A Curve Evolution Approach to Medical Image Magnification via the Mumford-Shah Functional*

Andy Tsai[1], Anthony Yezzi Jr.[2], and Alan S. Willsky[1]

[1] Department of Electrical Engineering and Computer Science,
Massachusetts Institute of Technology,
Cambridge, MA, USA
{atsai, willsky}@mit.edu
[2] School of Electrical and Computer Engineering,
Georgia Institute of Technology,
Atlanta, GA, USA
Anthony.Yezzi@ece.gatech.edu

Abstract. In this paper, we introduce a curve evolution approach to image magnification based on a generalization of the Mumford-Shah functional. This work is a natural extension of the curve evolution implementation of the Mumford-Shah functional presented by the authors in previous work. In particular, by considering the image magnification problem as a structured case of the missing data problem, we generalize the data fidelity term of the original Mumford-Shah energy functional by incorporating a spatially varying penalty to accommodate those pixels with missing measurements. This generalization leads us to a PDE-based approach for simultaneous image magnification, segmentation, and smoothing, thereby extending the traditional applications of the Mumford-Shah functional which only considers simultaneous segmentation and smoothing. This novel approach for image magnification is more global and much less susceptible to blurring or blockiness artifacts as compared to other more traditional magnification techniques, and has the additional attractive denoising capability.

1 Introduction

Image magnification is often required to aid in the display, manipulation, and analysis of medical images. It also plays a critical role in computer aided diagnosis and computer assisted surgery. Image magnification involves enlarging a small image to several times its size and often requires some sort of an interpolation scheme. The most straightforward approach for image enlargement is to use a zero-order interpolation technique, commonly known as replication, which may cause the resulting image to appear blocky [2]. Classical enlargement techniques such as bilinear or bicubic interpolation schemes tend to cause blurring across the edges when applied indiscriminantly to the image [2]. More sophisticated

* This work was supported by ONR grant N00014-00-1-0089, by AFOSR grant F49620-98-1-0349, and by subcontract GC123919NGD from Boston Univ. under the AFOSR Multidisciplinary Research Program on Reduced Signature Target Recognition.

S.L. Delp, A.M. DiGioia, and B. Jaramaz (Eds.): MICCAI 2000, LNCS 1935, pp. 246-255, 2000.

schemes may locate the edges first with local filters prior to interpolation so as to avoid the blurring artifacts [1,6]. Three important shortcomings are evident in these types of algorithms. One, the interpolation schemes used for magnification are local since they only utilize data values from neighboring pixels. These types of interpolation scheme become problematic when the observed image is noisy. Two, edge detection schemes employed prior to interpolation often only make use of local information (which are very susceptible to noise artifacts) and cannot guarantee continuous closed edge contours. Three, it is unclear in what order the three operations (smoothing, edge detection, and interpolation) should be performed for image magnification since they are not commutative. Our approach for image magnification addresses the first deficiency by using a PDE-based model for interpolating the data which incorporates the use of all data values within each homogeneous region, not just neighboring pixels, to determine the interpolant. As a result, this interpolation scheme is much more robust to noise. The second deficiency is addressed by the use of our active contour model for boundary detection which is more global in nature than local filters (and therefore not as sensitive to noise) and is curve-based (hence providing a continuous closed edge contour). The third deficiency is addressed by our use of the Mumford-Shah model which provides, in a single framework, a tight coupling for *simultaneous* image segmentation, denoising, and magnification. In this manner, the ordering of the different operations is no longer an issue.

Our approach to image magnification is based on the Mumford-Shah active contour model presented earlier by the authors in [8]. In that work, we approached the classic Mumford-Shah problem from a curve evolution perspective. In particular, by viewing an active contour as the set of discontinuities in the Mumford-Shah problem, we used the corresponding functional to determine gradient descent evolution equations to deform the active contour. Each gradient step involved solving an optimal estimation problem for the data within each region. The resulting active contour model offered a tractable implementation of the original Mumford-Shah model to simultaneously segment and smoothly reconstruct the data within a given image in a coupled manner. In this paper, we extend the application of this Mumford-Shah active contour model to one that also handles image magnification. By considering the image magnification problem as a structured case of the missing data problem, we generalize the data fidelity term of the original Mumford-Shah functional by substituting a spatially varying penalty for the traditional constant one (to accommodate those pixels with missing measurements). This generalization leads to a novel approach for simultaneous image magnification, segmentation, and smoothing, thereby providing a new application of the Mumford-Shah functional.

This paper is organized as follows. Section 2 describes the curve evolution implementation of the Mumford-Shah functional proposed in [8] for simultaneous image segmentation and smoothing. In Section 3, we describe the generalization of the Mumford-Shah active contour model for image magnification. Then Section 4 provides some experimental results to illustrate this magnification technique. Finally, we conclude this paper in Section 5 with a summary.

2 The Mumford-Shah formulation as a curve evolution problem

The point of reference for this paper is the Mumford-Shah functional

$$E(\mathbf{f},\vec{C}) = \beta \iint_{\Omega} (\mathbf{f}-\mathbf{g})^2 \; dA + \alpha \iint_{\Omega\setminus\vec{C}} |\nabla \mathbf{f}|^2 \; dA \; + \; \gamma \oint_{\vec{C}} ds \tag{1}$$

in which $\vec{C}$ denotes the smooth, closed segmenting curve, $\mathbf{g}$ denotes the observed data, $\mathbf{f}$ denotes the piecewise smooth approximation to $\mathbf{g}$ with discontinuities only along $\vec{C}$, and Ω denotes the image domain [4]. The parameters α, β, and γ are positive real scalars which control the competition between the various terms above and determine the "scale" of the segmentation and smoothing. The Mumford-Shah problem is to minimize $E(\mathbf{f},\vec{C})$ over admissible $\mathbf{f}$ and $\vec{C}$. In our approach, we constrain the set of discontinuities in the Mumford-Shah problem to correspond to evolving sets of curves, enabling us to tackle the problem via a curve-evolution-based approach.

2.1 Optimal image estimation

Any arbitrary closed curve $\vec{C}$ partitions the domain Ω of the image into two regions R and R^c, corresponding to the interior and exterior of curve $\vec{C}$ respectively. If we fix this partitioning, minimizing (1) corresponds to minimizing

$$\begin{aligned} E_{\vec{C}}(\mathbf{f}_R,\mathbf{f}_{R^c}) = {} & \beta \iint_{R} (\mathbf{f}_R-\mathbf{g})^2 \, dA + \alpha \iint_{R} |\nabla \mathbf{f}_R|^2 \, dA \\ & + \beta \iint_{R^c} (\mathbf{f}_{R^c}-\mathbf{g})^2 \, dA + \alpha \iint_{R^c} |\nabla \mathbf{f}_{R^c}|^2 \, dA. \end{aligned} \tag{2}$$

The estimates $\hat{\mathbf{f}}_R$ and $\hat{\mathbf{f}}_{R^c}$ that minimize (2) satisfy (decoupled) PDE's which can be obtained using standard variational methods [4]. This method gives the following damped Poisson equation with Neumann boundary condition for $\hat{\mathbf{f}}_R$:

$$\hat{\mathbf{f}}_R - \frac{\alpha}{\beta} \nabla^2 \hat{\mathbf{f}}_R = \mathbf{g} \qquad \text{on } R \tag{3a}$$

$$\frac{\partial \hat{\mathbf{f}}_R}{\partial \vec{\mathcal{N}}} = 0 \qquad \text{on } \vec{C}. \tag{3b}$$

In a similar fashion, $\hat{\mathbf{f}}_{R^c}$ is given as the solution to

$$\hat{\mathbf{f}}_{R^c} - \frac{\alpha}{\beta} \nabla^2 \hat{\mathbf{f}}_{R^c} = \mathbf{g} \qquad \text{on } R^c \tag{4a}$$

$$\frac{\partial \hat{\mathbf{f}}_{R^c}}{\partial \vec{\mathcal{N}}} = 0 \qquad \text{on } \vec{C}. \tag{4b}$$

Conjugate gradient (CG) method is used to solve the above *estimation PDE's*.

2.2 Gradient flows that minimize the Mumford-Shah functional

With the ability to calculate $\hat{\mathbf{f}}_R$ and $\hat{\mathbf{f}}_{R^c}$ for any given $\vec{C}$, we now wish to derive a curve evolution for $\vec{C}$ that minimizes (1). That is, as a function of $\vec{C}$, we wish to find $\vec{C}_t$ that minimizes

$$E_{\hat{\mathbf{f}}_R,\hat{\mathbf{f}}_{R^c}}(\vec{C}) = \beta\iint_R (\hat{\mathbf{f}}_R - \mathbf{g})^2 \; dA + \alpha\iint_R |\nabla\hat{\mathbf{f}}_R|^2 \, dA$$
$$+\beta\iint_{R^c} (\hat{\mathbf{f}}_{R^c} - \mathbf{g})^2 \; dA + \alpha\iint_{R^c} |\nabla\hat{\mathbf{f}}_{R^c}|^2 \, dA + \gamma\oint_{\vec{C}} ds. \tag{5}$$

The curve evolution that minimizes (5) is given by (see [8] for a derivation)

$$\vec{C}_t = \frac{\alpha}{2}\left(|\nabla\hat{\mathbf{f}}_{R^c}|^2 - |\nabla\hat{\mathbf{f}}_R|^2\right)\vec{\mathcal{N}} + \frac{\beta}{2}\left((\mathbf{g} - \hat{\mathbf{f}}_{R^c})^2 - (\mathbf{g} - \hat{\mathbf{f}}_R)^2\right)\vec{\mathcal{N}} - \gamma\kappa\vec{\mathcal{N}}. \tag{6}$$

This *Mumford-Shah flow* is implemented via the level set method [5,7] which offers a natural and numerically reliable implementation of these solutions within a context that handles topological changes in the interface without any additional effort.

2.3 Implementation of the Mumford-Shah active contour model

The algorithm described above requires solving two PDE's (equations (3) and (4)) at every evolution step of the curve making it computationally expensive and impractical. We propose an approximate gradient descent approach to calculate $\hat{\mathbf{f}}_R$, $\hat{\mathbf{f}}_{R^c}$, and $\vec{C}$ to alleviate some of the computational burdens. This approach consists of alternating between these two steps:

- Fix $\hat{\mathbf{f}}_R$ and $\hat{\mathbf{f}}_{R^c}$, and take several gradient descent curve evolution steps to move the curve $\vec{C}$.
- Fix $\vec{C}$, and perform just a few iterations of the conjugate gradients method–without taking it to convergence–to obtain a *rough* estimate of $\mathbf{f}_R$ and $\mathbf{f}_{R^c}$.

Only a rough estimate of $\mathbf{f}_R$ and $\mathbf{f}_{R^c}$ is required to direct the curve to move in the descent direction. The idea is to make the algorithm faster by reducing the number of times $\hat{\mathbf{f}}_R$ and $\hat{\mathbf{f}}_{R^c}$ are calculated and also the amount of time required to calculate each of them. CG method is carried to convergence in the last iteration to obtain an accurate estimate of $\mathbf{f}_R$ and $\mathbf{f}_{R^c}$.

Next, we propose an implementation of our active contour model, building on the preceding modification, to enable our model to handle images with multiple junctions without resorting to more sophisticated level set techniques [3]. Given an image, we apply our Mumford-Shah active contour model for segmentation and smoothing. After segmentation, if any of the resulting subregions require additional segmentation, apply our algorithm again, but this time, restricting the algorithm to operate only in that particular subregion. This approach has the

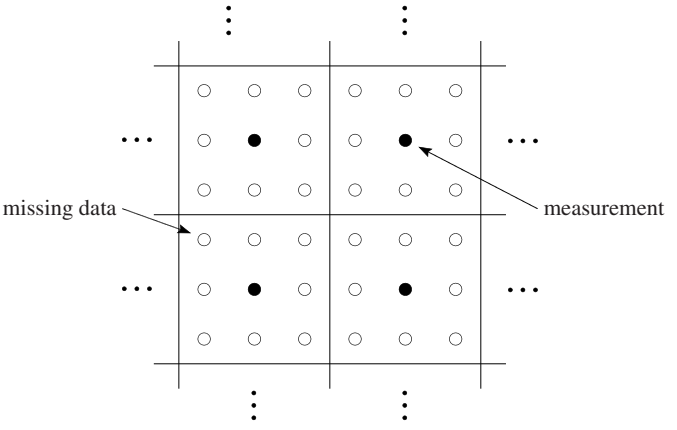

Fig. 1. Diagram showing the locations of missing data in relation to the measurements in our image magnification technique.

natural notion of starting with a crude segmentation and refining the segmentation by telescoping down to the different subregions in order to capture finer and finer details in the image. The attractive feature associated with this hierarchical implementation is that it allows us to handle images with multiple junctions by employing multiple curves to represent such junctions. This implementation also affords us better control as to what details we desire and what objects we would like to capture, in the segmentation and smoothing of our image.

3 Generalization of the Mumford-Shah active contour model for image magnification

Image magnification capability is weaved into the Mumford-Shah active contour model by considering the image magnification problem as a structured case of the missing data problem. Specifically, consider a new grid with three times as many pixels in each direction and assign the value of the original image to the "center" pixel in each 3×3 block on the grid and treat the remaining pixels as missing data points (see Figure 1). From an estimation-theoretic standpoint, we can view these "center" pixels as sparse measurements on a much larger image domain. We then employ the generalized Mumford-Shah curve evolution procedure (described below) to interpolate to this finer grid, using the curve evolution portion of this procedure to partition the domain of the magnified image into different homogeneous subregions so as to provide smooth interpolations where appropriate without blurring across regions of high contrast.

The Mumford-Shah model handles image magnification through the parameter β. In the standard Mumford-Shah formulation (1), β is a constant scalar parameter reflecting our confidence in the measurements. For image magnification in which there are pixels with missing measurements, we replace the constant parameter β by a spatially varying function $\boldsymbol{\beta}$ whose value at each pixel

is inversely proportional to the variance of the measured noise at that pixel. For example, in the situation where the data at pixel (x_o, y_o) is missing, we consider the variance of the data at that pixel as being infinite and accordingly set $\beta(x_o, y_o) = 0$. By introducing this spatially varying β, equation (1) becomes:

$$E(\mathbf{f}, \vec{C}) = \iint_{\Omega} \beta(\mathbf{f} - \mathbf{g})^2 \, dA + \alpha \iint_{\Omega \setminus \vec{C}} |\nabla \mathbf{f}|^2 \, dA + \gamma \oint_{\vec{C}} ds. \tag{7}$$

The gradient flow that minimizes the generalized Mumford-Shah functional, shown in (7), is given by

$$\vec{C}_t = \frac{\alpha}{2} \left(|\nabla \hat{\mathbf{f}}_{R^c}|^2 - |\nabla \hat{\mathbf{f}}_R|^2 \right) \vec{\mathcal{N}} + \frac{\beta}{2} \left((\mathbf{g} - \hat{\mathbf{f}}_{R^c})^2 - (\mathbf{g} - \hat{\mathbf{f}}_R)^2 \right) \vec{\mathcal{N}} - \gamma \kappa \vec{\mathcal{N}} \tag{8}$$

where the optimal estimates $\hat{\mathbf{f}}_R$ and $\hat{\mathbf{f}}_{R^c}$ of (8) satisfy

$$\begin{aligned} \beta \hat{\mathbf{f}}_R - \alpha \nabla^2 \hat{\mathbf{f}}_R &= \beta \mathbf{g} && \text{on } R \\ \frac{\partial \hat{\mathbf{f}}_R}{\partial \vec{\mathcal{N}}} &= 0 && \text{on } \vec{C} \end{aligned}$$

and

$$\begin{aligned} \beta \hat{\mathbf{f}}_{R^c} - \alpha \nabla^2 \hat{\mathbf{f}}_{R^c} &= \beta \mathbf{g} && \text{on } R^c \\ \frac{\partial \hat{\mathbf{f}}_{R^c}}{\partial \vec{\mathcal{N}}} &= 0 && \text{on } \vec{C}. \end{aligned}$$

Over each region of missing data D, the estimation equation reduces to the Laplace equation with the same Neumann boundary condition:

$$\begin{aligned} \nabla^2 \hat{\mathbf{f}}_D &= 0 && \text{on } D \\ \frac{\partial \hat{\mathbf{f}}_D}{\partial \vec{\mathcal{N}}} &= 0 && \text{on } \vec{C}. \end{aligned}$$

As solutions to the Laplace equation, the estimates obtained over any such missing data regions not containing part of $\vec{C}$ take the form of harmonic functions. As such, we can infer much about the smooth nature of these interpolated estimates as they are subject to both a maximum (and minimum) principle as well as the mean value property. However if the curve $\vec{C}$ intersects D, no such smoothing occurs across this boundary, allowing interpolation to be guided by the segmentation defined by $\vec{C}$.

4 Examples

We illustrate our magnification technique by applying it to a noisy microscopic image of red blood cells (Figure 2(a)). The initializing curve is shown in Figure 2(a). Figure 2(b) shows an intermediate step of the Mumford-Shah model.

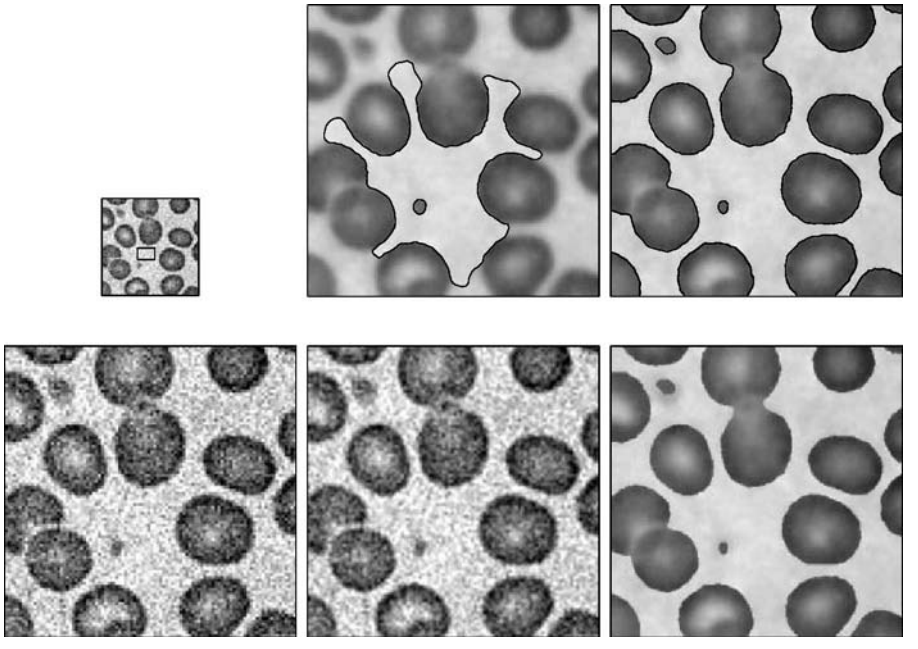

Fig. 2. Three-fold magnification of a noisy 100×100 microscopic image of blood cells. (a) Original image with location of initializing curve. (b) An intermediate step in the magnification and segmentation process. (c) Magnified image with the segmenting curve. (d) Magnified image based on zero-order interpolation. (e) Magnified image based on bilinear interpolation. (f) Magnified image based on the Mumford-Shah model.

Notice the curve automatically proceeds in the outward direction without the aid of any inflationary forces to direct this movement. This is one of the many attractive properties associated with the Mumford-Shah flow (see [8] for more details). The final segmenting curve and the smooth enlarged reconstruction of the image is shown in Figure 2(c). For comparison, in Figure 2(d), we show the blocky and noisy magnification of the original image based on zero-order interpolation. In Figure 2(e), we show a slightly smoother but blurrier magnification of the original image based on bilinear interpolation. In Figure 2(f), we again show the magnified image based on our Mumford-Shah approach but without the segmenting curve. The smooth boundaries of the magnified image are the direct result of the minimum length prior placed on the segmenting curve while the smooth spatially varying structures within each homogeneous regions of the observed image are successfully captured by our PDE-based model for interpolation. It is evident that our magnification approach is better than other conventional image magnification techniques by avoiding many of the processing artifacts such as blockiness and blurring while at the same time, denoising the image. Also, notice in this figure that the starting contour automatically splits into many contours in order to segment this image. This provides a compelling

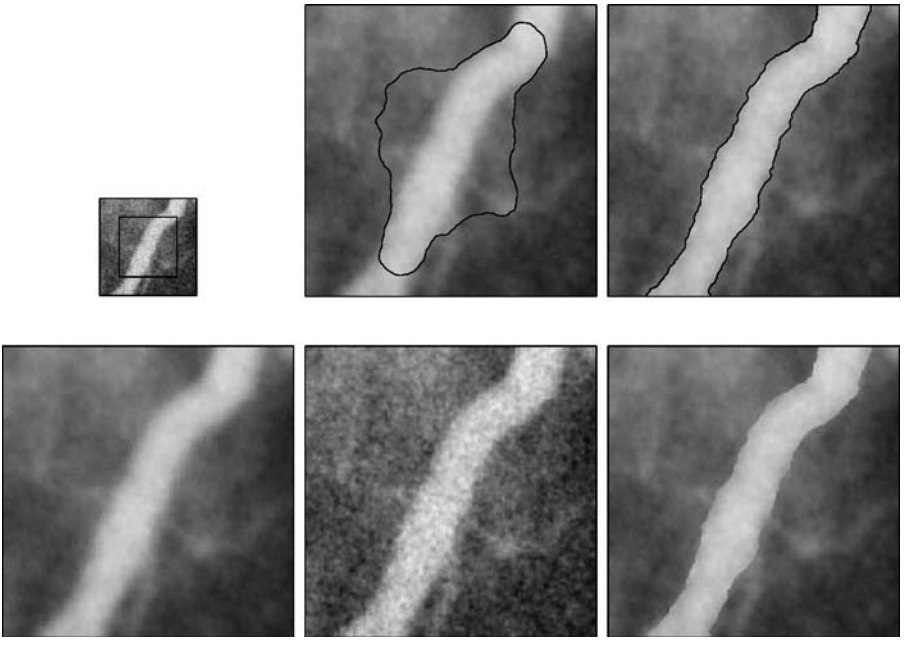

Fig. 3. Three-fold magnification of a noisy 110×110 angiogram. (a) Original image with location of initializing curve. (b) An intermediate step in the magnification and segmentation process. (c) Magnified image with the segmenting curve. (d) Magnified image obtained by isotropic smoothing then replication. (e) Magnified image obtained by replication then isotropic smoothing. (f) Magnified image based on the Mumford-Shah model.

demonstration of the topological transitions allowed by the Mumford-Shah active contour model's level set implementation [5, 7].

We next illustrate our magnification technique by applying it to a noisy angiogram (Figure 3(a)). The initializing curve is shown in Figure 3(a). Figure 2(b) shows an intermediate step of the Mumford-Shah model. Notice the bi-directional flow of the curve as it automatically proceeds toward the boundaries of the image. The final segmenting curve and smooth enlarged reconstruction of the image is shown in Figure 3(c). For comparison, in Figure 3(d), we show a magnified angiogram obtained by first applying isotropic smoothing (i.e. the same smoothing used in the Mumford-Shah framework but applied over the image as one single region) then replication. This image is blurry because the edges of the image were destroyed during the initial smoothing step. In Figure 3(e), we show a magnified angiogram obtained by first replicating the image then applying isotropic smoothing. Notice the magnified image is still noisy because the noise components within the original image have been exaggerated by the zero-order interpolation scheme. In Figure 3(f), we again show the magnified image based on our Mumford-Shah model but displayed without the segmenting curve.

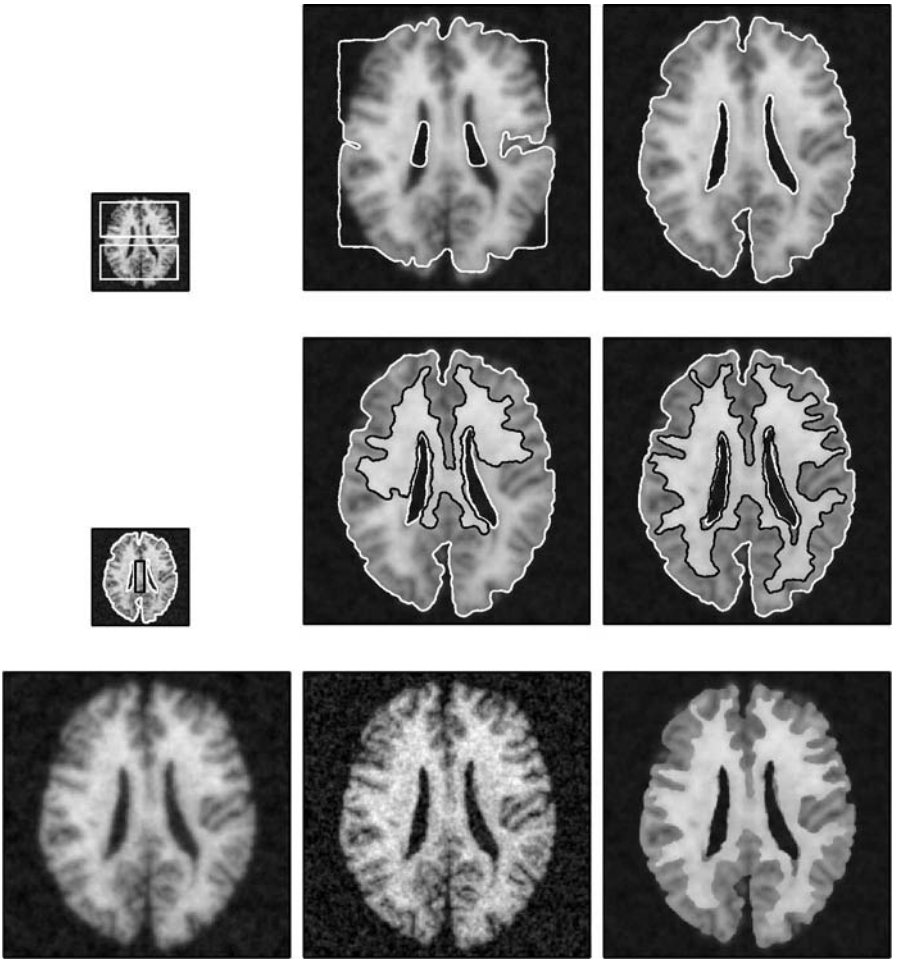

Fig. 4. Three-fold magnification of a noisy 225×225 axial brain pathology section. (a) Original image with the initializing contour for the first curve. (b) An intermediate step in the magnification and segmentation of the image based on the first curve. (c) Magnified image using one segmenting curve which is also displayed. (d) Original image with the segmentation result of the first curve and the initializing contour for the second curve. (e) An intermediate step in the magnification and segmentation of the image based on the second curve. The segmentation result of the first curve is also displayed. (f) Magnified image with the segmentation results from both curves. (g) Magnified image obtained by isotropic smoothing then replication. (h) Magnified image obtained by replication then isotropic smoothing. (i) Magnified image based on the Mumford-Shah model.

We use the noisy brain pathology image shown in Figure 4(a) to demonstrate the hierarchical implementation of our active contour model for image magnification. The segmentation and enlarged reconstruction of the image, shown in Figure 4(c), is obtained using a single curve (which is initialized as shown in Figure 4(a)) and based on the approximate gradient descent approach described earlier. The blurring across the boundary of the white and the gray matter shown in the reconstruction is due to the erroneous implication of this coarse segmentation, namely that the inside of the brain is one region over which smoothing is performed. To provide better details within the brain, we again applied our technique to the interior region of the brain. A second curve, initialized as shown in Figure 4(d), is used in conjunction with the first curve to obtain the segmentation and enlarged reconstruction of the image shown in Figure 4(f). For comparison, in Figure 4(g), we show the magnified image obtained by first isotropically smoothing the original noisy image then magnifying it using zero-order hold. We show in Figure 4(h) the magnified image obtained by first magnifying the original noisy image using zero-order hold then smoothing it isotropically. Finally, in Figure 4(i), we show the magnification results based on our approach without the segmenting curve.

5 Conclusion

The magnification technique we introduced in this paper constitutes a more global approach to interpolating magnified data than traditional bilinear or bicubic interpolation schemes, while still maintaining sharp transitions along region boundaries. The technique is also much less susceptible to blurring artifacts as compared to other more traditional techniques, and has the additional attractive denoising capability. In addition, the curve length penalty in our Mumford-Shah based flow tends to prevent the blocky appearance of object boundaries which is a symptom of replication-based schemes.

References

1. J. Allebach and W. Wong, "Edge directed interpolation," *IEEE International Conference on Image Processing*, vol. 3, pp. 707-711, 1996.
2. J. Lim, *Two-dimensional Signal and Image Processing*, Prentice Hall, 1992.
3. B. Merriman, J. Bence, and S. Osher, "Motion of multiple junctions: A level set approach," *Journal of Computational Physics*, vol. 112, no. 2, pp. 334–363, 1994.
4. D. Mumford and J. Shah, "Optimal approximations by piecewise smooth functions and associated variational problems," *Comm. Pure and Appl. Math.*, vol. 42, 1989.
5. S. Osher and J. Sethian, "Fronts propagation with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations," *J. of Comp. Physics*, vol. 79, pp. 12–49, 1988.
6. K. Ratakonda and N. Ahuja, "POC based adaptive image magnification," *ICIP*, 1998.
7. J. Sethian, *Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision, and Material Science*, Cambridge University Press, 1996.
8. A. Tsai, A. Yezzi, and A. Willsky, "A curve evolution approach to smoothing and segmentation using the Mumford-Shah functional," *CVPR*, 2000.