Using 3-D Shape Models to Guide Segmentation of MR Brain Images

Kevin P. Hinshaw, M.S.* and James F. Brinkley, M.D., Ph.D.†
* Dept. of Computer Science and Engineering, Box 352350
† Dept. of Biological Structure, Box 357420
University of Washington, Seattle, WA 98195
hinshaw@cs.washington.edu, brinkley@u.washington.edu

Accurate segmentation of medical images poses one of the major challenges in computer vision. Approaches that rely solely on intensity information frequently fail because similar intensity values appear in multiple structures. This paper presents a method for using shape knowledge to guide the segmentation process, applying it to the task of finding the surface of the brain. A 3-D model that includes local shape constraints is fitted to an MR volume dataset. The resulting low-resolution surface is used to mask out regions far from the cortical surface, enabling an isosurface extraction algorithm to isolate a more detailed surface boundary. The surfaces generated by this technique are comparable to those achieved by other methods, without requiring user adjustment of a large number of ad hoc parameters.

INTRODUCTION

The need for methods of segmenting medical images has been well-established. Despite steady advances, current solutions still have difficulty because of the inherent visual complexity present in such images.

One common approach to segmentation of volume datasets is region growing.[1,2] This method works by building regions from seed points scattered throughout the volume, with each region gradually adding adjacent voxels which satisfy an inclusion cost function. A common drawback among region growing methods is leakage. In the case of brain magnetic resonance (MR) images, there are frequently small areas where the cortical surface is very close to the scalp. As a result, when the expanding regions for the brain and the scalp reach such an area, the region grower may inadvertently treat the two separate regions as components of a single, larger object. Such problems occur because these methods rely heavily on local information.

Classification methods are also popular for segmentation, especially for brain MR images.[3] Statistical tools are used to group voxels into clusters based on the distribution of their intensities. This type of classification typically requires multiple image volumes of the same patient, each acquired with a different imaging modality. This additional information helps considerably, but voxels are still misclassified because intensity values alone are not always sufficient to differentiate the structures present in the data.

We hypothesize that these difficulties might be alleviated if the segmentation process made use of shape knowledge about the structures in the dataset. In most cases, the goal of medical image segmentation is to locate the boundaries for a small number of objects, each of which has certain characteristic shape attributes. The results of segmentation could be improved by incorporating this type of information.

The method we advocate is most similar to a deformable model. First introduced by Kass *et al.*,[4] the idea behind deformable models is to treat segmentation as an optimization problem, typically by minimizing an energy function that rewards boundaries that are locally smooth and pass through high-gradient image regions. These models have been used successfully for many applications, and several researchers are currently investigating ways to use them to locate the surface of the brain.[5,6] However, deformable models are easily attracted to local minima, so their success often depends on getting a starting point which is close to the actual surface. In addition, they are often sensitive to the settings of non-intuitive parameters.

This paper offers an alternate but related approach to model-based segmentation, in which the model constraints are defined by the expected shape of the object, rather than by a set of ad hoc parameters. The next section presents a method for modeling a structure's shape as a network of local constraints. This is followed by a description of how the model can be used to identify the surface of an object in a volume dataset. In particular, attention is focused on segmenting the surface of the brain from a series of MR images. The results section illustrates the various stages of this technique. We then discuss the strengths and weaknesses of our approach and describe extensions that are being pursued.

THE RADIAL SURFACE MODEL

The model we employ is called a *radial surface model*.[7] It stores constraints that describe the basic shape and range of variation for a training set of *radial surfaces*. A radial surface consists of a series of parallel slices, each containing evenly-spaced radials that extend outward from its center to the surface boundary. Each surface also includes a local coordinate system that describes its orientation in 3-D space. For brain surfaces, this coordinate system can be defined by the landmarks of the Talairach reference system, which is commonly used for registration purposes in neurological studies.[8] These landmarks include the intersection of the mid-sagittal plane with the anterior and posterior commissures, and a bounding box around the cerebrum. This local coordinate system makes it possible to determine correspondences between radial measurements across different surfaces.

A radial surface model can be constructed from a set of radial surfaces as follows. First, in order to acquire local shape features, ratios of radial lengths are computed for each surface. For radial i and j with lengths r_i and r_j, this ratio is simply $s_{ij} = r_i/r_j$, and it is computed between each radial and its four most immediate neighbors in 3-D. Second, for each pair of neighboring radials i and j, the model stores a lower bound, L_{ij}, and an upper bound, U_{ij}, which correspond to the minimum and maximum values of s_{ij} that occur in the training set. These bounds constitute the model's shape constraints, measuring the expected range of variation in the relative lengths for pairs of adjacent radials.

SHAPE-GUIDED SEGMENTATION

The method for finding the boundary of a structure consists of three basic steps: generation of an initial low-resolution surface, creation of a voxel mask from it, and extraction of a detailed surface from the masked dataset.

Generating an Initial Surface

The first step of the segmentation process is fitting a radial surface model to the volume data. The user initiates this step by loading a particular shape model and entering the landmarks for its coordinate system, along with one or more initial radials. By applying a simple constraint propagation algorithm, these radials are used in combination with the shape model's local constraints to bound the lengths for other radials on the surface. Suppose the user specifies that radial j has length r_j. For each neighbor i of j, the constraint $L_{ij} \leq s_{ij} \leq U_{ij}$ is used to infer that $L_{ij}r_j \leq r_i \leq U_{ij}r_j$. These newly bounded radials can be used to generate bounds for their own neighbors, creating a wave of updates that propagates across the entire surface. As the segmentation progresses, an *uncertainty interval* for each radial keeps track of the range of values that satisfies both the model's constraints and the observed data.

After the propagation has completed, the new uncertainty intervals are searched with a one-dimensional edge detector to identify likely boundary locations. As edges are found, they can be used to tighten the uncertainty intervals further. The system enters a search-and-propagate loop, alternating between finding edges and propagating constraints, until it can no longer find any good edge candidates. The result will be a low-resolution approximation of the object's surface. Since some radials may be positioned incorrectly during this process, the system lets the user fix mistakes in the radial surface before continuing to the next stage.

Creating the Voxel Mask

Once an initial surface has been found, it is used as a mask to exclude regions outside the brain's surface. The radial surface model is polyhedral, so it must be converted into a voxel-based mask. This is accomplished in two steps. First, a shell is constructed by locating all voxels within a certain distance, α, of at least one facet of the model. Using $\alpha = 0$ produces a strict voxel rasterization of the radial surface; increasing the tolerance gives the shell thickness and rounds out its corners. Second, a flood fill algorithm is used to add all voxels inside the shell to the mask. Since the model is closed by construction, this filling operation will not spill into the region outside the shell.

The goal of this strategy is to avoid doing an expensive point-in-polyhedron test for every voxel in the dataset, because that would require comparing each voxel to every facet in the initial surface. By including any voxel that is within some distance of the surface, we can bypass all further voxel/facet checks for a particular voxel as soon as one sufficiently close facet is found.

Volume datasets contain millions of voxels, so the amount of per-voxel computation must be kept to a minimum. Therefore the mask generation stage has been optimized in two ways. First, each voxel is checked against the bounding box of the radial surface; if it is not within the distance tolerance of that box, then it can be discarded immediately. For the datasets we have analyzed, this simple check has ruled out approximate two-thirds of the voxels. Second, before computing the distance between a voxel and a facet, the distance between the voxel and the facet's plane is checked. The voxel-to-plane distance, d_p, is simple to compute, and we know that it cannot exceed the voxel-to-facet distance, d_f. If it turns out that $d_p > \alpha$, then $d_f > \alpha$, so the more complicated computation of d_f can be avoided. This second check eliminates roughly two-thirds of the voxel/facet pairs that survive the bounding box test. Combined, these two optimizations reduce the amount of computation required to generate the voxel mask by nearly an order of magnitude.

Extracting the Final Surface

After the voxel mask has been computed, we are ready to extract the final surface. First a new volume dataset is created by setting the intensity of all voxels outside the masked region to zero. This masked dataset is then passed to an isosurface extraction algorithm to convert its voxel-based representation of the brain into a surface-based one. The technique is based on the marching cubes algorithm of Lorensen and Cline,[9] using the method described by Bloomenthal to resolve topological ambiguities.[10]

RESULTS

The images in this section provide snapshots of the various stages of this shape-guided segmentation process. These images were captured from Scanner, our interactive segmentation system.[7] (See figure 1.) The brain shape model used in this example was constructed from three hand-drawn radial surfaces; each surface consisted of 20 coronal slices with 30 radials per slice. The datasets used for this study were acquired using a whole body 1.5 Tesla MR scanner, which produced 124 T1-weighted sagittal slices at 1.2mm spacing,

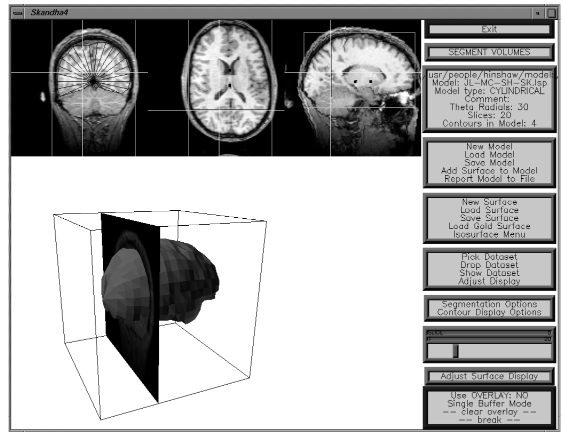

Figure 1: Snapshot of Scanner, our image segmentation system.

each with dimensions 256x256. Datasets were resampled to 128^3 prior to segmentation.

Figure 2 illustrates the process of fitting a radial surface model to a particular volume dataset. Note that the first two images show two surfaces — the darker, inner one connects the lower bounds of the radials; the lighter, outer one connects the upper bounds. Figure 2a shows the model's uncertainty region after propagating a single radial length through the shape constraint network. Figure 2b shows the region after the first pass through the search-and-propagate loop; the intervals on the central slice have been searched with a one-dimensional edge detector and the resulting edges propagated. After all of the slices have been searched, the system presents its best guess (figure 2c) to the user, who may then make corrections before continuing to the next stage. (Figure 2d.)

The results of converting the radial surface into a voxel mask appear in figure 3. Three orthogonal slices through the original volume are shown, along with their corresponding masked versions. Figure 4 illustrates the isosurface that was extracted from the masked dataset.

DISCUSSION

An appealing feature of our approach is that it is not limited to segmentation of a particular organ. Given sample surfaces for the kidney, for instance, we could use our method to construct a radial surface model for finding the surface of the kidney. Previous work has shown 2-D radial models to be useful for semi-automatic segmentation of CT images of the abdomen.[11] We intend to test the usefulness of our 3-D segmentation technique

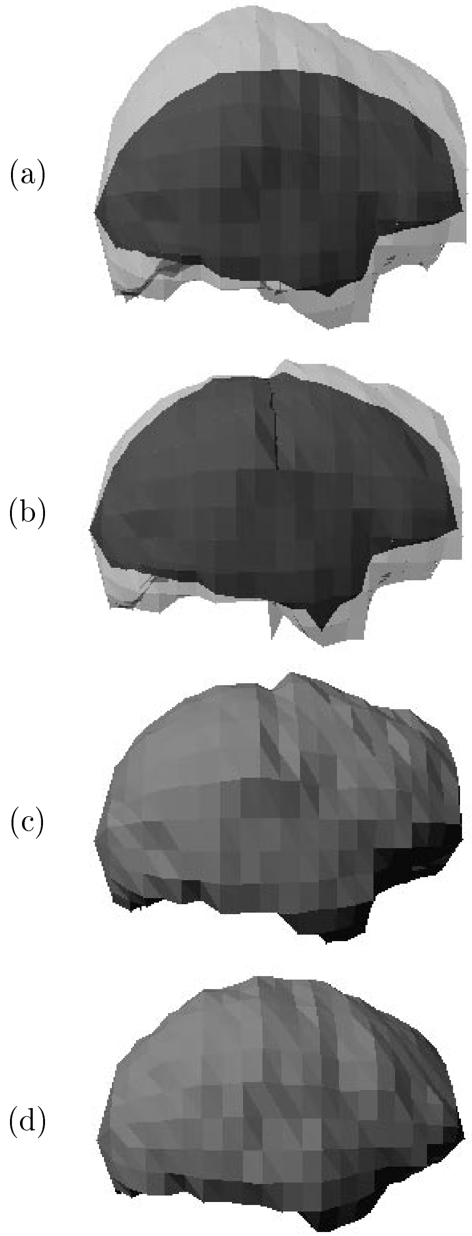

Figure 2: Progression of propagated shape constraints for a radial surface model of the cerebrum. (a) The inner (dark) and outer (light) uncertainty bounds after a single radial has been entered and propagated. (b) The intervals after an edge detector has searched the central slice and the results have been propagated. (c) The final surface found by the search-and-propagate phase. (d) The surface after it has been corrected by the user.

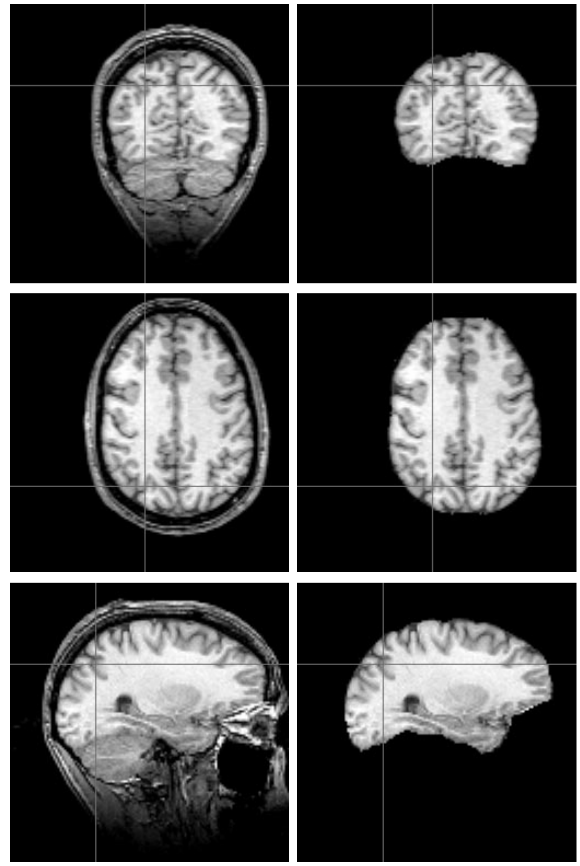

Figure 3: Three orthogonal slices through the original dataset and their corresponding masked versions.

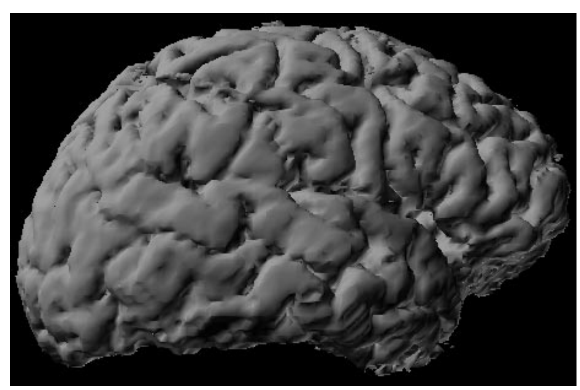

Figure 4: The final cortex segmentation achieved by extracting an isosurface from the masked dataset.

by building reconstructions of critical organs for radiation treatment planning.

We believe that fully automatic segmentation tools are still far from robust. We have therefore made a conscious choice to make our system interactive, so that the user can intervene to correct problems at key points in the segmentation process. The hope is that this approach will lead to a system that is truly useful in both research and clinical settings. We have already mentioned the user's ability to modify the radial surface before computing the voxel mask. Another interactive tool that we are currently developing will allow the user to remove the small bits of scalp that sometimes appear in the final surface when the voxel mask extends outward too far. This addition should lead to much cleaner renderings.

The current segmentation method depends on having a uniform intensity along the surface of the object. Some of the datasets we have analyzed have this property, but many do not — various artifacts can result during the imaging process, leading to nonuniform intensities for a single tissue. We can compensate for such problems by linking the threshold used by the isosurface routine to the intensities at the vertices of the initial surface. Since these vertices all lie on the surface of the brain, their intensities indicate the threshold that should be used at those points. By interpolating these values, we expect to derive better threshold values for other voxels in order to generate a more accurate final surface.

We have presented a method for using shape information to guide the segmentation of volume datasets. By incorporating models that contain local shape constraints, this technique deals with many of the problems of leakage and misclassification that occur with other segmentation methods. Furthermore, the integral role of user interaction in this system gives it the potential to be a valuable tool for both research and clinical applications.

Acknowledgements

This work was sponsored by National Cancer Institute Grant R29CA59070, Achievement Rewards for College Scientists (ARCS), and Human Brain Project grant LM/DC02310; co-sponsored by the National Library of Medicine and the National Institute on Deafness and other Communication Disorders. The project is part of the Digital Anatomist Program at the University of Washington. Thanks to Bharath Modayur for suggesting the initial idea behind the conversion from a radial surface to a voxel mask, and to Ken Maravilla for providing the MR data used for this study. We are also indebted to Jeff Prothero for continued software support, in particular the isosurface extraction routine.

References

1. Myers L and Brinkley JF. Visualization of brain surface features using registered partially segmented mri scans. In *Proc. SPIE Image Display*, volume 2431, pages 43–52, 1995.
2. Justice RK, Stokely EM, Strobel JS, Ideker RE, and Smith WM. Medical image segmentation using 3-d region growing. In *Proc. SPIE Image Processing*, volume 3034, 1997. In Press.
3. Choi HS, Haynor DR, and Kim Y. Partial volume tissue classification of multichannel magnetic resonance images — a mixel model. *IEEE Trans. Med. Imag.*, 10(3):395–407, 1991.
4. Kass M, Witkin A, and Terzopoulos D. Snakes: Active contour models. *Int. J. Comput. Vision*, 1(4):321–331, 1987.
5. MacDonald D, Avis D, and Evans AC. Multiple surface identification and matching in magnetic resonance images. In *Proc. SPIE Visual. Biomed. Comput.*, volume 2359, pages 160–169, 1994.
6. Schlesinger DJ, Snell JW, Mansfield LE, Brookeman JR, Ortega JM, and Kassell NF. Segmentation of volumetric medical imagery using multiple geodesic-based active surfaces. In *Proc. SPIE Image Processing*, volume 2710, pages 243–253, 1996.
7. Hinshaw KP and Brinkley JF. Shape-based interactive three-dimensional medical image segmentation. In *Proc. SPIE Image Processing*, volume 3034, 1997. In Press.
8. Talairach J and Tournoux P. *Co-planar stereotactic atlas of the human brain*. George Thieme Verlag, Stuttgart, 1988.
9. Lorensen WE and Cline HE. Marching cubes: a high resolution 3d surface construction algorithm. In *Computer Graphics*, volume 21, pages 163–170, 1987.
10. Bloomenthal J. Polygonization of implicit surfaces. *IEEE CG&A*, 5(4):341–355, 1988.
11. Brinkley JF. A flexible, generic model for anatomic shape: Application to interactive two-dimensional medical image segmentation and matching. *Comput. Biomed. Res.*, 26(2):121–142, 1993.