

Tunneling descent for m.a.p. active contours in ultrasound segmentation

Zhong Tao^a, Hemant D. Tagare^{b,c,*}

^a Department of Electrical Engineering, Yale University, New Haven, CT 06520, United States

^b Department of Diagnostic Radiology, Yale University, New Haven, CT 06520, United States

^c Department of Biomedical Engineering, Yale University, New Haven, CT 06520, United States

Received 17 October 2005; received in revised form 22 March 2006; accepted 21 November 2006

Abstract

Active contours that evolve in ultrasound images under gradient descent are often trapped in spurious local minima. This paper presents an evolution strategy called tunneling descent, which is capable of escaping from such minima. The key idea is to evolve the contour by a sequence of constrained minimizations that move the contour in to, and out of, local minima. This strategy is an extension of classical gradient descent.

Because tunneling descent does not terminate at a local minima an explicit stopping rule is required. Model-based and model-free stopping rules are presented and formulae for choosing the stopping threshold are given.

The algorithm is used to segment the endocardium in 44 short axis cardiac ultrasound images. The energy function of the active contour is derived from a m.a.p. formulation. All segmentations are achieved without tweaking either the energy function or numerical parameters. Experimental evaluation of the segmentations show that the algorithm overcomes multiple local minima to find the endocardium. The accuracy of the algorithm is comparable to that of manual segmentations and significantly better than classical gradient descent active contours. The sensitivity of the segmentation to initialization is also evaluated and it is shown that segmentations from quite different initializations are close to each other. Finally, some limitations of the m.a.p. formulation are discussed.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Ultrasound segmentation; M.A.P. segmentation; Active contours; Optimization

1. Introduction

Segmentation of B-mode cardiac ultrasound images using active contours is a challenging task. There are two reasons for this: First, ultrasound images contain speckle. Speckle creates spurious local minima in the segmentation energy function. An active contour that is initialized far away from the boundary is often trapped in spurious minima and gives obviously wrong segmentations.

Second, ultrasound images sometimes exhibit dropout. These are regions in the image where the signal is missing or greatly attenuated, and which appear dark. Dropout can create holes through which active contours can leak out.

In this paper, we address the first problem – we present a deterministic evolution strategy that lets active contours escape from spurious local minima. We call it *tunneling descent*. The key idea is to pose the contour evolution problem as a *sequence* of constrained minimizations instead of a single unconstrained minimization. The constraints guarantee that the algorithm can escape from any spurious local minima, no matter how deep or how steep its walls. Tunneling descent is especially useful in *maximum-a-posterior (m.a.p.) active contour* segmentation of the endocardium in cardiac ultrasound images.

* Corresponding author. Address: Department of Diagnostic Radiology, Yale University, New Haven, CT 06520, United States. Tel.: +1 203 737 4271.

E-mail addresses: zhong.tao@yale.edu (Z. Tao), hemant.tagare@yale.edu (H.D. Tagare).

Tunneling descent offers other advantages too. Other than step size (which can be fixed a priori), it does not require any manual interaction – all segmentations reported in this paper are achieved without any parameter tweaking. Moreover, the segmentations offered by tunneling descent are very robust to initialization. Initializing the active contour in very different positions gives essentially the same segmentation. All of this is borne out by the experiments in Section 8.

A level set version of tunneling descent was reported recently (Tao and Tagare, 2005). Finally, we note that we are only concerned with B-mode images in this paper, and henceforth we will refer to them simply as “ultrasound images.”

1.1. Organization

This paper is organized as follows. Section 2 is a brief review of the literature on active contours and ultrasound segmentation. Section 3 introduces m.a.p. active contours and illustrates the problem of local minima in ultrasound images. It also introduces the terminology and notation for the rest of the paper. Sections 4–6 introduce tunneling descent. Section 7 contains the log-likelihood function for ultrasound segmentation, and Section 8 describes experimental results. Section 9 discusses some of the limitations of the likelihood function. Section 10 concludes the paper. Appendix A contains implementation and numerical details.

2. Previous work

The literature on active contours is vast and we only mention work that has a direct bearing on the local minima problem. Comprehensive general reviews of active contours and deformable models are available in Xu et al. (2000) and McInerney and Terzopoulos (1996).

Previous active contour strategies for escaping from local minima fall under two classes: *multi-resolution minimization* and *addition of extra forces*. In the multi-resolution approach, the energy minimum is sought by first smoothing the image, finding a local minimum in the smoothed image, and then tracking the local minimum back to the original image (Kass et al., 1987; Terzopoulos et al., 1988; Muzzolini et al., 1993). Smoothing eliminates spurious local minima by merging them with nearby local maxima.

In the second approach, the active contour evolution is modified by adding extra forces. For example, *balloons* overcome local minima by adding a constant expansion force to the active contour (Cohen and Cohen, 1991). The idea is that the constant expansion force drags the active contour out of small spurious local minima. Other forces have also been proposed for extending the attraction range (Cohen and Cohen, 1993) or for deforming the contour into boundary concavities (Xu and Prince, 1998).

Multi-resolution minimization and extra forces work well with many images, but have limitations that preclude their use as a completely general purpose strategy. Both approaches have free parameters which have to be carefully set by the user. In the multi-resolution approach, the right amount of smoothing is critical. Over-smoothing can annihilate not just the spurious minima, but also the desired minimum. And under-smoothing can leave some spurious local minima intact, ready once more to trap gradient descent. Similarly, an excessive balloon force can force the active contour out of the desired minimum, while a weak balloon force can prevent it from escaping a spurious minimum. Guessing just the right amount of smoothing or force is difficult, even for an expert. Further, the right amount can vary from image to image, and has to be manually adjusted.

Thus there is a need for an alternate minimization strategy, especially one that does not require frequent adjustment of scales and forces. This is our motivation for developing tunneling descent.

2.1. Ultrasound statistics

An ultrasound image is a record of the backscatter from a propagating ultrasound wave. Random scatterers in the acoustic medium give rise to *speckle* in the ultrasound image and fully organized interfaces give rise to *specularity*. Speckle is manifest as a spatial random process in the image. It contributes to the “grainy” look of ultrasound images.

There has been much theoretical work regarding the statistics of speckle. These studies show that the distribution of speckle is significantly non-Gaussian. Goodman (1975) and Butkhardt (1978) showed that the first-order distribution of speckle is Rayleigh if the scatterers are uniformly distributed in the acoustic medium. Wagner et al. (1983) showed that the speckle distribution is Rician for quasi-periodic scatterers. Jakeman and Tough (1987) and Weng et al. (1991), suggest the K distribution for speckle statistics, while Shankar (2000) proposed the Nakagami distribution. Abyratne et al. (1996) and CramBlitt and Parker (1999) proposed more complex models for structured scatterers. A comprehensive overview of statistical speckle models is published by Insana et al. (2000). These models give the statistics of speckle at the transducer. Using these models in image processing tasks is not always appropriate because the transducer signal undergoes complex processing before it is displayed as an image.

An alternative is to create empirical models for image statistics from real-world ultrasound images, e.g. Zimmer et al. (2000). The authors of this paper undertook such a study in Tao et al. (2002) whose results are summarized in Section 7. This empirical model is used to in Section 8 to segment B-mode short-axis cardiac ultrasound images.

2.2. Ultrasound segmentation

Because speckle makes ultrasound segmentation difficult, a number of researchers have developed speckle suppression techniques by using, for instance, order statistical filters (Belohlacek and Greenleaf, 1997), adaptive filters (Bamber and Daft, 1986; Karaman et al., 1995), and wavelets (Rakotomamonjy et al., 2000).

Classical computer-vision approaches, such as multiresolution texture analysis (Muzzolini et al., 1993) and edge and line detectors (Czerwinski et al., 1999) have also been proposed for ultrasound segmentation.

As mentioned before, active contour models perform poorly with ultrasound images. Many attempts have been made to improve their performance, e.g. multi-resolution schemes are used in Aston and Parker (1995), and point distribution, Active Shape, and Active Appearance Models in Parker et al. (1994), Hill and Taylor (1992), Bosch et al. (2002). Anatomical and other expert information has been integrated into ultrasound segmentation algorithms (Han et al. (1991)). In a series of publications (Geiser et al., 1998; Wilson et al., 2000; Geiser and Wilson, 1996), an elaborate segmentation algorithm that uses knowledge of anatomical and image features is reported. Multi-frame spatio-temporal approaches to ultrasound segmentation are reported in Mulet-Parada and Noble (1998), Guofang et al. (2002), Jacob et al. (2002), and Sanchez-Ortiz et al. (2002).

In contrast to the above approaches, we pursue an active contour maximum-a-posteriori estimation (m.a.p.) approach to ultrasound segmentation. The m.a.p. approach is appealing because the ultrasound signal is a random process and segmenting an ultrasound image can be naturally viewed as an estimation problem.

The literature on m.a.p. segmentation is vast. We mention two representative approaches, one using active contours (Staib and Duncan, 1992) and the other using Markov random fields (Winkler, 2006).

3. M.a.p. active contours

We now turn to describing m.a.p. active contours. We begin by fixing some notation and terminology.

We denote curves in the image by capital letters such as C, D, \dots . We assume that the curves are closed and arc-length parameterizable. \mathcal{E} denotes the set of all such curves

in the image. Any curve C has an *interior* region Ω_C (Fig. 1a). The complement of the interior, $\tilde{\Omega}_C$, is the *exterior* of the curve. We consistently use this notation, so that for a curve D , the interior and exterior regions are Ω_D and $\tilde{\Omega}_D$, for curve E they are Ω_E and $\tilde{\Omega}_E$, and so on.

3.1. Log-likelihood and energy

Suppose that the image has two regions separated by a boundary curve B (Fig. 1b). This is the boundary we wish to find. The gray levels in the inside and outside regions of the boundary come from different probability models. Let $l_1(I_\Sigma|\theta_1)$ be the log-likelihood that gray levels I_Σ of any region Σ in the image come from the region inside the boundary and let $l_2(I_\Sigma|\theta_2)$ be the log-likelihood that I_Σ come from the region outside the boundary. In these formulae, θ_1 and θ_2 are free parameters of the models. Then, the conditional log-likelihood of observing the image I assuming that C is the boundary is

$$L_{\text{cond}}(I|C, \theta_1, \theta_2) = l_1(I_{\Omega_C}|\theta_1) + l_2(I_{\tilde{\Omega}_C}|\theta_2).$$

If $L_{\text{prior}}(C, \theta_1, \theta_2)$ is the log-likelihood prior of the curve and the free parameters, then the posterior log-likelihood of C, θ_1, θ_2 is

$$L(C, \theta_1, \theta_2|I) = L_{\text{cond}}(I|C, \theta_1, \theta_2) + L_{\text{prior}}(C, \theta_1, \theta_2). \quad (1)$$

The *m.a.p. active contour* uses the negative posterior log-likelihood as its energy, i.e. the energy is $E(C, \theta_1, \theta_2) = -L(C, \theta_1, \theta_2|I)$. The active contour seeks to minimize its energy, hence maximize the posterior likelihood. The usual minimization strategy is to conduct gradient descent on the energy function,

$$\begin{aligned} \frac{\partial C}{\partial t} &= -\nabla_C E(C, \theta_1, \theta_2), \\ \frac{\partial(\theta_1, \theta_2)}{\partial t} &= -\nabla_{\theta_1, \theta_2} E(C, \theta_1, \theta_2), \end{aligned} \quad (2)$$

where $\nabla_C E(C, \theta_1, \theta_2)$ is the gradient (first variation) of $E(C, \theta_1, \theta_2)$ with respect to C , and $\nabla_{\theta_1, \theta_2}$ is the gradient of $E(C, \theta_1, \theta_2)$ with respect to θ_1, θ_2 .

Under gradient descent, the m.a.p. active contour becomes stationary at the first local minimum it encounters. Unfortunately, this minimum may not occur at the true boundary. Fig. 2 illustrates how serious the problem can be. Fig. 2a and b show the result of using the m.a.p. gradient descent active contour on a cardiac ultrasound




Fig. 1. (a) The curve C , interior Ω_C , exterior $\tilde{\Omega}_C$ regions. (b) The image, boundary, and likelihood functions. Notation and Likelihood.




Fig. 2. M.a.p. active contour evolving under gradient is trapped away from the endocardium. The contour is initialized as the diamond and is trapped as the larger curve in the figure. The parts (a) and (b) show different initializations.

image. Both figures contain the same image. The central dark region in each image is the *blood pool* of the left ventricle and the surrounding bright annulus is the *myocardium*. The task is to find the *endocardium* which is the boundary between the blood pool and the myocardium. The a and b parts of Fig. 2 show a m.a.p. active contour (the log-likelihood functions are given in Section 7) initialized in the blood pool at two different locations as a diamond shaped curve. The figures also show the local minima where the active contour has become stationary – in each figure, it is the larger curve surrounding the diamonds. It is quite clear from the figures that the contour is trapped far from the endocardium. The figure also indicates that initializing the active contour at different locations does not help – it is simply trapped at another spurious local minimum.

4. Escaping from local minimum

To overcome this problem we need a strategy for the active contour to escape from spurious local minima. The following line of thought suggests one possibility:

- (1) When an active contour is initialized within a region and evolving by gradient descent, it is almost always trapped within the region before reaching the boundary. This suggests growing the active contour monotonically from the local minimum to escape from the minimum.
- (2) To understand which monotonic growth makes sense, consider Fig. 3 which shows the curve C trapped at a local minimum and two possible finger-like monotonic extensions of C . One extension causes the curve to penetrate through the true boundary, while the other extension keeps the curve inside. The extension that crosses the boundary adds pixels from the outside region to Ω_C and is likely to cause a greater increase in the energy than the extension




Fig. 3. Extending the curve from a local minimum.

that says in the region and adds pixels from within the region. Because we do not want the curve to cross the boundary (we want it to fill the inner region), we choose the latter extension. In other words, *from all possible monotonic extensions at the local minima we choose the one that gives the least increase in energy.*

- (3) The final step is to realize that descending into a minima and climbing out of it in the least steep direction can be realized with a single strategy. At all points in the evolution we simply ask the curve to grow monotonically with the least increase in energy. Since least increase is mathematically equal to most decrease, the curve will go downhill in the steepest direction into a local minima and uphill from there in the least steep direction. This is tunneling descent.

5. Tunneling descent

We next derive an algorithm that achieves the above by modifying gradient descent. We consider the discrete time version of gradient descent where it generates a sequence of curves C_1, C_2, \dots from an initial curve C_0 .

Let $v(s)$ be a normal vector field on C , where s is the arc-length of C . Define a neighborhood $\mathcal{G}(C)$ of C as the set of

all curves of the type $C(s) + v(s)$ for which $\|v\| = \int |v(s)|^2 ds \leq \Delta_1$, where $\Delta_1 > 0$, i.e.

$$\mathcal{G}(C) = \{D \in \Xi | D(s) = C(s) + v(s), \|v\| \leq \Delta_1\}. \quad (3)$$

Discrete gradient descent generates the sequence of curves C_1, C_2, \dots by choosing C_{n+1} as that curve in $\mathcal{G}(C_n)$ which has the least energy (Fig. 4a).

To modify gradient descent, define two more neighborhoods of C . Let $\mathcal{A}(C)$ be the set of curves whose area is greater than the area of C by a positive quantity $\Delta_2 > 0$:

$$\mathcal{A}(C) = \{D \in \Xi | \text{Area}(\Omega_D) - \text{Area}(\Omega_C) \geq \Delta_2\}. \quad (4)$$

Note that $\mathcal{A}(C)$ does not contain C .

Let $\mathcal{D}(C)$ be the set of curves which are “outside” C , i.e. if $D \in \mathcal{D}(C)$, then $\Omega_D \supseteq \Omega_C$,

$$\mathcal{D}(C) = \{D \in \Xi | \Omega_D \supseteq \Omega_C\}. \quad (5)$$

Then, $\mathcal{M}(C) = \mathcal{G}(C) \cap \mathcal{A}(C) \cap \mathcal{D}(C)$ is the set of all curves that are (1) in the gradient descent neighborhood of C , (2) have an area that is at least Δ_2 greater than the area of C , and (3) are “outside” C . Thus, $\mathcal{M}(C)$ is the set of curves in $\mathcal{G}(C)$ that are monotonically larger than C by Δ_2 .

The idea is to generate the sequence of curves by taking C_{n+1} as the energy minimizing curve in $\mathcal{M}(C_n)$, rather than $\mathcal{G}(C_n)$. Assume for the time being that the parameters θ_1, θ_2 of the probability distributions are known, and drop the dependence of the energy function on them. Then, starting from the initial curve C_0 , we create a sequence of curves $C_1, C_2, \dots, C_n, C_{n+1} \dots$ by the minimizations

$$C_{n+1} = \arg \min_{C \in \mathcal{M}(C_n)} E(C). \quad (6)$$

This is *tunneling descent*. A typical step of tunneling descent is illustrated in Fig. 5a and a typical sequence in Fig. 5b.

Taking C_{n+1} as the energy minimizing curve in $\mathcal{M}(C_n)$ has the following consequences:

- (1) Because all curves in $\mathcal{M}(C_n)$ have an area greater than the area of C_n , the sequence of curves generated by this is forced to grow monotonically, even at a local minimum.

- (2) Suppose C_n is not at a local minimum and $\mathcal{M}(C_n)$ contains curves whose energy is lower than the energy of C_n . Then, by our strategy, C_{n+1} is the curve in $\mathcal{M}(C_n)$ that has the least energy. That is, when a decrease in energy is compatible with monotonic growth, the curve will evolve in the direction of most energy decrease.
- (3) If C_n is at a local energy minimum, then energies of all curves in $\mathcal{M}(C_n)$ are greater than the energy of C_n . Now our strategy will ensure that C_{n+1} will grow from C_n with the *least increase* in energy. Thus the sequence of minimizations of Eq. (6) gives the evolution strategy we sought in Section 4.
- (4) Because the sequence of curves is required to grow monotonically, it will ultimately fill out the inside region and grow beyond the true boundary. A *stopping rule* is necessary to stop the evolution when this happens. As we show in Section 6, it is possible to create many stopping rules.
- (5) At the $n + 1^{st}$ iteration, define

$$B_{n+1} = C_k, \quad \text{where } k = \arg \min_{i=1, \dots, n+1} E(C_i). \quad (7)$$

That is, B_{n+1} is the curve with the least energy amongst all curves C_0, \dots, C_{n+1} . The curve B_{n+1} is the best estimate of the boundary at the $n + 1$ st iteration. Note that since the curves C_0, \dots, C_{n+1} are monotonically increasing, $B_{n+1} (= C_k \text{ for some } k \leq n + 1)$ is always inside C_{n+1} . Also, if $i \geq j$, then the least energy curve B_i must have energy that is less than or equal to the least energy curve B_j , i.e. $i \geq j$ implies $E(B_i) \leq E(B_j)$.

5.1. Tunneling descent with parameters

Tunneling descent with parameters is simply tunneling descent of Eq. (6) with parameters estimated simultaneously. The descent is initialized with curve C_0 and parameters $\theta_{1,0}, \theta_{2,0}$. It generates the sequence of curves and parameters $\{C_1, \theta_{1,1}, \theta_{2,1}\}, \{C_2, \theta_{1,2}, \theta_{2,2}\}, \dots, \{C_n, \theta_{1,n}, \theta_{2,n}\}, \{C_{n+1}, \theta_{1,n+1}, \theta_{2,n+1}\}$ according to

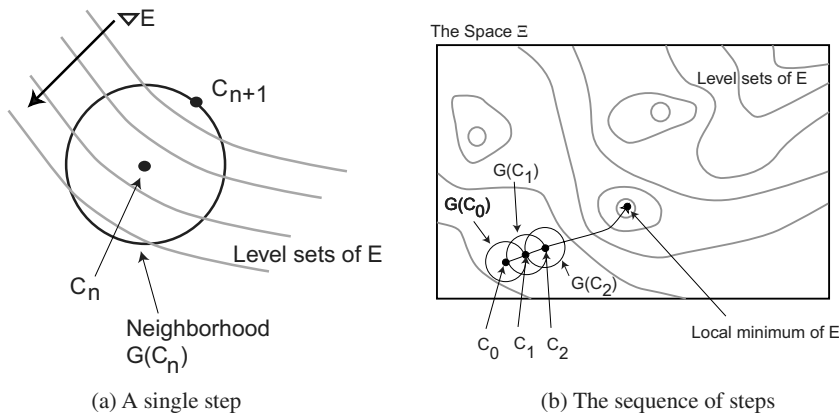


Fig. 4. (a) A single step. (b) The sequence of steps gradient descent in Ξ .




Fig. 5. (a) A single step. (b) The sequence of steps tunneling descent in Ξ .

$$C_{n+1} = \arg \min_{C \in \mathcal{M}(C_n)} E(C, \theta_{1,n}, \theta_{2,n}),$$

$$(\theta_{1,n+1}, \theta_{2,n+1}) = \arg \min_{\theta_1, \theta_2} E(C_n, \theta_1, \theta_2).$$

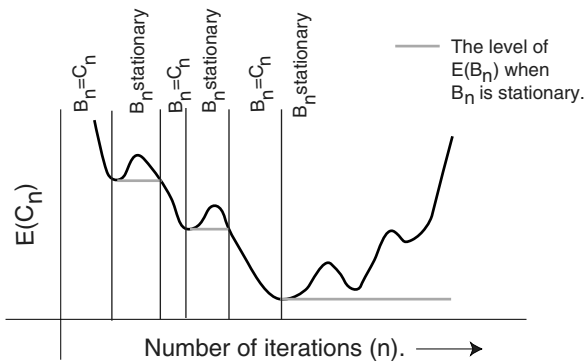
As before,

$$B_{n+1} = C_k, \quad \text{where } k = \arg \min_{i=1, \dots, n+1} E(C_i, \theta_{1,i}, \theta_{2,i})$$


is the best estimate of the boundary at the $n + 1$ st iteration.

5.2. Behavior of the best estimate

The estimate B_n has a characteristic which is illustrated in Fig. 6a. The figure shows a plot of the energy $E(C_n)$ as a



(a) Different regimes of evolution



(b) C_n inside the true boundary

(c) C_n outside the true boundary

Fig. 6. (a) Different regimes of evolution. (b) C_n inside the true boundary. (c) C_n outside the true boundary. The behavior of B_n and its relation to C_n .

function of n . Initially, as the energy decreases, B_n equals C_n , since C_n is the least energy curve until the n th iteration. Once C_n passes through a local minimum of the energy function, $E(C_n)$ increases but B_n stops at the local minimum (since $E(B_n)$ cannot increase). After $E(C_n)$ decreases once more to $E(B_n)$, B_n again equals C_n . Thus, the evolution has distinct regimes, depending on whether $B_n = C_n$ or B_n is stationary. Fig. 6a shows these regimes as regions between vertical lines. In the regimes where B_n is stationary, the value of $E(B_n)$ is shown as a gray horizontal line.

If the energy function is properly designed, then its minima near the true boundary are deeper than minima elsewhere. Therefore, as C_n approaches the boundary, B_n evolves through a series of local minima each having lower energy than the previous. But, once C_n crosses the boundary, B_n no longer follows C_n ; instead it tends to remain stationary near the true boundary. This behavior is useful in designing stopping rules.

6. Stopping rules for tunneling descent

We now turn to examining stopping rules for terminating tunneling descent. Two obvious stopping rules are: First, terminate the descent when the curve reaches the border of the entire image. Second, terminate the descent after the area of C_n has increased past a threshold. The latter stopping rule is useful when the approximate sizes of organs and upper bounds on the size are known.

A stopping rule that is faster than the above rules can be created based on the behavior of B_n discussed above:

- (1) Test whether the gray levels $I_{\Omega_{C_n - \Omega_{B_n}}}$ come from log-likelihoods l_1 or l_2 . If the gray levels appear to come from l_1 , then C_n is still inside the true boundary. If they come from l_2 , then C_n has crossed the boundary. This is a *model-based stopping rule*, since it uses the probability models implicit in l_1 and l_2 for stopping.

- (2) Test whether the gray levels $I_{\Omega_{C_n} - \Omega_{B_n}}$ and gray levels $I_{\Omega_{B_n}}$ come from the same distribution. If they do, then C_n is inside the true boundary. If not, C_n has evolved outside. This does not require use of the likelihoods and gives *model-free stopping rules*.

6.1. Model-based stopping

The classic model-based stopping rule is the generalized likelihood ratio test (GLRT):

$$l_2(I_{\Omega_{C_n} - \Omega_{B_n}} | \theta_{2,n}) - l_1(I_{\Omega_{C_n} - \Omega_{B_n}} | \theta_{1,n}) > T, \quad (10)$$

where $T > 0$ is a threshold, and the left-hand side is set to be 0 if $B_n = C_n$ (i.e. when $\Omega_{C_n} - \Omega_{B_n}$ is empty). If the left hand side exceeds the threshold, then the gray levels $I_{\Omega_{C_n} - \Omega_{B_n}}$ come from l_2 , and tunneling descent is terminated.

6.2. Model-free stopping

Rosenbaum's test (Neave and Worthington, 1988): Let $I_{\Omega_{B_n}}$ have a maximum gray level X_1 and $I_{\Omega_{C_n} - \Omega_{B_n}}$ have a maximum gray level X_2 . If $X_1 \geq X_2$, then the Rosenbaum test statistic is the number of pixels in Ω_{B_n} with gray levels greater than X_2 . If $X_2 > X_1$, then the test statistic is the number of pixels in $\Omega_{C_n} - \Omega_{B_n}$ whose gray levels are greater than X_1 . The test accepts the null hypothesis when the test statistic is less than a threshold $T > 0$. Else, the null hypothesis is rejected and tunneling descent terminated.

Two-sample χ^2 test: This test bins the gray levels $I_{\Omega_{B_n}}$ and $I_{\Omega_{C_n} - \Omega_{B_n}}$ separately into histograms R and S with the same number of bins, k . The test statistic is

$$\chi^2 = \sum_{i=1}^k \frac{(K_1 R_i - K_2 S_i)^2}{R_i + S_i}, \quad (11)$$

where the sum is over all bins of the histogram, R_i and S_i are the observed frequencies in the i th bins of the two histograms. Also, $K_1 = \sqrt{\frac{n_2}{n_1}}$ and $K_2 = 1/K_1$, where n_1 and n_2 are the number of pixels in Ω_{B_n} and $\Omega_{C_n} - \Omega_{B_n}$. The null hypothesis is accepted when $\chi^2 < T$, for a threshold $T > 0$. Else, tunneling descent is terminated.

6.3. The threshold T

All of the stopping rules require a threshold T . This threshold determines the false alarm rate of the test. We want to set the threshold T so that this rate is low.

GLRT: A good rule of thumb for setting T for the GLRT comes from observing an analogy between the stopping rule of Eq. (10) and the sequential probability ratio test (SPRT) of classical decision theory (Berger, 1980; Ghosh, 1971; Basseville and Nikiforov, 1993). In classical SPRT, data are incrementally added to the test statistic until the test statistic exceeds the threshold. The analogy is that the growing curve C_n adds more data to $I_{\Omega_{C_n} - \Omega_{B_n}}$ from the outside region till the test statistic exceeds T . Classically, the false alarm rate for the SPRT is estimated by Wald's approximation, which gives the false alarm rate, f_a , as a function of the threshold T :

$$f_a = e^{-T}. \quad (12)$$

This approximation is a reliable heuristic for selecting T in the stopping rule of Eq. (10).

Rosenbaum's test: The false alarm rate of Rosenbaum's test is (Neave and Worthington, 1988)

$$f_a = \frac{n_1!(n_1 + n_2 - T)!}{(n_1 + n_2)!(n_1 - T)!}, \quad (13)$$

where n_1 is the total number of pixels in Ω_{B_n} , n_2 is the total number of pixels in $\Omega_{C_n} - \Omega_{B_n}$, and T is the threshold.

χ^2 test: The false alarm rate of the χ^2 test is

$$f_a = 1 - \int_0^T \chi_{k-c}^2(x) dx, \quad (14)$$

where $k - c$ is the degrees of freedom in χ^2 distribution with $c = 0$ for unequal sample size and $c = 1$ otherwise, and T is the threshold.

6.4. The complete algorithm

To sum up, the complete tunneling descent algorithm is as follows:

- (1) Initialize the curve inside the region as C_0 and the parameters as $\theta_{1,0}, \theta_{2,0}$. Set $n = 0$.
- (2) Generate $C_{n+1}, \theta_{1,n+1}, \theta_{2,n+1}$ using Eq. (8).
- (3) Find B_{n+1} according to Eq. (9), and apply a stopping rule. If the rule indicates termination, stop with B_{n+1} as the boundary. Else, set $n = n + 1$ and go to 2.

6.5. Shrinking tunneling descent

So far, we assumed that the contour was initialized within a region and monotonically grown to find its boundary. However, there are no obstructions to creating an algorithm that monotonically shrinks the contour instead. The contour is now initialized outside the desired boundary and shrinks by reducing area. We call this *shrinking tunneling descent*. We do not pursue shrinking tunneling descent in this paper. The interested reader may refer to (Tao, 2005).

6.6. Numerical techniques

Details of the numerical techniques used in implementing tunneling descent are given in Appendix A. We give a brief outline here: First, we discretize the curve using a finite number of vertices (knot points) along the curve. The tunneling descent constraints are expressed using the discretization and are seen to be convex. Then, the curve minimization in tunneling descent (Eq. (8)) is carried out by the gradient projection algorithm which is a classic minimization techniques for convex constraints. The gradient projection onto the convex constraints is carried out by Dykstra's algorithm. The parameter minimization in Eq. (8) is carried out by ordinary gradient descent (or a closed form minimization, if possible).

7. Ultrasound segmentation

We now turn to our main application – segmentation of the endocardium in short-axis cardiac ultrasound images (Fig. 2). As mentioned in Section 2.1, we carried out an empirical investigation of the statistics of ultrasound images in Tao et al. (2002). The results of that study are summarized as follows:

- (1) As an approximation, the first-order gray levels in blood and tissue in short-axis cardiac images can be modeled by Gamma distributions. The gamma distribution has two parameters, a shape parameter which we denote by α and a scale parameter which we denote by β . The probability density function of the Gamma distribution is

$$p_\gamma(x|\alpha, \beta) = \frac{x^{\alpha-1}}{\Gamma(\alpha)\beta^\alpha} e^{-\frac{x}{\beta}}.$$

The shape parameters for blood and tissue can be fixed a priori to known constants. We denote the shape parameters of blood and tissue by α_1 and α_2 . Empirical evidence (Tao et al., 2002) suggests that their values are $\alpha_1 = 3.2$, $\alpha_2 = 7.8$. The scale parameters are estimated from the image.

- (2) Inside the endocardium is the blood pool and outside is tissue. The log-likelihood functions (as introduced and defined in Section 3.1) for the inside and outside of the endocardium are given by

$$l_1(I_\Sigma|\beta_1) = \int_\Sigma \log p_\gamma(I|\alpha_1, \beta_1) da, \text{ and} \quad (15)$$

$$l_2(I_\Sigma|\beta_2) = \int_\Sigma \log p_\gamma(I|\alpha_2, \beta_2) da, \quad (16)$$

where da is the differential area, and α_1 , α_2 and β_1 , β_2 are the shape and scale parameters for the gamma distributions for blood and tissue.

- (3) A log-normal prior serves well for the scale parameters (Tao et al., 2002). The log-likelihood of this prior on β_1 , β_2 is

$$L_{\text{prior}}(\beta_1, \beta_2) = \frac{-1}{2\sigma_\beta^2} \left(\log \frac{\beta_1}{\beta_2} - \mu_\beta \right)^2 \quad (17)$$

where we have dropped all terms that are independent of β_1 , β_2 , and μ_β and σ_β are constants (the mean and standard deviation of the log normal). Based on empirical evidence (Tao et al., 2002), their values were set to $\mu_\beta = 2.35$ and $\sigma_\beta = 1.18$.

- (4) We use the standard arc-length prior on C , so that the log-likelihood is

$$L_{\text{prior}}(C) = -\lambda \text{Length}(C), \quad (18)$$

where $\lambda > 0$ is a constant. This term is more commonly known as the “arc-length internal energy” of the active contour.

The value of λ is set to 1 in all experiments.

Thus, the total log-likelihood prior for segmenting ultrasound images is

$$L_{\text{prior}}(C, \beta_1, \beta_2) = L_{\text{prior}}(C) + L_{\text{prior}}(\beta_1, \beta_2). \quad (19)$$

This log-likelihood prior along with the likelihoods of Eqs. (15) and (16) define the posterior log-likelihood of Eq. (1).

8. Experiments

We extensively tested tunneling descent on clinical B-mode short-axis cardiac images. The images were acquired from different subjects with an Acuson Sequoia C256 imaging system. All images with significant data drop out were excluded; 44 images remained, which were processed. These images were of different subjects obtained under conditions that we did not have any control over (the images themselves came from a small archive of images that had already been created). The active contour was initialized in the blood pool and propagated outwards by tunneling descent. The GLRT stopping rule was used. The reason for choosing the GLRT is explained below in Section 8.1.

The values of the tunneling descent parameters are given in Table 1. Informal experimentation showed little dependence on the performance of the algorithm with respect to Δ_1 and Δ_2 . A detailed study of the algorithm with respect to T is reported below in Section 8.1.

Tunneling descent segmented all images successfully. That is to say, in no case did the stopping rule terminate the descent early and a visual inspection of all results showed that the segmentation was very close to the endocardium. A more quantitative evaluation of the performance is given below in Section 8.2.

To assess the quality of results visually, Fig. 7 shows four typical segmentation results with the active contour initialized in blood. In each row, the first figure shows the image, the initial diamond shaped contour, and the final contour found by tunneling descent. The second figure in each row plots the energy $E(C_n, \beta_{0,n}, \beta_{1,n})$ as a function of n . The local minima in the energy function are indicated by vertical lines. These figures show that tunneling descent escaped through multiple local minima to find the endocardium. Gradient descent would have been trapped in any one of these.

No parameter tweaking was involved in any of the 44 segmentations. Quite literally, the same algorithm worked on all images.

8.1. Stopping rule performance

Three different stopping rules were proposed in Section 6 and we experimentally evaluated their performance.

We evaluated stopping rules for all 44 images. The false alarm rates were chosen as $f_a = 10^{-4}$, 10^{-6} , 10^{-8} , 10^{-10} , 10^{-12} and the stopping thresholds calculated according to

Table 1
Values of constants used in experiments

	Δ_1	Δ_2	Stop. rule threshold T
Tun. Desc.	10	20	500

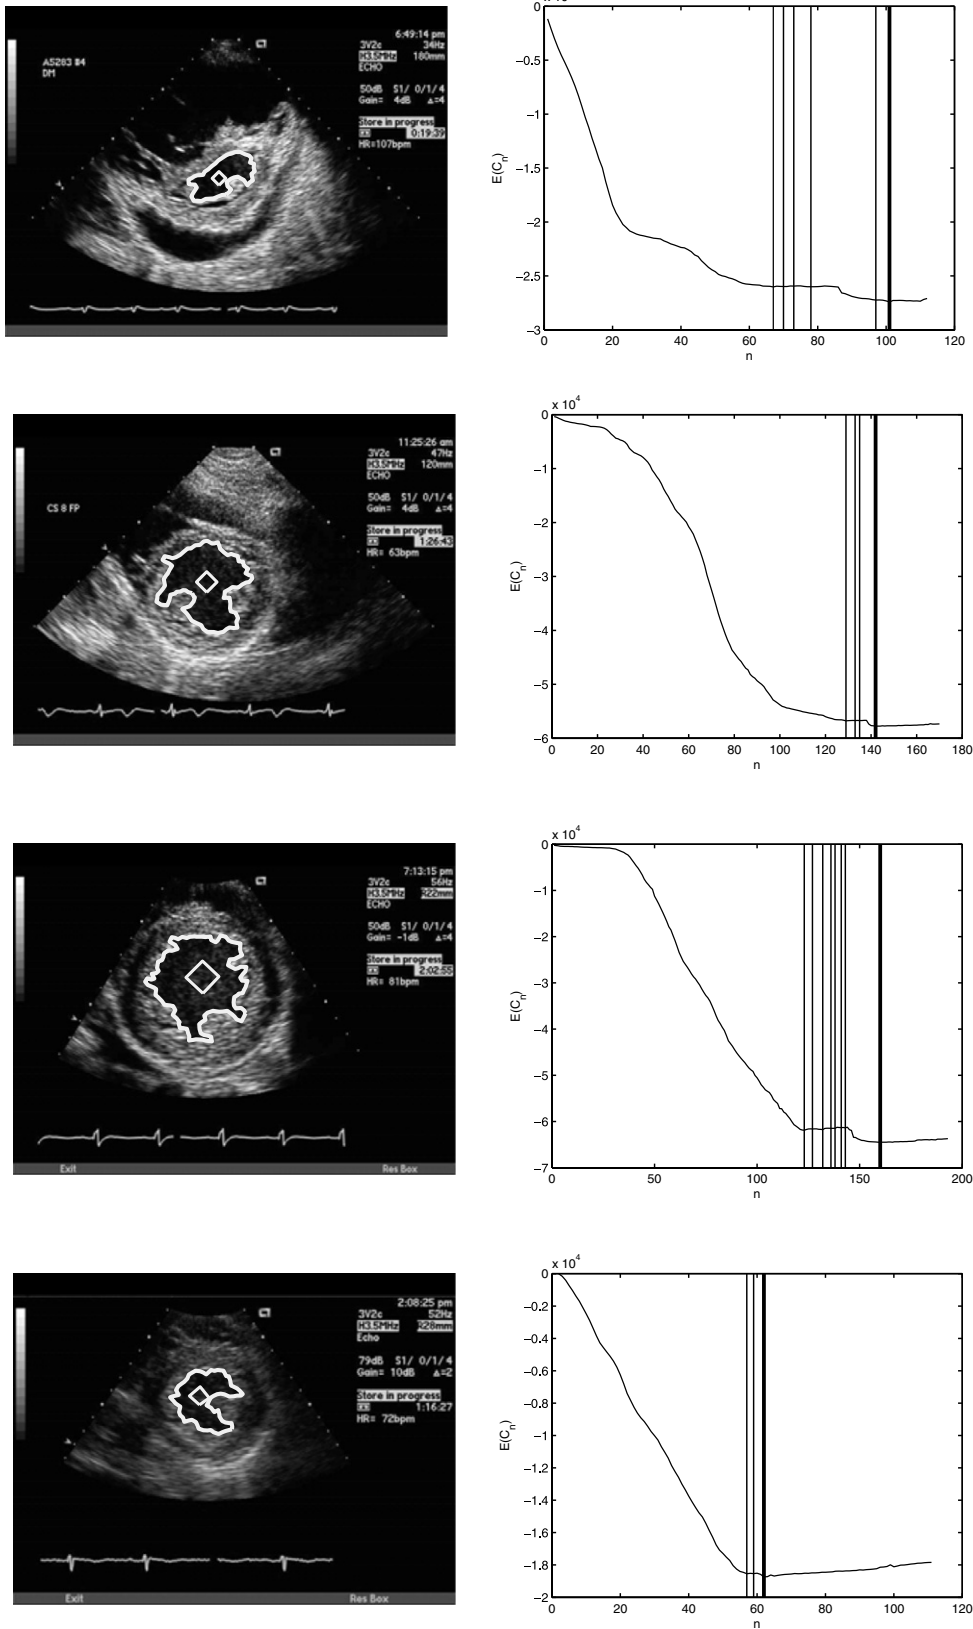


Fig. 7. Segmentation by tunneling descent.

Eqs. (12)–(14), where we used 20 bins for χ^2 test. For low values of threshold (high f_a), the active contour very clearly failed to reach the endocardium. For very high values of

threshold (low f_a), sometimes the contour completely escaped from the myocardium. These very obvious failures were counted as such. The failure rates of the three stopping

Table 2
Failure rates of stopping rules

f_a	10^{-4}	10^{-6}	10^{-8}	10^{-10}	10^{-12}
GLRT (%)	4.5	2.3	0.0	0.0	0.0
Rosenbaum (%)	6.8	6.8	9.1	11.4	11.4
χ^2 (%)	4.5	2.3	4.5	6.8	9.1

rules are shown in Table 2 as a function of f_a . Note that of the three stopping rules GLRT is the best behaved. Its failure rate monotonically decreases with decreasing f_a .

For the images where the stopping rule did not fail, the terminating C_n penetrated into the myocardium while the terminating B_n stayed close to the boundary. We measured the area between the terminating C_n and B_n , and called it *penetration area*. This evaluated the extent to which the stopping rule allowed C_n to penetrate into the tissue before termination. The average and standard deviation of the penetration area (measured in pixels) is shown in Fig. 8 as a function of f_a . This figure shows that penetration of GLRT is small (of the order of a few hundred pixels), and is lower than the penetration of Rosenbaum’s test and χ^2 test. Also the penetration area of the GLRT is stable with respect to f_a . Clearly, GLRT outperforms the other two tests and is our choice of the stopping rule. Except for this experiment, all tunneling descents reported in this paper use the GLRT stopping rule.

8.2. Comparison with manual segmentation and balloons

Next, we compared the algorithm segmentation with manual segmentation. Manual segmentations are laborious and notoriously prone to fatigue, and multiple manual segmentations of the same image are usually desirable. For a subset of 19 images, we were able to obtain two manual segmentations which we compared to tunneling descent. To assess the advantage of tunneling descent over classical gradient descent, we also created a m.a.p. active contour

with a balloon force which was initialized in the same location as tunneling descent. Since it is difficult to estimate a priori what the right amount of balloon force is, three different balloon forces were used bracketing a value that appeared to get the contours out of most local minima.

Thus, for each image, we had two manual segmentations, one tunneling descent segmentation, and three gradient descent m.a.p. with balloon force segmentations. The two manual segmentations were taken as the gold standard and other segmentations were compared with them.

Any two segmentation curves C_1 and C_2 were compared by two measures. The first, denoted ϵ , measured the relative area of the non-overlapping regions of the curves as one minus the ratio of the area of overlap of the two curves to the average area of the interiors of the two curves (Fig. 9):

$$\epsilon = 1 - \frac{\text{Area}(\Omega_{C_1} \cap \Omega_{C_2})}{(\text{Area}(\Omega_{C_1}) + \text{Area}(\Omega_{C_2}))/2}. \tag{20}$$

The second, denoted δ , measured the maximum absolute deviation of one curve from the other:

$$\delta = \max(d(C_1, C_2), d(C_2, C_1)), \tag{21}$$

where $d(\dots)$ measures the maximum distance of the curve in the first argument from the curve in the second:

$$d(E, F) = \max_{f \in F} \min_{e \in E} \|e - f\|,$$

where $\| \cdot \|$ is the Euclidean distance.

For every image, we first measured ϵ_m , the extent of non-overlap of the two manual segmentations. Then, for each machine segmentation, we measured its non-overlap ϵ with both manual segmentations and took the average. For tunneling descent, this average non-overlap is denoted ϵ_t . For the three balloon active contours, the average non-overlaps are denoted $\epsilon_{b1}, \epsilon_{b2}, \epsilon_{b3}$.

Similarly, we measured δ_m between the two manual segmentations. And for each machine segmentation, we measured the δ with both manual segmentations and took at the average. For tunneling descent the average is denoted




Fig. 8. Penetration area for different stopping rules.

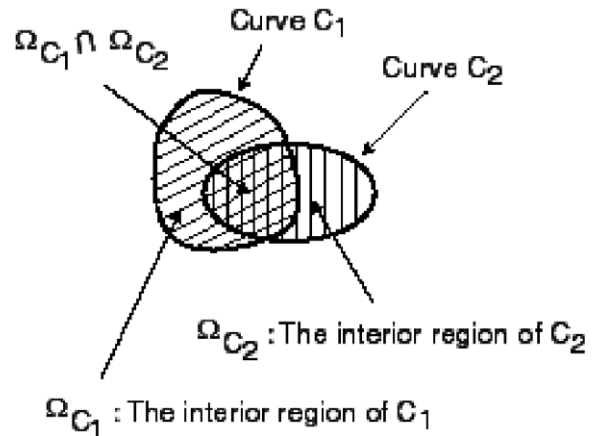


Fig. 9. Curve overlap.

δ_t . For the balloon active contours, the averages are denoted $\delta_{b1}, \delta_{b2}, \delta_{b3}$.

Fig. 10a shows the cumulative distribution of ϵ for all 19 images. For each value of ϵ on the x -axis, the plot gives the corresponding fraction of segmentations for which the measured ϵ was less than or equal to this value. In the figures, the thick solid lines are for ϵ_m , and the thin solid lines for ϵ_t and the other lines, dashed or dotted, for $\epsilon_{b1}, \epsilon_{b2}$ and ϵ_{b3} . As seen in Fig. 10a, all manual segmentations have ϵ_m less than 0.125. Thus, roughly speaking, all manual segmentations were within 12.5% of each other. The cumulative error curve of tunneling descent in Fig. 10a stays close to ϵ_m . In fact, tunneling descent results are only slightly less similar to the manual segmentations than the manual segmentations are to each other. Also tunneling descent significantly outperforms all of the active contours with balloon forces.

Fig. 10b shows the cumulative distribution of δ . Here too it is clear that δ_t stays consistently close to δ_m . Further δ_t is closer to δ_m than any of the balloon δ 's. The only exception occurs for the two rightmost values of δ_{b2} . However this does not imply that δ_{b2} has a superior performance to δ_t , because δ_{b2} is noticeably worse elsewhere. This only shows that it is difficult to choose a good balloon force, and that the correct balloon force can vary from image to image.

8.3. Sensitivity to initialization

Finally, we evaluated the sensitivity of tunneling descent segmentation to initialization. Figs. 11a–j show the same ultrasound image with the active contour initialized as the diamond in the central blood pool. In Fig. 11a the contour is initialized in the middle of the blood pool. In Figs. 11b–e the contour with initialized at the same y co-ordinate as Fig. 11a, but with the x coordinate varying by

$\Delta x = -40, -20, 20, 40$ pixels, respectively. In Figs. 11d–j the contour with initialized at the same x co-ordinate as Fig. 11a, but with the y coordinate varying by $\Delta y = -40, -20, 20, 40, 60$ pixels respectively (the blood pool is slightly longer in the vertical direction than the horizontal direction in the center). The larger curves in figures are the final segmentations obtained by tunneling descent. Inspection of Figs. 11a–j reveals that there are only minor variations in the final segmentations.

For a more objective assessment, we turned to the measures ϵ and δ as defined in Eqs. (20) and (21). The segmentations of Figs. 11b–j were compared with the segmentation of Fig. 11a. The results are shown in Table 3. Note that all ϵ 's are in the 1–3% range, and all the δ 's in the 6–7 pixels range. For an idea of the magnitude of δ , note that 7 pixels is the size (linear dimension) of one grain of speckle in the images in Fig. 11, and that the bounding box for the segmentation in Fig. 11a is 136 pixels by 128 pixels. This confirms the informal observation that the segmentations in Figs. 11a–j are insensitive to initialization. Experimentation with other images shows that these numbers are typical.

9. Limitations of the likelihood function

As mentioned in Section 1, escaping from spurious local minima is only one of many problems that have to be addressed before a robust ultrasound segmentation algorithm can be developed. In particular, dropout needs to be addressed. Although dealing with dropout is not the aim of this paper, it is interesting to see the performance of tunneling descent on images with dropout. Fig. 12 shows an example. The active contour is initialized as the small diamond shaped curve in the center of the image. It then evolves and terminates at the larger curve as the final segmentation. In the absence of any other information, the

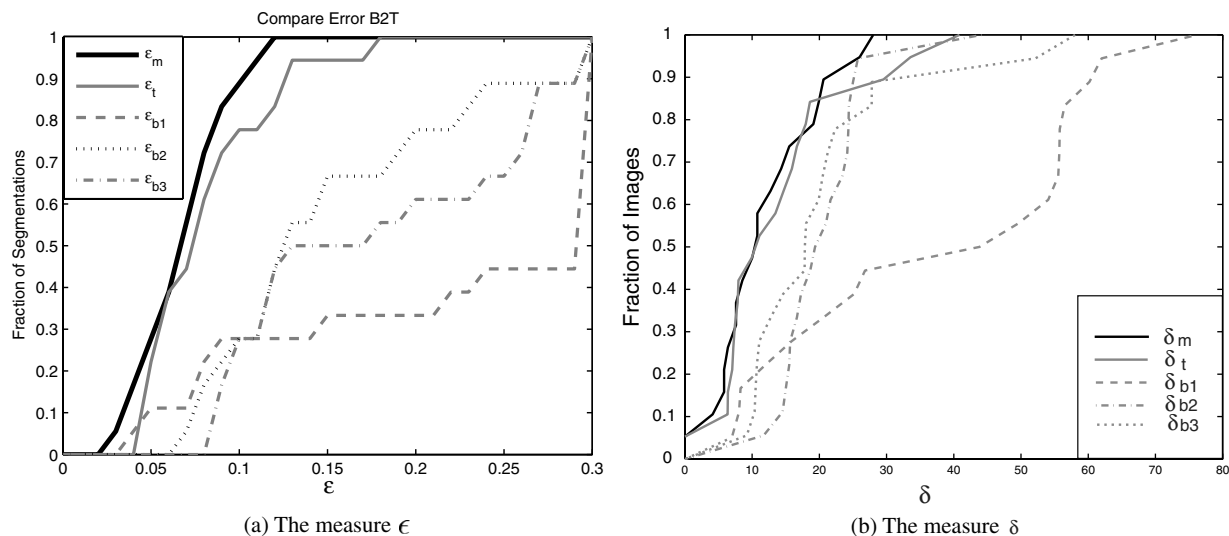


Fig. 10. (a) The measure ϵ . (b) The measure δ Comparison with manual segmentation and balloons.

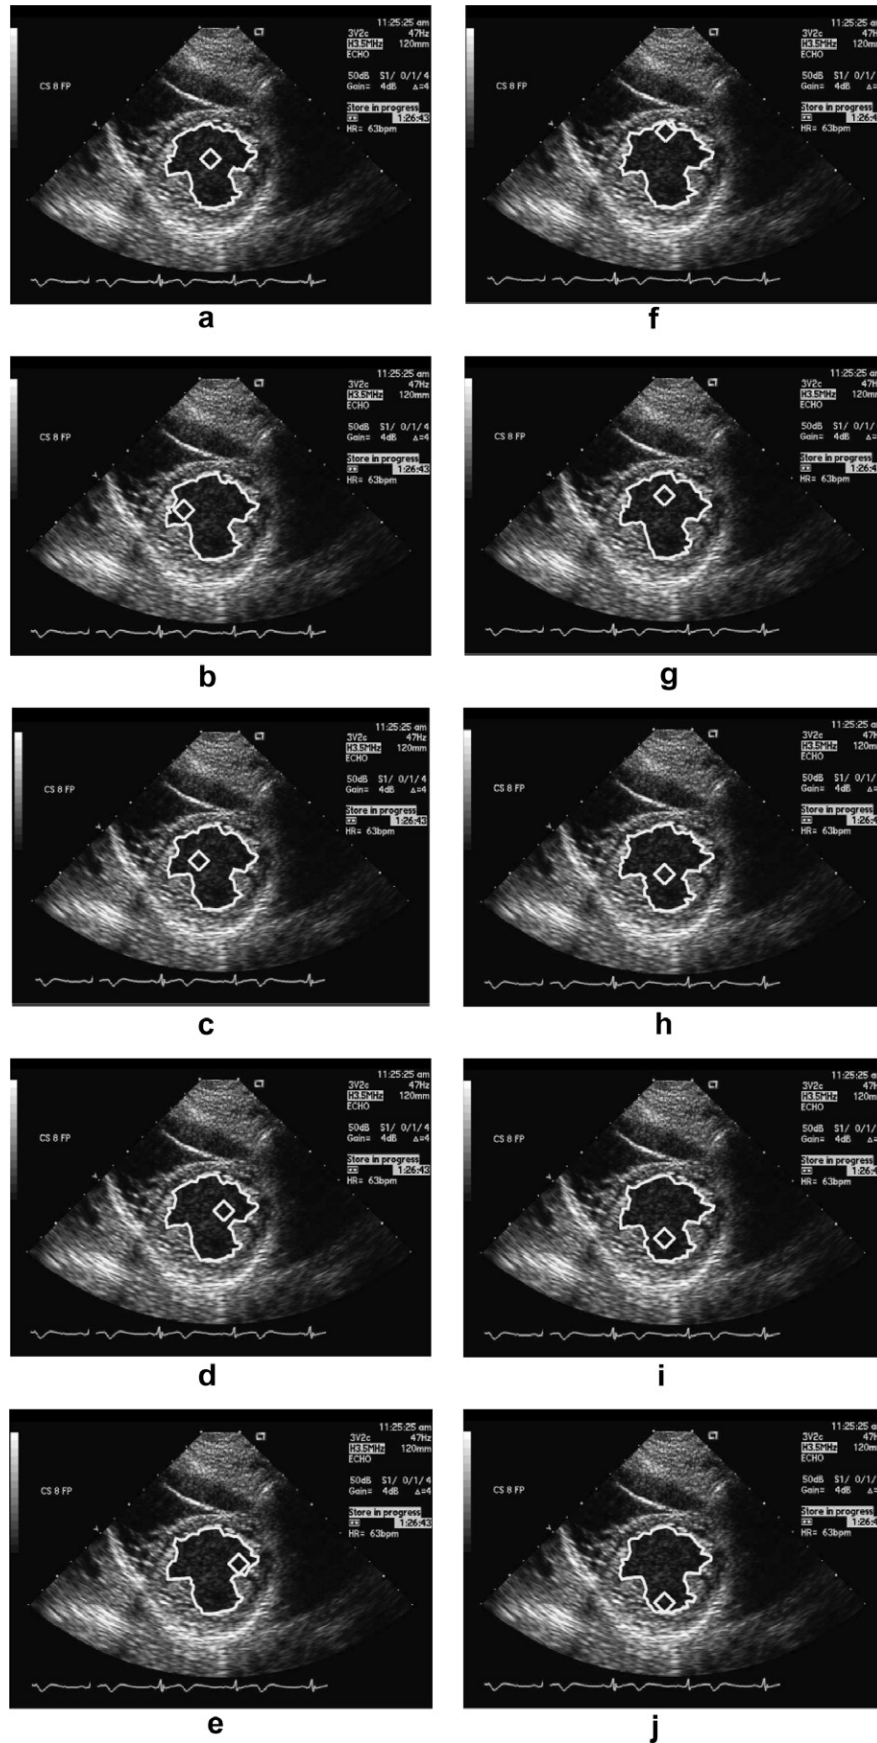


Fig. 11. Sensitivity to initialization.

Table 3
Sensitivity to initialization

	Fig. 11b	Fig. 11c	Fig. 11d	Fig. 11e	Fig. 11f
ϵ	0.021	0.023	0.016	0.023	0.023
δ (pixels)	6.71	7.00	6.33	7.07	7.00
	Fig. 11g	Fig. 11h	Fig. 11i	Fig. 11j	
ϵ	0.023	0.024	0.022	0.019	
δ (pixels)	6.71	7.07	7.07	6.08	




Fig. 12. Leakage due to dropout.

segmentation of Fig. 12 into dark and bright regions is quite good. On the other hand, anatomical knowledge and familiarity with ultrasound imaging tells us that there is significant dropout around the 9 o'clock region of the endocardium and the active contour has leaked through the dropout.

This is clearly not the fault of tunneling descent, but of the simple likelihood function which does not model dropout and hence is unable to understand it. The problem can be fixed by modelling the dropout as an image multiplicative term that goes to zero in the dropout region and by adding a shape prior (e.g. as in Chen et al. (2002)) that informs the contour how to bridge the dropout region. We are currently developing such a modified likelihood function, and hope to report it in forthcoming publications. Modifying the likelihood function has no effect on tunneling descent as the evolution strategy.

10. Conclusions

We proposed a new strategy called *tunneling descent* for evolving active contours. The strategy is based on the idea that if a contour is evolved monotonically in such a way that it increases its area at a non-zero rate while maximally decreasing its energy at each step, then it will enter into a local minimum and escape from it in the right direction. A set of stopping rules can terminate the evolution once the contour has grown past the true boundary.

The capacity to overcome spurious minima is evident in experiments where the algorithm is used to find the endocardium in ultrasound images. For segmenting short-axis

cardiac ultrasound images, the performance of the algorithm is comparable to manual segmentation and superior to gradient descent m.a.p. active contours with balloon forces. Further, tunneling descent can be used without any parameter tweaking. Finally, tunneling descent segmentations are insensitive to initialization.

Acknowledgements

We gratefully acknowledge the many discussions with, and support and data from Dr. C. Carl Jaffe of the National Cancer Institute. This research was supported by the Grant R01 LM6911 from the National Library of Medicine and the Grant R01 HL077810 from the National Heart Lung and Blood Institute.

Appendix A. Numerical techniques

A.1. Discretization

We discretize the curve as a polyhedron with k vertices and enumerate the vertices in a clockwise manner (Fig. 13a) with the i th vertex of a curve denoted by $v_i = (x_i, y_i)$, and the entire set of vertices denoted by $\mathbf{v} = [v_1 \dots v_k]^T$. The curve itself is denoted as $C(\mathbf{v})$ to explicitly show the dependence on \mathbf{v} .

Moving the vertices \mathbf{v} moves the curve. Set $\mathbf{w} = [w_1, \dots, w_k]$ and let $\mathbf{v} + \mathbf{w}$ represent the vertices of a curve $D(\mathbf{v} + \mathbf{w})$ that is displaced from C by \mathbf{w} (Fig. 13a). Then, all of the neighborhoods that define tunneling descent can be expressed in terms of \mathbf{w} as shown below. In these formulae, all index addition and subtraction is modulo k .

[1] $\mathcal{G}(C)$ is the set of all curves $D(\mathbf{v} + \mathbf{w})$ for which

$$\sum_{i=1}^k \|w_i\|^2 \leq \Delta_1. \quad (\text{A.1})$$

[2] $\mathcal{A}(C)$ is the set of all curves $D(\mathbf{v} + \mathbf{w})$ for which

$$\sum_{i=1}^k \{(v_{i-1} - v_{i+1})^\perp\}^T w_i \geq \Delta_2, \quad (\text{A.2})$$

where the operator $^\perp$ indicates rotation by 90° in the counter-clockwise direction. This constraint is the first-order




Fig. 13. Curves and vertices.

approximation to $\text{Area}(\Omega_D) - \text{Area}(\Omega_C) \geq \Delta_2$, valid when w_i are small.

[3] $\mathcal{D}(C)$ is the set of all curves $D(\mathbf{v} + \mathbf{w})$ for which

$$\begin{aligned} \{(v_i - v_{i-1})^\perp\}^T w_i &\geq 0, \\ \{(v_{i+1} - v_i)^\perp\}^T w_i &\geq 0 \end{aligned} \quad (\text{A.3})$$

for all $i = 1, \dots, k$. This constraint expresses $\Omega_C \subseteq \Omega_D$ (Fig. 13b).

Given the vertices \mathbf{v} , let $\mathcal{M}(\mathbf{v})$ denote the set of \mathbf{w} that satisfy the constraints of Eqs. (A.1)–(A.3). A sufficient condition for $\mathcal{M}(\mathbf{v})$ to be not empty is:

$$\frac{\Delta_2}{\sum_i^k (v_{i-1} - v_{i+1})^2} \leq \Delta_1. \quad (\text{A.4})$$

All constraints in Eqs. (A.1)–(A.3) are convex, hence the set $\mathcal{M}(\mathbf{v})$, which is the intersection of the constraints, is also convex.

When discretized tunneling descent is initialized at $C_0(\mathbf{v}_0)$, it generates the sequence of the curves and parameter estimates $(C_1(\mathbf{v}_1), \theta_{1,1}, \theta_{2,1}), (C_2(\mathbf{v}_2), \theta_{1,2}, \theta_{2,2}), \dots, (C_{n+1}(\mathbf{v}_{n+1}), \theta_{1,n+1}, \theta_{2,n+1})$ according to

$$\mathbf{v}_{n+1} = \mathbf{v}_n + \arg \min_{\mathbf{w} \in \mathcal{M}(\mathbf{v}_n)} E(C(\mathbf{v}_n + \mathbf{w}), \theta_{1,n}, \theta_{2,n}), \quad (\text{A.5})$$

$$(\theta_{1,n+1}, \theta_{2,n+1}) = \arg \min_{\theta_1, \theta_2} E(C(\mathbf{v}_{n+1}), \theta_1, \theta_2). \quad (\text{A.6})$$

The minimization in Eq. (A.6) is unconstrained and carried out by gradient descent. The minimization in Eq. (A.5) is the minimization of the nonlinear function E in the convex set $\mathcal{M}(\mathbf{v}_n)$. There are many standard techniques for minimization in a convex set (Bertsekas, 1996); we use the gradient projection method.

A.2. Minimization by gradient projection

The gradient projection method works as follows: starting from an initial value, it proceeds in the negative gradient direction. If it hits the boundary of the convex set, it proceeds by projecting the gradient onto the convex set.

$$\mathbf{w}^k = \mathbf{w}^{k-1} + (\bar{\mathbf{w}}^{k-1} - \mathbf{w}^{k-1}),$$

where

$$\bar{\mathbf{w}}^{k-1} = [\mathbf{w}^{k-1} - s \nabla_{\mathbf{w}} E(C(\mathbf{v}_n + \mathbf{w}), \theta_{1,n}, \theta_{2,n})]^\perp.$$

Here $[\cdot]^\perp$ denotes projection on the set $\mathcal{M}(\mathbf{v}_n)$, and s is a positive scalar. The sequence terminates when $\|(\bar{\mathbf{w}}^{k-1} - \mathbf{w}^{k-1})\|$ falls below a small positive number, similar to the termination of gradient descent.

Because the convex set $\mathcal{M}(\mathbf{v}_n)$ is defined by the intersection of three convex sets given by Eqs. (A.1)–(A.3), its projection operator is obtained by Dykstra’s algorithm (Deutsch, 2000), which is a classic algorithm for projecting onto an intersection of convex sets.

A.2.1. Dykstra’s algorithm

Given any vector $\mathbf{u}^0 = [u_1^0, \dots, u_k^0]$, Dykstra’s algorithm generates a sequence of iterates \mathbf{u}^m that converge to $[\mathbf{u}^0]^\perp$, the projection of \mathbf{u}^0 onto $\mathcal{M}(\mathbf{v}_n)$. Defining $\mathbf{e}^{-3} = \mathbf{e}^{-2} = \mathbf{e}^{-1} = \mathbf{0}$, the iterations of Dykstra’s algorithm are:

$$\begin{aligned} \mathbf{u}^{m+1} &= [\mathbf{u}^m + \mathbf{e}^{m-3}]_{(m \bmod 3)}^\perp, \\ \mathbf{e}^m &= (\mathbf{u}^m - \mathbf{e}^{m-3}) - \mathbf{u}^{m+1}, \text{ for } m = 0, 1, 2, \dots, \end{aligned}$$

where $[\cdot]_0^\perp, [\cdot]_1^\perp, [\cdot]_2^\perp$, denotes projections on to the convex sets $\mathcal{G}(C_{\mathbf{v}_n}), \mathcal{A}(C_{\mathbf{v}_n}), \mathcal{D}(C_{\mathbf{v}_n})$. The iterations are terminated when $\|\mathbf{u}^{m+1} - \mathbf{u}^m\|$ falls below a small number.

The projection operators $[\cdot]_0^\perp, [\cdot]_1^\perp$ are given by

$$[\mathbf{u}]_0^\perp = \begin{cases} \mathbf{u} & \text{if } \|\mathbf{u}\|^2 < \Delta_1, \\ \sqrt{\frac{\Delta_1}{\sum_i^k \|w_i\|^2}} \mathbf{u} & \text{otherwise.} \end{cases} \quad (\text{A.7})$$

$$[\mathbf{u}]_1^\perp = \begin{cases} \mathbf{u} & \text{if } \sum_i^k \{(v_{i-1} - v_{i+1})^\perp\}^T u_i \geq \Delta_2, \\ \left(\mathbf{1} + \frac{\Delta_2 - \sum_i^k \{(v_{i-1} - v_{i+1})^\perp\}^T u_i}{\sum_i^k (v_{i-1} - v_{i+1})^2} \right) \mathbf{u} & \text{otherwise.} \end{cases} \quad (\text{A.8})$$

The projection operator $[\cdot]_2^\perp$ is obtained by applying each of the following projection operators $[\cdot]_{2,i}^\perp$ to the corresponding co-ordinate u_i of \mathbf{u} for $i = 1, \dots, k$ respectively:

$$[u_i]_{2,i}^\perp = \begin{cases} u_i & \text{if } \{(v_i - v_{i-1})^\perp\}^T u_i \geq 0, \text{ and } \{(v_{i+1} - v_i)^\perp\}^T u_i \geq 0 \\ \left(1 - \frac{\{(v_i - v_{i-1})^\perp\}^T u_i}{\|v_i - v_{i-1}\|^2}\right) u_i & \text{if } \{(v_i - v_{i-1})^\perp\}^T u_i < 0, \text{ and } \{(v_{i+1} - v_i)^\perp\}^T u_i \geq 0 \\ \left(1 - \frac{\{(v_{i+1} - v_i)^\perp\}^T u_i}{\|v_{i+1} - v_i\|^2}\right) u_i & \text{if } \{(v_i - v_{i-1})^\perp\}^T u_i \geq 0, \text{ and } \{(v_{i+1} - v_i)^\perp\}^T u_i < 0 \\ 0 & \text{if } \{(v_i - v_{i-1})^\perp\}^T u_i < 0, \text{ and } \{(v_{i+1} - v_i)^\perp\}^T u_i < 0, \end{cases}$$

The method terminates when the gradient or the projected gradient is zero.

When applied to the minimization in Eq. (A.5), gradient projection is initialized with $\mathbf{w}^0 = \mathbf{0}$ and generates a sequence of $\mathbf{w}^k, k = 1, \dots$ according to

where, as before, $i + 1, i - 1$ are modulo k .

Dykstra’s algorithm might appear to be computationally expensive, but, in practice, we find that only a few iterations are needed to get a numerically reliable projection.

A.3. Vertex deletion and insertion

To prevent the active contour from self crossing, we add and delete the vertices dynamically during curve evolution. This is required of all active contour implementations that discretize the curve with knot points. In our implementation, when two neighboring vertices v_i and v_{i+1} are greater than 5 pixels, a vertex is added at the mid point of these two vertices. If two vertices from a limited neighborhood v_i, v_{i+k} with $1 \leq k \leq n/2$ are less than 5 pixels apart, we delete all vertices between them.

This completes the description of the numerical algorithm for tunneling descent.

References

- Abyratne, U.R., Petropulu, A.P., Reid, J.M., 1996. On modeling the tissue response from ultrasonic B-scan images. *IEEE Trans. Med. Imag.* 15 (4).
- Aston, E.A., Parker, K.J., 1995. Multiple resolution Bayesian segmentation of ultrasound images. *Ultrason. Imag.* 17 (4), 291–304.
- Bamber, J.C., Daft, C., 1986. Adaptive filtering for reduction of speckle in ultrasonic pulse-echo images. *Ultrasonics*, 41–44.
- Basseville, M., Nikiforov, I.V., 1993. *Detection of Abrupt Changes: Theory and Application*. Prentice-Hall.
- Belohlavec, M., Greenleaf, J.F., 1997. Detection of cardiac boundaries in echocardiographic images using a customized order statistics filter. *Ultrasonic Imaging* 19 (2), 127–137.
- Berger, J.O., 1980. *Statistical Decision Theory and Bayesian Analysis*. Springer-Verlag.
- Bertsekas, D.P., 1996. *Nonlinear Programming*. Athena Scientific.
- Bosch, J.G., Mitchell, S.C., Lelieveldt, B.O.F., Nijland, F., Kamp, O., Sonka, M., Reiber, J.H.C., 2002. Automatic segmentation of echocardiographic sequences by active appearance motion models. *IEEE Trans. Med. Imag.* 21 (11), 1374–1383.
- Butkhardt, C.B., 1978. Speckle in ultrasound B-mode scans. *IEEE Trans. Son. Ultra.* SU-25 (Jan), 1–6.
- Chen, Y., Tagare, H.D., Thiruvankadam, S.R., Huang, F., Wilson, D., Geiser, A., Gopinath, K., Briggs, R., 2002. Using prior shapes in geometric active contours in a variational framework. *Intl. J. Computer Vision* 50 (3), 315–328.
- Cohen, L.D., Cohen, I., 1991. On active contour models and balloons. *CVGIP: Image Understand.* 53 (2), 211–218.
- Cohen, L.D., Cohen, I., 1993. Finite element methods for active contour models and balloons for 2D and 3D images. *IEEE Trans. Pattern Anal. Machine Intelligence* 15 (Nov).
- Cramblitt, R.M., Parker, K.J., 1999. Generation of non-Rayleigh speckle distribution using marked regularity models. *IEEE Trans. Ultra. Ferr. Freq. Contr.* 46 (4).
- Czerwinski, N., Jones, D.L., O'Brien Jr., W.D., 1999. Line and boundary detection in speckle images-application to medical ultrasound. *IEEE Trans. Med. Imag.* 8 (2), 126–136.
- Deutsch, F., 2000. Best approximation in inner product spaces. In: *CMS Books in Mathematics*. Springer.
- Geiser, E.A., Wilson, D.C., 1996. Two-dimensional echocardiographic image analysis by computer. In: John Sutton, M.G. St., Oldershaw, P.J., Kotler, M.N. (Eds.), *Textbook of Echocardiography and Doppler in Adults and Children*. Blackwell Science (Chapter 21).
- Geiser, E.A., Wilson, D.C., Wang, D., Conetta, D.A., Murphy, J.D., Hutson, A.D., 1998. Autonomous epicardial and endocardial boundary detection in echocardiographic short axis images. *J. Am. Soc. Echocardiogr.* 11, 338–348.
- Ghosh, B.K., *Sequential Tests of Statistical Hypotheses*, Addison-Wesley Publishing Company, 1971.
- Goodman, J.W., 1975. Laser speckle and related phenomenon. In: Dainty, J.C. (Ed.). Springer-Verlag.
- Guofang, X., Brady, M., Noble, J.A., Zhang, Y., 2002. Segmentation of ultrasound B-mode images with intensity inhomogeneity correction medical imaging. *IEEE Trans. Med. Image* 21 (1).
- Han, C.Y., Lin, K.N., Wee, W.G., Mintz, R.M., Porembka, D.T., 1991. Knowledge-based image analysis for automated boundary extraction of transesophageal echocardiographic left-ventricular images. *IEEE Trans. Med. Imag.* 10 (4), 602–609.
- Hill, A., Taylor, C.J., 1992. Model-based image interpretation using genetic algorithms. *Image Vis. Comput.* 10 (5).
- Insana, M.F., Myers, K.J., Grossman, L.W., 2000. Signal modeling for tissue characterization. In: Sonka, M., Fitzpatrick, M.J. (Eds.). In: *Handbook of Medical Imaging*, vol. 2. SPIE Press.
- Jacob, G., Noble, J.A., Behrenbruch, C., Kelion, A.D., Banning, A.P., 2002. A shape-space-based approach to tracking myocardial borders and quantifying regional left-ventricular function applied in echocardiography. *IEEE Trans. Med. Imag.* 21 (March).
- Jakeman, E., Tough, R.J.A., 1987. Generalized K distribution: a statistical model for weak scattering. *J. Opt. Soc. Am.* 4 (Sept), 1764–1772.
- Karaman, M., Kutay, M.A., Bozdagi, G., 1995. An adaptive speckle suppression filter for medical ultrasonic imaging. *IEEE Trans. Med. Imag.* 14 (2), 283–291.
- Kass, M., Witkin, A., Terzopoulos, D., 1987. Snake: active contour models. *Proc. Intl. Conf. Comput. Vis.*, 259–268.
- McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey. *Med. Image Anal.* 1 (2), 91–108.
- Mulet-Parada, M., Noble, J.A., 1998. Intensity-invariant 2D+T acoustic boundary detection. In: *Proceedings of the Workshop on Biomedical Image Analysis*, June 1998.
- Muzzolini, R., Yang, Y., Pierson, R., 1993. Multiresolution texture segmentation with application to diagnostic ultrasound images. *IEEE Trans. Med. Imag.* 12 (1).
- Neave, H.R., Worthington, P.L., 1988. *Distribution-free tests* London, Unwin Hyman.
- Parker, A.D., Hill, A., Taylor, C.J., Cootes, T.F., Jin, X.Y., Gibson, D.G., 1994. Application of point distribution models to the automated analysis of echocardiograms. *Comput. Cardiol.*, 25–28.
- Rakotomamonjy, A., Deforge, P., Marche, P., 2000. Wavelet-based speckle noise reduction in ultrasound B-scan images. *Ultrason. Imag.* 22, 73–94.
- Sanchez-Ortiz, G.I., Wright, G.J.T., Clarke, N., Declerck, J., Banning, A.P., Noble, J.A., 2002. Automated 3-D echocardiography analysis compared with manual delineations and SPECT MUGA. *IEEE Trans. Med. Imag.* 21 (9).
- Shankar, P.M., 2000. A general statistical model for ultrasonic scattering from tissues. *IEEE Trans. Ultrasonic. Ferroelect. Freq. Contr.* 47 (3), 727–736.
- Staib, L.H., Duncan, J.S., 1992. Boundary finding with parametrically deformable models. *IEEE Trans. Pat. Recog. Imag. Anal.* 14 (11), 1061–1075.
- Tao, Z., 2005. Tunneling descent: a new strategy for active contour segmentation of ultrasound images, Ph.D. Thesis, Yale University.
- Tao, Z., Tagare, H.D., 2005. Level set tunneling descent segmentation of ultrasound images. *Proceedings of the 19th International Conference on Inform. Processing Med. Images*. In: LNCS, vol. 3565. Springer-Verlag, pp. 750–761.
- Tao, Z., Beaty, J., Jaffe, C.C., Tagare, H.D., 2002. Gray level models for segmenting myocardium and blood in cardiac ultrasound images. *Proceedings ISBI2002*, 265–268.
- Terzopoulos, D., Witkin, A., Kass, M., 1988. Constraints on deformable models: recovering 3D shape and nonrigid motion. *Art. Intell.* 36 (1), 91–123.
- Wagner, R.F., Smith, S.W., Sandrick, J.M., Lopez, H., 1983. Statistics of Speckle in Ultrasound B-scans. *IEEE Trans. Son. Ultra.* 30 (Jan), 156–163.
- Weng, L., Reid, J.M., Shankar, P.M., Soetanto, K., 1991. Ultrasound speckle analysis based on K-distribution. *J. Acoust. Soc. Am.* 89, 2992–2995.
- Wilson, D.C., Geiser, E.A., Larocca, J.J. 2000. Automated analysis of echocardiographic apical 4-chamber images. In: *Proceedings of the*

- SPIE, *Mathematical Modelling, Estimation, and Imaging*, San Diego, CA, vol. 4121, pp. 128–139.
- Winkler, G., 2006. Image analysis. In: *Random Fields and Markov Chain Monte Carlo Methods: A Mathematical Introduction*. Springer.
- Xu, C., Prince, J.L., 1998. Shapes, snakes and gradient vector flow. *IEEE Trans. Image Process.* (March), 359–369.
- Xu, C., Pham, D.L., Prince, J.L. 2000. Image segmentation using deformable models. *Handbook Med. Imag. 2* (Chapter 3).
- Zimmer, Y., Tepper, R. Akselrod, S. 2000. A lognormal approximation for the gray level statistics in ultrasound images. In: *Proceedings of the World Congress on Medical Physics and Biomedical Engineering*. Chicago, USA, 23–28 July 2000 (short paper no. 3304–50240 on CDROM).