

From Sinograms to Surfaces: A Direct Approach to the Segmentation of Tomographic Data

Vidya Elangovan and Ross T. Whitaker

School of Computing, University of Utah
Salt Lake City, UT, USA

Abstract. Under ideal circumstances the problem of tomographic reconstruction is well-posed, and measured data are sufficient to obtain accurate estimates of volume densities. In such cases segmentation and surface estimation from the *reconstructed volume* are justified. In other situations the reconstructed volumes are not suitable for subsequent segmentation. This can happen in the case of incomplete sinograms, noise in the measurement process, or misregistration of the views. This paper presents a direct approach to the segmentation of incomplete and noisy tomographic data. The strategy is to impose a fairly simple model on the data, and treat segmentation as a problem of estimating the *interface* between two substances of somewhat homogeneous density. The segmentation is achieved by simultaneously deforming a surface model and updating density parameters in order to achieve a best fit between the projected model and the input sinograms. The deformation is implemented with level-set surface models, calculated at the resolution of the input data. Several computational innovations make the approach feasible with state-of-the-art computers. The usefulness of the approach is demonstrated by reconstructing the shape of spiny dendrites from electron microscope tomographic data.

1 Introduction

Certain kinds of tomographic reconstruction problems are ill-posed due to incomplete sinogram data. Difficulties in reconstructing volumes from such data are aggravated by noise in the measurements and misalignments among projections. This paper addresses the problem of segmentation in the context of such difficult tomography problems in biology and medicine. The usual approach for segmentation (indeed, any kind of post processing) in such cases is to reconstruct the volume, as best as possible, from the measured data using standard techniques such as filtered backprojection. Segmentation is then performed based on the grey-scale properties of that volume. However, the ill-posed nature of the reconstruction problem tends to produce various kinds of grey-scale artifacts, which state-of-the-art segmentation techniques cannot overcome. This paper presents a *direct* approach to segmentation, which uses the information in the sinograms instead of working with the reconstructed volumes and their associated artifacts. This direct strategy alleviates the effects of noise and misregistration and provides surface models of objects directly from incomplete sinograms.

The application studied in this paper is electron microscope tomography (EMT). While this data serves to motivate and demonstrate the proposed method, the characteristics of EMT are not unique, and the principles developed in this paper are applicable

to a variety of clinical and biological problems. EMT data, which deals with structures of very small dimensions (on the order of a few micrometers), has several inherent problems. First, there are technical limits to the projection angles from which data can be measured. These limits are due to the mechanical apparatus used to tilt the specimens and the tradeoff between the destructive effects of electron energy and the effective specimen thickness, which increases with tilt angle. Usually, the maximum tilt angle is restricted to about ± 60 – 70 degrees. Figure 1 (a) shows an illustration of the geometry of this limited-angle scenario. The second problem with EMT data is the degree of electron scattering, which results in projection images (sinograms) that are noisy relative to many other modalities, e.g. X-ray CT. Finally, due to the flexible nature of biological objects and the imperfections in the tilting mechanism, the objects undergo some movements while being tilted. Manual alignment procedures used to account for this tend to produce small misregistration errors. The segmentation approach described in this paper works well despite these limitations in the data.

Typically EMT data sets are produced using some kind of a contrast-enhancing mechanism (a dye) that highlights regions of interest, e.g. the interior of a cell. Thus, an ideal measurement and reconstruction would present a volume with two distinct densities (or electron opacities); one for the undyed regions (the background) and another for the dyed region (the foreground, or object). In practice, the densities are not homogeneous for either the object or the background. However, to the extent that the inhomogeneities are random and small relative to the contrast between the object and the background, the problem can be posed as estimating the interface between the object and the surrounding tissue and the densities of those two materials.

Our strategy, therefore, is to estimate the interface between two relatively homogeneous materials directly from the sinogram data, by deforming a surface model so that its forward projection is a best fit to the input data. For this work, the “best fit” is the sinogram that minimizes the mean square difference to the input. Other statistically-based metrics can be easily incorporated into this formulation, but that subject is beyond the scope of this paper. The relationship between *surface shape* and the associated projections is nonlinear, and we therefore use an iterative approach that deforms a surface model to achieve incrementally better fits. This requires a surface representation, which is sufficiently expressive to capture complex shape while allowing incremental, local deformations. For this we use a level-set surface model, which is a volume with grey-scale values that change according to a partial differential equation that controls surface movements. We consider the densities of the object and background as unknowns and estimate them simultaneously with the surface fitting. Refer to figure 1b for a schematic of the approach. We apply the proposed method to limited-angle, noisy, misregistered EMT data of a small section of a *spiny dendrite*. However, the basic assumptions of the approach are consistent with tomographic imaging modalities that are characterized by contrasting foreground and background, such as SPECT or PET.

The remainder of this paper is as follows. Section 2 gives a brief account of related work. Section 3 lays out the mathematical formulations of our segmentation strategy, the resulting surface deformation, the level-set evolution equations and the computational innovations that result in a fast and accurate implementation. Sections 4 and 5 give experimental results and conclusions.

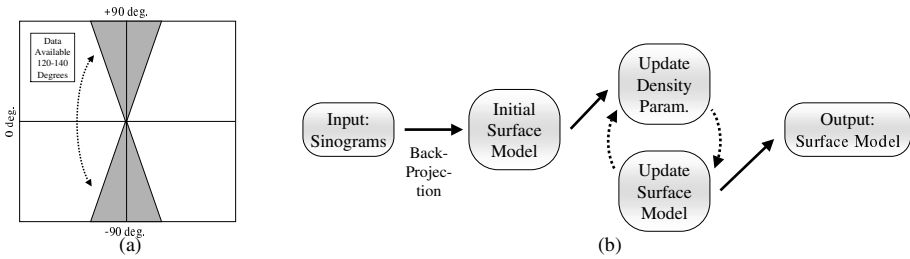


Fig. 1. (a) Illustration of limited-angle tomography, data missing in the shaded region (b) Schematic of the proposed method

2 Related Work

Several tomographic reconstruction methods are described in the literature [8], [13], and the method of choice depends on the quality of projection data. Filtered backprojection (FBP), the most widely used approach, works well in the case of the full-data reconstruction problem, where one is given *enough high-quality projections over 180 degree angular range*. Statistical, iterative approaches such as maximum-likelihood (ML) and maximum a posteriori (MAP) estimation have been proved to work well with *noisy* projection data, but generally rely on complete data sets. Some hybrid approaches [9], [12] have been specifically designed to deal with *limited-angle* tomography by “guessing” missing sinogram data via some extrapolation technique. A few researchers have proposed direct surface estimation using the projection data. For instance, Battle *et al.* [1], [2] use geometric deformable models for 3D tomographic reconstruction of objects of constant interior density in the context of SPECT imaging. They pose the inverse problem in a Bayesian framework and used a MAP estimate to compute the reconstruction of the shape of synthetic heart and lung phantoms represented as triangulated surface meshes. The fitting is accomplished by progressive surface refinement, but is limited in fidelity by the formulation and the surface representation. They demonstrate robustness to noise, but show results only for simulated data and do not study the effects of incomplete data.

The literature describes many examples of level-sets as curve and surface models for image segmentation [15], [10], [17], [3]. The authors have examined their usefulness for 3D segmentation of EMT *reconstructions* [19], the results of which are limited in quality due to reconstruction artifacts. Recently, several authors [4], [16] have proposed level-set formulations for finding piecewise smooth approximations to images in the context of segmentation. We assert that while level-sets are an interesting technology for modeling deformable surfaces, they are only as good as the data that drives their deformation—better results can be obtained by better formulations of the relationship between the models and the data.

Several authors have proposed solving *inverse problems* using level sets, with a more direct relationship between the data and the model. Santosa [14] lays out this basic strategy and demonstrates the idea on some small problems such as de-blurring 2D, binary images. Dorn *et al.* [5], [6] use this strategy to solve for permittivity using a

sequence of electromagnetic measurements, and show examples for simulated data in a 2D domain.

This paper makes several important contributions; first we describe a specific formulation of the direct surface estimation problem for tomographic data, second we present a level-set implementation of this formulation that is practical on large data sets, and third we demonstrate, using real data sets, the effectiveness of direct surface estimation for a specific class of tomographic problems that are subject to noise, limited-angle sinograms, and misregistration.

3 Mathematical Formulation

For brevity we give the basic formulation in two dimensions, the extension to three dimensions is different only in the mechanism for incorporating priors, or surface smoothing, which we describe later. The projection of a 2D signal $f(x, y)$ produces a sinogram given by the radon transform as

$$p(s, \theta) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(x, y) \delta(R_{\theta} \mathbf{x} - s) dx dy, \tag{1}$$

where $R_{\theta} \mathbf{x} = x \cos(\theta) + y \sin(\theta)$ is a rotation and projection of a point $\mathbf{x} = (x, y)$ onto the imaging plane associated with θ . The 3D formulation is the same, except that the signal $f(x, y, z)$ produces a collection of images, which we denote by $p(r, s, \theta)$.

The goal is to estimate the shapes of regions that have a density β_1 , which is different from the background, with density β_0 . We denote the (open) set of points in those regions as Ω , the closure of that set, the surface, as \mathcal{S} , and the projection of the model and background as $\widehat{p}(s, \theta)$. The goal is to find Ω , β_0 , and β_1 to minimize

$$E_{\text{data}} = \sum_{i=1}^M \sum_{j=1}^N \left[\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(x, y) \delta(R_{\theta_i} \mathbf{x} - s_j) dx dy - p(s_j, \theta_i) \right]^2, \tag{2}$$

where $\theta_1, \dots, \theta_M$ are the angles at which projections are taken and s_1, \dots, s_N are the samples taken for each projection angle. If we let $g(x, y)$ be a binary inside-outside function on Ω , then we have the following approximation to $f(x, y)$:

$$f(x, y) \approx \beta_0 + (\beta_1 - \beta_0)g(x, y), \tag{3}$$

and the energy becomes

$$\sum_{i=1}^M \sum_{j=1}^N \left[\beta_0 K(s_j, \theta_i) + (\beta_1 - \beta_0) \int_{\Omega} \delta(R_{\theta_i} \mathbf{x} - s_j) d\mathbf{x} - p(s_j, \theta_i) \right]^2, \tag{4}$$

where $K(s_j, \theta_i)$ depends on the geometry of the region over which the data is taken, and is independent of the surface estimate.

We use a gradient descent minimization approach for finding local deformation of the surface with respect to an evolution parameter t , which gives Ω_t and \mathcal{S}_t , and we

denote the movement of a point \mathbf{x} on the surface as $\frac{\partial \mathbf{x}}{\partial t}$. The first variation of equation (4) with respect to a small movement of the surface $\delta \mathbf{x}$ is

$$\delta E_{\text{data}} = \left(\sum_{i=1}^M [\widehat{p}(R_{\theta_i} \mathbf{x}, \theta_i) - p(R_{\theta_i} \mathbf{x}, \theta_i)] \right) \mathbf{N}(\mathbf{x}) \cdot \delta \mathbf{x}, \quad (5)$$

where $\mathbf{N}(\mathbf{x})$ is the surface normal at a point $\mathbf{x} \in \mathcal{S}_t$. Thus the surface deformation is

$$\frac{\partial \mathbf{x}}{\partial t} = \left(\sum_{i=1}^M [\widehat{p}(R_{\theta_i} \mathbf{x}, \theta_i) - p(R_{\theta_i} \mathbf{x}, \theta_i)] \right) \mathbf{N}(\mathbf{x}). \quad (6)$$

Because the data are noisy and incomplete it is useful to impose some other considerations on the surface estimate to maintain smoothness. This can be considered as a kind of regularization or as a prior—if one prefers a statistical formulation. One choice of a prior that is described in the literature is to penalize the surface area, in which case the first variation is the mean curvature. An alternative prior [17], used in this paper, is the weighted sum of principal curvatures that does not penalize linear structures such as cylinders. Curvature motion is in the normal direction, and the surface motion becomes

$$\frac{\partial \mathbf{x}}{\partial t} = \left(\sum_{i=1}^M [\widehat{p}(R_{\theta_i} \mathbf{x}, \theta_i) - p(R_{\theta_i} \mathbf{x}, \theta_i)] + C\kappa(\mathbf{x}) \right) \mathbf{N}(\mathbf{x}), \quad (7)$$

where C is a user-specified constant that controls the influence of the smoothness term relative to the data, and κ denotes a curvature-based smoothing term.

3.1 The Level-Set Deformation

The level-set method, proposed by Osher and Sethian [15], [11], relies on an implicit representation of a deformable surface model \mathcal{S}_t :

$$\mathcal{S}_t = \{\mathbf{x} | \phi(\mathbf{x}, t) = k\}, \quad (8)$$

where $\mathbf{x} \in \mathbb{R}^3$, $\phi(\mathbf{x}, t) : \mathbb{R}^3 \times \mathbb{R}^+ \mapsto \mathbb{R}$, and k , usually set to be zero, is arbitrary and does not affect the formulation. Models evolving in this manner are not limited by topological considerations and can take a wide range of possible shapes. The evolution of ϕ that corresponds to the surface motion is

$$\frac{\partial \phi}{\partial t} = -\nabla \phi \cdot \frac{\partial \mathbf{x}}{\partial t}. \quad (9)$$

Substituting the expression for $\frac{\partial \mathbf{x}}{\partial t}$, with $\mathbf{N} = \nabla \phi / |\nabla \phi|$, gives

$$\frac{\partial \phi}{\partial t} = -|\nabla \phi| \left(\sum_{i=1}^M [\widehat{p}(R_{\theta_i} \mathbf{x}, \theta_i) - p(R_{\theta_i} \mathbf{x}, \theta_i)] + C\kappa(\mathbf{x}) \right). \quad (10)$$

This equation is solved by finite differences on a discrete grid, i.e. a volume. Osher and Sethian [11] have proposed an up-wind method for solving these equations in

a stable fashion, and a variety of researchers have proposed computationally efficient algorithms. We use the sparse-field algorithm [18], which computes updates on a (relatively) small set of grid points, called the *active layer*, which is only one point wide, and can position level-set surface models to sub-voxel accuracy.¹ The curvature, κ , is computed as described in the literature [15], [17] using central differences on the discrete approximation to ϕ .

3.2 Initialization

The deformable model fitting approach requires an initial model, i.e. $\phi(\mathbf{x}, t = 0)$. One reasonable way to obtain this is by thresholding the “best” information available prior to our solution, which is a volume reconstructed by filtered backprojection. In practice we do not require the initial model to be particularly close to the desired solution, and thus the threshold is not a critical parameter.

3.3 Estimation of Density Parameters

We consider the object and background to be homogeneous substances of *unknown* densities β_1 and β_0 respectively, and these density parameters affect the error term in equation (4). We update the estimate of the surface model iteratively, and at each iteration we reestimate the quantities β_1 and β_0 in such a way that the energy, E_{data} is minimized. Treating Ω as constant, equation (4) is quadratic in the density parameters. Thus, β_1 and β_0 are computed from the following linear system:

$$\frac{\partial E_{\text{data}}}{\partial \beta_0} = 0, \quad \frac{\partial E_{\text{data}}}{\partial \beta_1} = 0. \quad (11)$$

Variations in instrumentation can cause variations in the brightness levels of the images taken at different angles. In such cases we estimate sets of such parameters, i.e., $\beta_1(\theta_i)$ and $\beta_0(\theta_i)$ for $i = 1 \dots M$.

3.4 Speed and Accuracy Considerations

Because we are combining the reconstruction problem (which is known to be time consuming) with the level-set deformation, our approach is computationally intensive. Specifically, computing $\hat{p}(r, s, \theta)$ is a major bottleneck. Computing this term involves re-computing the sinogram of our model as it moves. In the worst case, we would re-project the entire model every time the surface moves.

To address this computational concern, we have designed and implemented a method, that we call *incremental projection update* (IPU). Rather than fully recompute \hat{p} at every iteration, we maintain a current running version of \hat{p} and update it to reflect the changes in the model as it deforms. Changes in the model are computed only on a small

¹ The implementation in this paper is built on VISPACk, a C++ open-source software library for processing images, volumes and level-set surface models. The library is available at <http://www.cs.utah.edu/~whitaker>.

set of grid points in the volume, and therefore the update time is proportional to the area of the surface, rather than the size of the volume it encloses.

The IPU strategy works with the the sparse-field algorithm [18] as follows. At each iteration, the sparse-field algorithm updates only the active layer (one voxel wide) and modifies the set of active grid points as the surface moves. The incremental projection update strategy takes advantage of this to selectively update the model projection to reflect those changes. At each iteration, the amount of change in an active point's value determines the direction of motion of that particular surface point. This quantitative measure determines what portions of the projection \hat{p} change (and by how much), and we can update it accordingly. Thus, the IPU maintains sub-voxel accuracy at a relatively low computational cost.

4 Experimental Results

In order to verify the correctness of the proposed method we first present results of digitally simulated slices (2D) for comparison against a ground truth, $f(x, y)$. The synthesized image is a homogeneous object on a homogeneous background. Projections of 128 samples each span 134 degrees (limited-angle) at two degree increments for a total of 67 views. To simulate misregistration errors, we randomly translate the image by plus/minus one pixel before taking the projections, and we corrupt each projection with additive, independent, Gaussian noise. Figure 2 shows the segmentation results for one simulated 2D slice. The input image (figure 2a) has three ellipses with two of them connected to each other. The incomplete, noisy, misaligned sinogram from this simulation (figure 2b) results in a rather poor backprojection (figure 2d). Figures 2c and 2f show the initial model and the final model after 150 iterations respectively overlaid on the input data. The initial model obtained from the backprojection is not a good estimate, it contains parts that do not belong to the object and it is missing some other pieces. The deformation process corrects these errors, and the final model captures all these connections and accurately estimates the contour corresponding to the input data. The final estimated sinogram, shown in figure 2e, demonstrates that the proposed method fills in the missing information in a reasonable way.

We also applied our algorithm to 3D EMT data obtained from a 3 MeV UHVEM. This 3D data set consists of 67 tilt series images, each corresponding to one view of the projection. Each tilt series image is of size 424x334. The volume reconstructed by FBP is of size 424x424x334. Figure 4a and 4b show the sinogram corresponding to a single slice of this data set and the estimate of the same sinogram created by the method. Figure 4e shows the surface estimate intersecting this slice overlaid on the backprojected slice. Some structures not seen in the backprojection are introduced in the final estimation, but the orientation of the structures introduced suggests that these structures were in the original object but were lost in the reconstruction artifacts during the backprojection. This hypothesis is consistent with the results of the simulated data. Also, the proposed method captures line-by-line brightness variations in the input sinogram (as explained in section 3.3). This suggests that the density estimation procedure is correct. Figure 5 shows the 3D initialization and the final 3D surface estimate. The figure also




Fig. 2. Results of a 2D simulation: (a) Digitally simulated input image (b) limited-angle, noisy, misaligned sinogram created by projecting the input image (c) Initial model obtained by thresholding the backprojection (d) Backprojection showing artifacts (e) Sinogram estimated by the proposed method (f) Final model showing the correct segmentation of the input (Note: The initial and final contours are overlaid on the input data)

shows enlarged initial and final versions of a small section of the surface. The enlarged versions clearly illustrate the missing structures getting filled in.

Figure 3a shows a plot of the percentage error, relative to the RMS magnitude of the input sinogram, versus number of iterations. The error converges to a constant value within approximately 50 iterations for the simulated data, whereas it takes about 150 iterations for the real dendrite data. This is justified by the complexity of the real data compared to the simulated data. Also, the percentage error is lower for simulated data, which suggests the higher degree of noise and inhomogeneities in the EMT data. Figure 3b shows the convergence of the average value of the density parameter $\beta_1(\theta_i)$ for a single slice of the EMT dendrite data. The final surface estimate for the EMT dendrite data required 250 iterations, which took approximately 5 hours on a single processor of Silicon Graphics Onyx2 workstation.

5 Conclusions

We have demonstrated direct segmentation and surface modeling using sinograms for difficult tomography problems. Our results show that this approach is better than working with reconstructed volumes and their associated artifacts. In the particular case EMT dendrite data, the method appears superior to the current practice of hand segmentation, which requires more time and fails to capture the same level of detail.

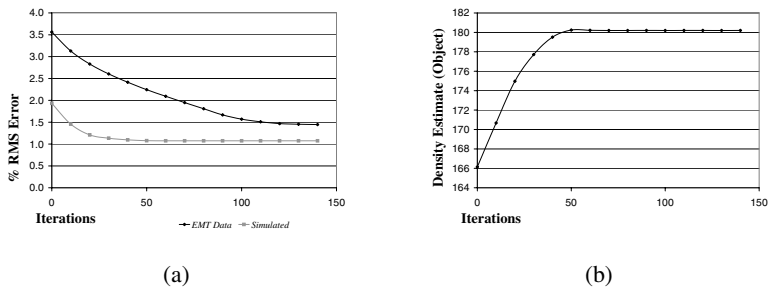


Fig. 3. Convergence plots: (a) Percentage error versus number of iterations for simulated and dendrite data. (b) The estimated object density parameter (average of $\beta_1(\theta_i)$ over all angles) for the EMT dendrite data versus the number of iterations

This work promises a number of interesting future directions. One direction is to extend the method to work for non homogeneous substances or multiple densities. In principle, one can modify the level-set equations to accommodate non-homogeneous density functions, that could also be estimated simultaneously with the surface estimate. Also, one could have several, interacting surface models, each one enclosing a different substance. Another direction is to apply the approach to other kinds of tomography, which have the similar problems. For example, a limited-angle problem occurs in cardiac CT imaging, where the carriage containing X-ray emitters and detectors can only travel part of the way through the full angular range before significant heart motion occurs. Finally, preliminary results show that the method can reconstruct surface shapes even in the extreme case of very few projections. This suggests that the approach could have applications in angiography and fluoroscopy.

6 Acknowledgments

This work is supported by the National Science Foundation under grants 0089915 and 9982273, the Office of Naval Research under grant N00014-01-10033, and the National Library of Medicine *Insight* project. The authors thank the National Center for Microscopy and Imaging Research (NCMIR) at San Diego for providing the EMT spiny dendrite data and Gordon Kindlmann for his valuable suggestions and volume analysis tools.

References

1. Battle, X.L., Cunningham, G.S., Hanson, K.M. (ed.): 3D tomographic reconstruction using geometrical models. Proc. SPIE Medical Imaging: Image Processing **3034** (1997) 346–357
2. Battle, X.L., Bizais, Y.J., Le Rest, C., Turzo, A., Hanson, K.M. (ed.): Tomographic reconstruction using free-form deformation models. Proc. SPIE Medical Imaging: Image Processing **3661** (1999) 356–367




Fig. 4. 2D slice of dendrite data: (a) Sinogram of one slice (b) Sinogram estimated by the method for the same slice (c) Backprojection showing artifacts (d) Initial model obtained by thresholding the backprojection overlaid on the backprojection (e) Final model overlaid on the backprojection




Fig. 5. 3D results: (a) Surface initialization (b) Final surface estimated after 250 iterations (c) A portion of the initial surface enlarged (d) The corresponding portion in the final surface

3. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic Active Contours. International Conference on Computer Vision, IEEE Computer Society Press (1995) 694–699
4. Chan, T.F., Vese, L.A.: A Level Set Algorithm for Minimizing the Mumford-Shah Functional in Image Processing. UCLA, Department of Mathematics, Technical report, CAM **00-13** (2000)
5. Dorn, O., Miller, E.L., Rappaport, C.: A shape reconstruction method for electromagnetic tomography using adjoint fields and level sets. *Inverse Problems: Special issue on Electromagnetic Imaging and Inversion of the Earth's Subsurface* **16** (2000) 1119–1156
6. Dorn, O., Miller, E.L., Rappaport, C.: Shape reconstruction in 2D from limited-view multi-frequency electromagnetic data. *to appear AMS series contemporary mathematics* (2001)
7. Frank, J.: *Electron Tomography: Three-Dimensional Imaging with the Transmission Electron Microscope*. New York Plenum Press (1992)
8. Herman, G.T.: *Image Reconstruction from Projections: The Fundamentals of Computerized Tomography*. Academic Press, New York (1980)
9. Inouye, T.: Image Reconstruction with Limited Angle Projection Data. *IEEE Transactions on Nuclear Science* **NS-26** (1979) 2666–2684
10. Malladi, R., Sethian, J.A., Vemuri, B.C.: Shape Modeling with Front Propagation: A Level Set Approach. *IEEE PAMI* **17(2)** (1995) 158–175
11. Osher, S., Sethian, J.A.: Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics* (1988) 12–49.
12. Prince, J.L., Willsky, A.S.: Hierarchical Reconstruction Using Geometry and Sinogram Restoration. *IEEE Transactions on Image Processing* **2(3)** (1993) 401–416
13. Roerdink, J.B.T.M.: Computerized tomography and its applications: a guided tour. *Nieuw Archief voor Wiskunde* **10(3)** (1992) 277–308
14. Santosa, F: A level set approach for inverse problems involving obstacles. *European Series in Applied and Industrial Mathematics: Control Optimization and Calculus of Variations* **1** (1996) 17–33
15. Sethian, J.A.: *Level Set Methods: Evolving interfaces in Geometry, Fluid Mechanics, Computer Vision, and Material Sciences*. Cambridge University Press (1996)
16. Tsai, A., Yezzi, A., Jr., Willsky, A.: A Curve Evolution Approach to Smoothing and Segmentation Using the Mumford-Shah Functional. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2000) 119–124
17. Whitaker, R.T., Robb, R.A. (ed.): *Volumetric Deformable Models: Active Blobs*. SPIE Visualization In Biomedical Computing (1994) 122–134
18. Whitaker R.T.: A Level-Set Approach to 3D Reconstruction From Range Data. *International Journal of Computer Vision* **29(3)** (1998) 203–231
19. Whitaker, R.T., Breen, D.E., Museth, K., Soni, N.: A Framework for Level Set Segmentation of Volume Datasets. *to appear ACM Volume Graphics Workshop* (2001)