this equation forward through time using an explicit first-order Euler method on an element-by-element basis, making the model fitting process efficient and easily parallelizable.

An isotropic linear elastic material is characterized by the Lamé constants, λ and μ . These constants are also related to Young’s modulus of elasticity E and Poisson’s ratio ν :

$$E = \frac{\mu(2\mu + 3\lambda)}{\mu + \lambda}, \quad \nu = \frac{\lambda}{2(\mu + \lambda)}, \tag{5}$$

where E relates tension of the object and its stretch in the longitudinal direction and ν is the ratio of lateral contraction to longitudinal stretch. We typically set λ to zero and allow E to range from 0.25 to 0.75, producing a range of stable elastic behavior from relatively stretchy to relatively rigid.

3.3 Applied Forces

As mentioned earlier, forces are computed for each of the surface vertices and are then distributed to the 8 nodes of the containing element using the element shape functions \mathbf{N}^j . We use a weighted pressure force, where the weights are derived from precomputed atlas object intensity statistics, to deform the object surface:

$$F(I(\mathbf{x}_i)) = +1, |I(\mathbf{x}_i) - \mu| \leq k\sigma, \quad F(I(\mathbf{x}_i)) = -1, |I(\mathbf{x}_i) - \mu| > k\sigma, \quad (6)$$

where μ is the mean image intensity of the target object, σ the standard deviation of the object intensity and k is a user defined constant.

We also use a functional F based on intensity gradients computed along a surface vertex normal. It is often the case that the intensity of an object varies considerably over its extent, limiting the usefulness of functionals based on absolute image statistics. However, the intensity gradients between an object and its neighbors is often fairly consistent when computed over a large enough surface region surrounding a surface vertex:

$$F(I(\mathbf{x}_i)) = -1, |\nabla I(\mathbf{x}_i) \cdot \mathbf{n}_i| > C, \quad F(I(\mathbf{x}_i)) = +1, |\nabla I(\mathbf{x}_i) \cdot \mathbf{n}_i| < C, \quad (7)$$

where C is the known average difference in intensity between two objects. The signs are reversed in the functional if C is negative.

4 Experiments

The deformable atlas is based on an MR brain atlas developed in our laboratory [9]. To match the atlas to the target MRI scan, we first apply a rigid registration to the atlas followed by a generic volumetric elastic match [7] to initially deform the atlas and provide good initial positions of atlas objects ¹. We then use the deformed atlas to generate smooth surfaces of objects and apply our model to deform the objects.

We have used our technique in a set of preliminary experiments to deform several subcortical structures and match them onto a target MRI scan. Although the technique has not yet been validated with a large number of datasets, the results of our experiments are extremely promising - the model appears robust to noise and in regions containing sufficient image feature information, generates visually accurate results. In these preliminary experiments, we use a merged left and right lateral ventricle, the left and right caudate nucleus, the corpus callosum, and the left and right putamen. Two deformation phases were used with 30 steps in the first phase and an element size of 4 voxels and 30 steps in the second with an element size of 2 voxels. This unoptimized version of our system is still quite efficient and each deformation step takes from 1 to 10 seconds (depending on the number of objects deformed concurrently and the number of elements). In the first experiment we deform the merged lateral ventricles, the

¹ Although a separate elastic matching program is currently used to initially deform the atlas, we will eventually incorporate this stage into our model framework.

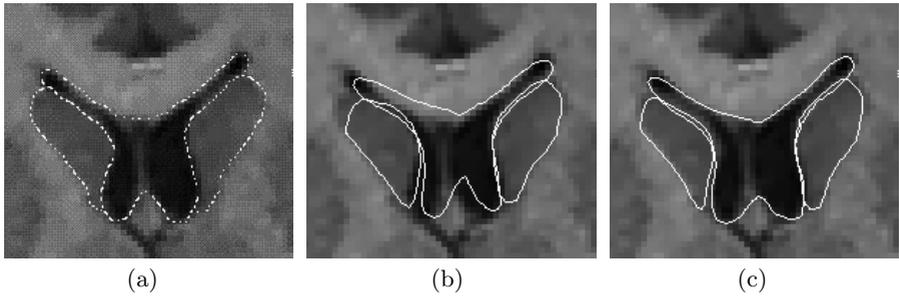


Fig. 3. (a) Manual tracing of left and right caudate nucleus and merged left and right lateral ventricle for slice 72, (b), (c) cross-section of initial and deformed surfaces.

left and right caudate nucleus and the surrounding volume (4 element layers are used). A significant segmentation improvement is obtained (figure 3(b, c)). In the second example, we deform the right putamen (Figure 4(a)(b)). The image intensity of the putamen varies considerably over its extent and its boundary is very noisy. For this reason, we integrate image feature information over a surface region centered at each surface vertex and then average this information to compute more reliable applied forces. In the final experiment we deform the corpus callosum (Figure 4(c)(d)). The strong edge features of this object results in a very accurate localization near the center of the object.

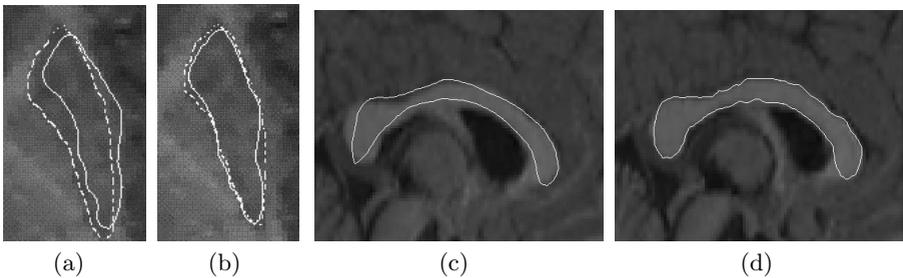


Fig. 4. Cross-section of initial and deformed surface of (a)(b) right putamen, (c) (d) corpus callosum. The dashed line in (a)(b) is a manual tracing.

5 Conclusion

We have created a 3D finite element-based object-oriented model to control the elastic deformation of a neuroanatomical atlas. The model framework provides us with the ability to deform not only the entire atlas or subvolumes of the atlas, but individual objects, separately or in combination. This ability allows us to accurately localize neuroanatomical structures in a target scan. The model has demonstrated considerable potential in preliminary experiments. We are

currently applying our model to a large number of datasets to validate its effectiveness. We are also exploring several model improvements with the goal of creating a single unified framework for the precise control of atlas deformation, ranging from the elastic deformation of deep subcortical structures, to a constrained viscoelastic deformation of the cortical surface. In particular, we intend to replace the cubical mesh lattice with a 3D finite element mesh that conforms to the geometry of each atlas object. The goal is to generate an axial-based or skeleton-based object parameterization so that the shape of an object can be constrained to maintain a feasible shape and its deformation can be controlled in a scheduled, global-to-local manner. This will allow us to make optimal use of object shape, symmetry, and image feature information.

References

1. R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Computer Vision, Graphics, and Image Processing*, 46:1–21, 1989.
2. K.J. Bathe. *Finite Element Procedures*. Prentice Hall, 1996.
3. F.L. Bookstein. Thin-plate splines and the atlas problem for biomedical images. In *Information Processing in Medical Imaging: Proc. 12th Int. Conf. (IPMI'91)*, Wye, UK, July, pages 326–342, 1991.
4. C. Broit. *Optimal Registration of Deformed Images*. PhD thesis, Computer and Information Science Dept., University of Pennsylvania, Philadelphia, PA, 1981.
5. C.A. Davatzikos, J.L. Prince, and R.N. Bryan. Image registration based on boundary mapping. *IEEE Trans. on Medical Imaging*, 15(1):112–115, Feb. 1996.
6. G. Christensen et al. Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing*, Sept. 1996.
7. Iosifescu et al. An automated measurement of subcortical brain mr structures in schizophrenia. *Neuroimage*, 6:13–25, 1997.
8. J.W. Snell et al. Model-based boundary estimation of complex objects using hierarchical active surface templates. *Pattern Recognition*, 28(10):1599–1609, 1995.
9. R. Kikinis et al. A digital brain atlas for surgical planning, model driven segmentation, and teaching. *IEEE: Visualization and Computer Graphics*, 2(3):232–241, 1996.
10. A.C. Evans, W. Dai, L. Collins, P. Neelin, and S. Marrett. Warping of a computerized 3D atlas to match brain image volumes for quantitative neuroanatomical and functional analysis. In *Medical Imaging V: Image Processing*, volume 1445 of *SPIE Proc.*, pages 236–246, 1991.
11. J. Gee, M. Reivich, and R. Bajcsy. Elastically deforming 3D atlas to match anatomical brain images. *Journal of Computer Assisted Tomography*, 17(2):225–236, March-April 1993.
12. D. McDonald, D. Avis, and A. Evans. Multiple surface identification and matching in magnetic resonance images. In *Proc. Third Conf. on Visualization in Biomedical Computing (VBC'94)*, Rochester, MN, October, 1994, pages 160–169, 1994.
13. P. Thompson and A.W. Toga. A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans. on Medical Imaging*, 15(4):402–417, August 1996.