

Segmentation of Bone in Clinical Knee MRI Using Texture-Based Geodesic Active Contours^{*}

Liana M. Lorigo¹, Olivier Faugeras^{1,2}, W.E.L. Grimson¹, Renaud Keriven^{2,3},
and Ron Kikinis^{4,5}

¹ MIT AI Laboratory, Cambridge MA, USA
liana@ai.mit.edu

² INRIA, Sophia Antipolis, France

³ Cermics, ENPC, France

⁴ Dept. Of Radiology, Brigham and Womens Hospital, Boston MA, USA

⁵ Harvard Medical School, Boston MA, USA

Abstract. This paper presents a method for automatic segmentation of the tibia and femur in clinical magnetic resonance images of knees. Texture information is incorporated into an active contours framework through the use of vector-valued geodesic snakes with local variance as a second value at each pixel, in addition to intensity. This additional information enables the system to better handle noise and the non-uniform intensities found within the structures to be segmented. It currently operates independently on 2D images (slices of a volumetric image) where the initial contour must be within the structure but not necessarily near the boundary. These separate segmentations are stacked to display the performance on the entire 3D structure.

1 Introduction

We address the problem of automatically segmenting clinical MRI of knees. There are many applications of this capability, including diagnosis, change-detection, as a pre-cursor to registration with a model, and in the building of an initial model for surgical simulation. Moreover, the segmentation system can be used as a tool to replace or expedite the tedious process of manual segmentation.

There are two primary strategies for detecting image boundaries. Region growing uses local properties of the region of interest along with techniques for splitting and merging regions [2, 1]. Deformable or “active” contour models (“snakes”) use only the boundary of the object which deforms over the image to minimize some energy function based on properties of the curve (such as smoothness) and of the image (such as gradients) [12, 6, 18]. Several methods explicitly combine both the region and contour approaches [23, 16]. Active boundary methods are commonly applied to medical images because they can capture the irregular shapes and shape deformations found in anatomical structures [14]. Several

^{*} This report describes research supported in part by NSF under contract 1R1-9610249 and in part by MERL, A Mitsubishi Electric Research Laboratory.

approaches to knee segmentation have been explored, including a model based approach to segmenting both bone and cartilage using region-growing, snakes, and spatial information [11], and a knowledge-based approach which uses information about the shape and the imaging process to localize the femur [19]. While such approaches are important for utilizing anatomical knowledge in guiding segmentation, we believe that new measurements of similarity can improve basic segmentation methods and provide a stronger basis for knowledge-rich segmentation. We use a deformable boundary model to capture the shape, but we extend it to include texture information for better performance.

In particular, trabecular bone has an apparent visual texture in MRI data which is approximately constant; it is this texture pattern that we wish to exploit. This implies that the method is best-suited to handle the portions of the bones near joints, since these portions contain much trabecular structure. The shafts of the bones contain much less of this structure and may not be as readily segmented by our method, depending on their appearance in the MR images. For this work, we address the problem of segmenting the trabecular bone in the portions of the femur and tibia near their respective joints. The method has not been applied to other areas of the bones. It should be further noted that we do not find cortical bone directly in this work. Should that segmentation be desired, our method may be useful as a pre-processing step, since it can be viewed as providing a rough segmentation of the interior boundary of cortical bone.

Segmenting the trabecular region of the bone can also be viewed as classifying the pixels in that region, since the boundary is initialized to contain intensity and texture corresponding to trabecular bone, then grows outward to find the true boundary of that bone region. However, no classification is performed on the rest of the image, and the classification of trabecular bone is performed locally only.

For the deformable model, we use geodesic active contours (“geodesic snakes”) which are described below [4, 3, 13]. The advantage of geodesic snakes over classical snakes [12] is that the former are independent of the parameterization of the curve and can handle topological changes automatically. Moreover, implementation by level-set methods provides accuracy and stability [15]. This segmentation method has been applied to various medical imaging domains in [20].

The energy function used by either active contour method is normally based on the intensity gradients in the image so the snake will lock onto strong edges. MRI images, however, are often too complex for gradient information alone to be reliable. Intensities often vary non-uniformly throughout a single structure and the boundary between neighboring structures may be noisy. This is where a less local approach, moving toward region-growing, is of benefit. Specifically, texture information can be incorporated to model these image attributes.

Many different distributions can be used to represent or learn texture classes [21, 22]. Additionally, filters at different scales can be used to decompose an image into low-level texture features [7]. Texture modeling is used for denoising, texture synthesis, classification, and segmentation. We have chosen to incorporate texture information directly into the weighting function of the geodesic

snakes. We add only variance at this time, but more statistics such as directional filters or multiscale filters could be incorporated easily if needed. Specifically, we define the snake to be attracted to intensity gradients as well as variance gradients, where variance is computed in a small neighborhood around each pixel. We use the vector-valued snakes approach to combine these measures in a manner formally consistent with the original snakes formulation [17].

This paper presents a system which automatically detects closed boundaries in 2D magnetic resonance images of knees. Both the femur and the tibia are segmented separately from each image of a volumetric set. Also input must be a small patch that is known to be within the desired structure but does not need to be anywhere near the true boundary; the snake flows outward from this patch until convergence. This is an improvement over many similar optimization approaches where the initial contour must be near the boundary, and the user may need to guide the contour. Each 2D image is segmented independently at this time, with the results stacked to give the volumetric segmentation. Future work will generalize the method to 3D. No user-interaction is required.

2 Geodesic Snakes

This work uses the framework of geodesic active contours as in [4, 3, 13, 17].

2.1 . . . Basic Formulation

The task of finding the curve that best fits the boundary of the object is posed as a minimization problem over all closed planar curves $C(q) : [0, 1] \rightarrow R^2$. The objective function is

$$\int_0^1 g(|\nabla I(C(q))|) |C'(q)| dq$$

where $I : [0, a] \times [0, b] \rightarrow R^+$ is the image and $g : [0, \infty) \rightarrow R^+$ is a strictly decreasing function such that $g(r) \rightarrow 0$ as $r \rightarrow \infty$. That is, we are looking for the minimal distance curve where distance is weighted by the function g which acts on intensity gradients, so that the curve is attracted to intensity edges in the image. For example, one common choice is $g(|\nabla I|) = \frac{1}{1+|\nabla I|^2}$.

To minimize this objective function by steepest descent, consider C to be a function of time t as well as parameter q . Then one can compute the Euler-Lagrange equation of the objective function to determine the evolution of the curve, i.e., its derivative with respect to time. This yields the curve evolution equation

$$\frac{\partial C(t)}{\partial t} = g\kappa \mathbf{N} - (\nabla g \cdot \mathbf{N}) \mathbf{N}$$

where κ is the Euclidean curvature and \mathbf{N} is the unit inward normal.

In order to make this flow intrinsic to the curve (independent of its parameterization), a surface $u : [0, a] \times [0, b] \rightarrow R$ can be defined to give the distance

from any image point to the curve $C(\cdot)$. Then u is evolved instead of C , which is identically the zero level-set of u . The evolution equation is

$$u_t = g\kappa|\nabla u| + \nabla g \cdot \nabla u.$$

An additional term c can be used either to increase the speed of the flow or to force the contour to flow outward, similar to the balloon force in [6], to yield

$$u_t = g(c + \kappa)|\nabla u| + \nabla g \cdot \nabla u.$$

2.2 ... on Vector-Valued Images

As described above, I is assumed to have one value at each pixel. The same minimization can be achieved on multi-valued images [17]. Let $\Phi(u_1, u_2) : R^2 \rightarrow R^m$ be an m -valued image, so the value of the image at any point is a vector in R^m . The quadratic form $d\Phi^2(\mathbf{v})$ gives the rate of change of the image in the \mathbf{v} direction:

$$d\Phi^2 = \begin{bmatrix} du_1 \\ du_2 \end{bmatrix}^T \begin{bmatrix} g_{11} & g_{12} \\ g_{21} & g_{22} \end{bmatrix} \begin{bmatrix} du_1 \\ du_2 \end{bmatrix}$$

where $g_{ij} := \frac{\partial \Phi}{\partial u_i} \cdot \frac{\partial \Phi}{\partial u_j}$. The extrema of $d\Phi^2$ are obtained in the directions of the eigenvectors of $[g_{ij}]$, the attained values are the corresponding eigenvalues λ_{\pm} .

We then want to define our function g to be used in the evolution equations according to the steepest-descent method. One approach is to make g be a function of either λ_+ or $(\lambda_+ - \lambda_-)$, still requiring that $g(r) \rightarrow 0$ as $r \rightarrow \infty$.

3 Segmentation Algorithm

Our segmentation algorithm uses these evolution equations in conjunction with local variance [10]. The algorithm also uses an image-dependent balloon force [9], requires an initial contour, and detects convergence automatically. Finally, a post-processing step is added to counteract the effects of smoothing and windowing on the energy function g .

3.1 Incorporating Texture

Although the trabecular bone regions of the MR images vary in intensity, there is a fairly uniform texture throughout those regions. This led us to use variance along with intensity for segmentation. In particular, local variance $S : [0, a] \times [0, b] \rightarrow R^+$ is computed over a fixed-sized window. We now treat the original image as a two-valued image with intensity and local variance as the values at each pixel. We choose $g = \frac{1}{1 + \sqrt{\lambda_+}}$, computed at each pixel in the image, and use the geodesic snakes formulation. An example of g on an image is shown in Figure 1.

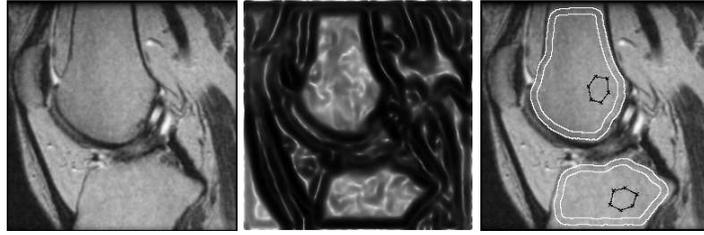


Fig. 1. Intermediate steps in algorithm. (1) An example image slice. (2) The weighting function g used in minimization, operating on that slice. (3) The output of the segmentation algorithm: the inner boundary is the true boundary based on g , the outer boundary is the result after compensation for windowing and smoothing, the starting contour is shown in black.

3.2 Image-Dependent Balloon Force

Our system performs outward flow with the balloon force proportional to g at each point on the curve. This means that the curve will push outward quickly when g is high (far from the boundary) and slowly when g is near zero. Note that this is an important deviation from the definition of geodesic active contours which may change the behavior of the partial differential equation, affecting the existence and uniqueness of the solution.

3.3 Initial Contour

The algorithm requires an initial closed curve which must be within the structure to be segmented. Our system does not require it to be anywhere near the actual boundary as shown in Figure 1. The initial contour does not need to be connected but could be a set of “bubbles” within the object or objects to be segmented.

3.4 Convergence

Convergence is detected automatically when the curve has not moved with respect to the image resolution over a fixed number of iterations. For the parameter settings used in our experiments, this number could be set between 10 and 20 with good results.

3.5 Windowing Effect

Because g is computed over a window of the image, the final contour more closely matches the centers of the respective windows which is a fixed distance (the radius of the window) inside of the true boundary. Additionally, the smoothing used in the computation of the gradients of the intensity and variance images shrinks the apparent boundary of the structure related to the amount of smoothing and to the structure’s curvature. For these reasons the final contour is expanded by

a constant k at each point as shown in Figure 1, where k depends on the window size used in the variance computation and the amount of smoothing. This is only an approximation, as k should also vary locally depending on the curvature of the boundary and on the appearance of neighboring structures. These factors are not currently incorporated, and k is constant over the curve.

4 Experimental Results

This algorithm has been run on one clinical volumetric knee image. This image consists of 45 2D slices (greyscale images), each with resolution 256x256. The slices on the ends were not used in the experiments since the femur and tibia were not prominent in them, leaving 37 slices used. The program is run on each slice independently. The same initial contours were used in each image, except for slices early in the sequence and late in the sequence where separate initial curves had to be defined so that the initial curve was within the structure. The femur and tibia were segmented separately, and are overlaid on the same image for display purposes only. The window-size for variance computation was 7x7, sigma of 4 was used in the gradient computations, the compensation factor k was 8, and convergence was defined to be movement less than image resolution over 12 iterations. Most segmentations required approximately 200 iterations for convergence, dependent primarily on the size of the structure in the image, and took approximately 30 seconds on a 250MHz UltraSparc.

4.1 Performance

Segmentations of the 35 central slices are pictured in Figure 2. The segmentation was performed separately for the femur and tibia, although this is not a requirement of the algorithm. Notice that the boundary is well approximated by most segmentations, and the starting contours are not near the boundaries. These segmentations were fully automatic, with fixed initial contours used for the early, middle, and late slices respectively, and all other parameters constant on all images. The exact settings were not crucial; many settings yield qualitatively equivalent results. The convergence criteria can be tightened to allow more outward flow, and this change would fix the error in the femur segmentation in one image in Figure 2 where a dark region within the femur causes the flow to stop. If the criteria is made too strict, however, the contours would leak into neighboring structures. The segmentations are rendered together to show the total segmentation of the bones in Figure 3.

4.2 Comparison to Intensity Only

This algorithm was run using only the intensity gradients, as in the traditional active contour definition. In this case, the energy function was defined as $g(|\nabla I|) = \frac{1}{1+|\nabla I|^2}$. This method worked for some images and settings but was much less stable than the two-valued method. Areas of low gradients (high

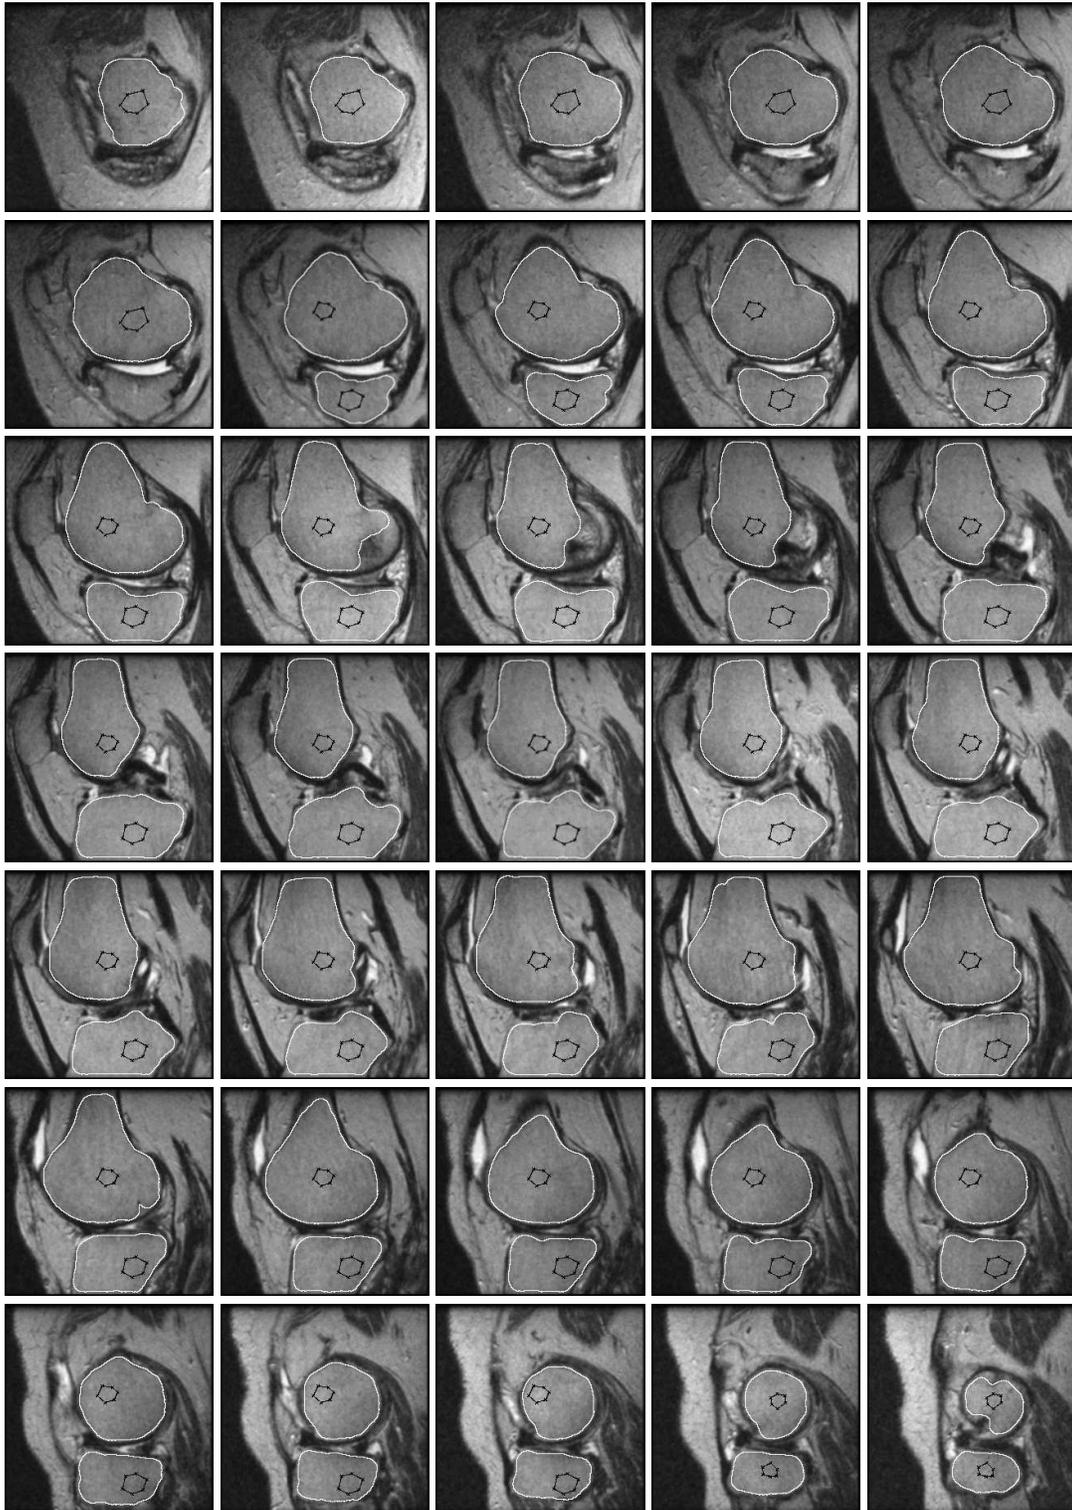


Fig. 2. 35 image slices shown with segmentations overlaid in white. Initial contour is shown in black.

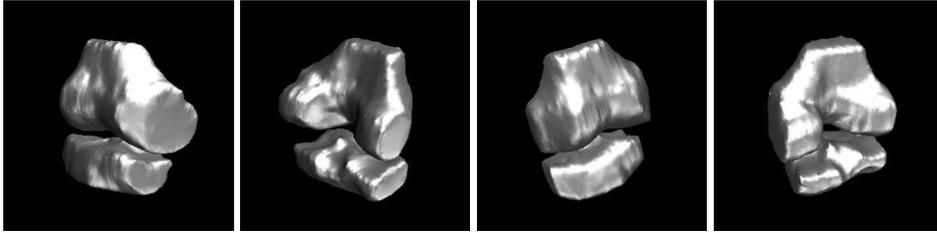


Fig. 3. Segmentations from 37 slices rendered together to show volumetric information.

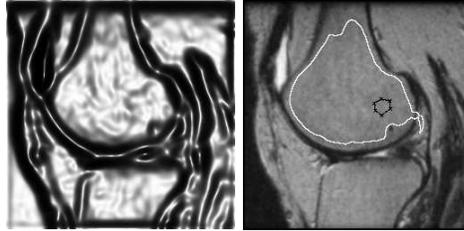


Fig. 4. Segmentation using only intensity gradients. (1) The energy function $g(|\nabla I|) = \frac{1}{1+|\nabla I|^2}$. (2) Curve flow showing leaking into neighboring structures, convergence not yet detected.

values of g) near the boundary often caused the contour to leak into other areas of the image. Figure 4 shows an example run using this energy function. The snake was stopped after a fixed number of iterations; it did not yet detect convergence. However, it had already leaked beyond the boundary of the femur, and subsequent iterations would cause it to leak much more. No compensation for smoothing was applied, but there is less of a need than in the two-valued case. Note again that the use of a balloon force that depends on g has changed the partial differential equations from the derivation of geodesic snakes, so the existence and uniqueness of a solution may not be assured as in the original formulation [4].

4.3 Comparison to Variance Only

Variance was also tested alone using the energy function $g(|\nabla S|) = \frac{1}{1+|\nabla S|^2}$, where S is the variance image. It attained much better results than intensity alone, rivaling the two-valued approach. It almost always converged to a reasonable boundary. Compared to the two-valued approach, it appeared slightly more likely to leak through a region boundary but was better able to handle large intensity variation within the bone region. Three experiments in which its results differed from those of the two-valued algorithm are shown in Figure 5. Again, the inner boundary is the result before compensation for windowing. The first experiment was stopped after some number of iterations: convergence was not detected. The second segmentation is better than that achieved in the two-valued

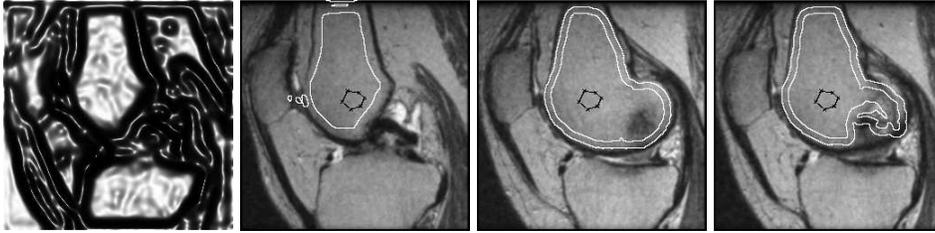


Fig. 5. Segmentation using only variance component of image. (1) The energy function $g(|\nabla S|) = \frac{1}{1+|\nabla S|^2}$, where S is the variance image. (2) Curve flow showing leaking into neighboring structures, convergence not yet detected (and stretching by k not applied). (3) A difficult segmentation. (4) A third segmentation where the result differs from the two-valued segmentation.

case for the same image and initial contour (Figure 2). Whether or not the third segmentation is better than the corresponding two-valued segmentation depends on what a radiologist would consider ground truth for that particular image, as the boundary is not clear. Overall the results for variance alone were good, and one could use this approach for segmentation but may need to be more careful when choosing parameters settings.

5 Future Work

One direction for future work is to extend the system to operate on 3D images directly, so that a surface is evolved instead of curves on individual slices. The geodesic snakes method was generalized to volumetric data in [5], and the vector-valued extension will hold in 3D as well. In that case, the data would be a vector-valued volume where one would consider mappings from R^3 to R^m where m is the number of (3D) texture features used.

A second direction is the incorporation of other texture measures which may be necessary for other applications or may improve the results for this application. Such measures include orientational filters and gray-level co-occurrence [8]. We also plan to consider textures across multiple scales to be able to handle a wide variety of textures and scenarios.

Acknowledgments

The authors thank Polina Golland for the program used to render the 2D slices as a volume. Also, thanks to Tina Kapur for discussions on this project.

References

1. R. Adams and L. Bischof. Seeded region growing. *IEEE Trans. Patt. Analysis and Mach. Intell.*, 16(6):641–647, 1994.

2. J. R. Beveridge, J. Griffith, R. R. Kohler, A. R. Hanson, and E. M. Riseman. Segmenting images using localized histograms and region merging. *Int'l Journal Comp. Vision*, 2:311–347, 1989.
3. V. Caselles, F. Catte, T. Coll, and F. Dibos. A geometric model for active contours. *Numerische Mathematik*, 66:1–31, 1993.
4. V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *Int'l Journal Comp. Vision*, 22(1):61–79, 1997.
5. V. Caselles, R. Kimmel, G. Sapiro, and C. Sbert. Minimal surfaces: A three dimensional segmentation approach. Technical Report 973, Technion, Dept. EE, June 1995.
6. L. D. Cohen. On active contour models and balloons. *CVGIP: Image Understanding*, 53(2):211–218, 1991.
7. J. De Bonet and P. Viola. Structure driven image database retrieval. In *Neural Information Processing Systems*, 1997.
8. R. M. Haralick and L. G. Shapiro. *Computer and Robot Vision*. Addison-Wesley Publishing Company, Inc., 1992.
9. T. Kapur. Segmentation of brain tissue from magnetic resonance images. Technical Report 1566, MIT Artificial Intelligence Laboratory, Jan. 1995.
10. T. Kapur, P. Beardsley, and S. Gibson. Segmentation of bone in clinical mri. Technical report, MERL, A Mitsubishi Electric Research Company, Aug. 1997.
11. T. Kapur, P. Beardsley, S. Gibson, W. Grimson, and W. Wells. Model-based segmentation of clinical knee MRI. In *Proc. IEEE Int'l Workshop on Model-Based 3D Image Analysis*, pages 97–106, 1998.
12. M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int'l Journal Comp. Vision*, 1(4):321–331, 1988.
13. A. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Gradient flows and geometric active contour models. In *Proc. IEEE Int'l Conf. Comp. Vision*, pages 810–815, 1995.
14. T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: a survey. *Medical Image Analysis*, 1(2):91–108, 1996.
15. S. Osher and J. Sethian. Gradient flows and geometric active contour models. *Journal of Computational Physics*, 79:12–49, 1988.
16. R. Ronfard. Region-based strategies for active contour models. *Int'l Journal Comp. Vision*, 13(2):229–251, 1994.
17. G. Sapiro. Vector-valued active contours. In *Proc. IEEE Conf. Comp. Vision and Patt. Recog.*, pages 680–685, 1996.
18. L. Staib and J. Duncan. Boundary finding with parametrically deformable contour models. *IEEE Trans. Patt. Analysis and Mach. Intell.*, 14(11), 1992.
19. M. Wolf, P. Weierich, and H. Niemann. Automatic segmentation and 3d-registration of the femoral bone in mr-images of the knee. *Pattern Recognition and Image Analysis: Advances in Mathematical Theory and Application, Russian Academy of Sciences*, 7, 1997.
20. A. Yezzi, S. Kichenassamy, A. Kumar, P. Olver, and A. Tannenbaum. A geometric snake model for segmentation of medical imagery. *IEEE Trans. Medical Imaging*, 16(2):199–209, 1997.
21. S. C. Zhu and D. Mumford. Prior learning and gibbs reaction-diffusion. *IEEE Trans. Patt. Analysis and Mach. Intell.*, 19(11), Nov. 1997.
22. S. C. Zhu, Y. Wu, and D. Mumford. Frame: Filters, random fields, and minimax entropy. In *Proc. IEEE Conf. Comp. Vision and Patt. Recog.*, pages 686–693, 1996.
23. S. C. Zhu and A. Yuille. Region competition. *IEEE Trans. Patt. Analysis and Mach. Intell.*, 18(9), 1996.