

Deformable Boundary Finding in Medical Images by Integrating Gradient and Region Information

Amit Chakraborty, Lawrence H. Staib and James S. Duncan

Abstract— Accurately segmenting and quantifying structures is a key issue in biomedical image analysis. The two conventional methods of image segmentation, region-based segmentation and boundary finding, often suffer from a variety of limitations. Here we propose a method which endeavors to integrate the two approaches in an effort to form a unified approach that is robust to noise and poor initialization. Our approach uses Green's theorem to derive the boundary of a homogeneous region-classified area in the image and integrates this with a grey level gradient-based boundary finder. This combines the perceptual notions of edge/shape information with grey level homogeneity. A number of experiments were performed both on synthetic and real medical images of the brain and heart to evaluate the new approach and it is shown that the integrated method typically performs better when compared to conventional gradient based deformable boundary finding. Further, this method yields these improvements with little increase in computational overhead, an advantage derived from the application of the Green's theorem.

I. INTRODUCTION

SEGMENTATION of pertinent structures is of utmost importance for a variety of image analysis and visualization tasks. Examples of medical image analysis tasks that are of interest to us are multimodality image registration, structural measurement of anatomy, deriving priors for image reconstruction in another modality and cardiac motion tracking. However, robust identification and measurement of such naturally occurring deformable structures/objects is not always achievable by using a single technique that depends on a single image-derived source of information. The basic premise of this paper is that we believe that to be able to do this reliably, one needs to make use of integrated methods that make use of both boundary and region information.

Most segmentation methods can be divided primarily into region-based and boundary-based approaches. Region-based methods [15], [14], [30], rely on the homogeneity of spatially localized features such as grey level intensity, texture and other pixel statistics. Homogeneity does not necessarily mean identical pixel values within a particular region, rather it means that the variation within a region is of a smaller extent than that between regions. The advantage of such methods is that since they rely directly on the grey level image, they are less susceptible to noise than methods

that use derivative information. Also, if the high frequency information in an image is either missing or is unreliable, the segmented images remain relatively unaffected. However, the problem with typical region-based segmentation methods, such as the split and merge techniques [18], is that the resulting segmentation depends considerably on the choice of seed points and the region's shape is too dependent on the choice of the actual algorithm used. In addition, such methods often result in an over-segmented image. Rule based systems [6] do better, but are extremely application-specific. Other region-based methods either use probabilistic techniques [15], [27], [21], [5], [14] or use non-linear diffusion methods [30], [1] (see [14] or [30] for the exact mathematical relationship between these methods). The main idea behind these methods is to do what can be termed as *edge-preserved smoothing*. However, isolating objects from the resulting image still requires considerable effort as they also suffer from the problems of poor localization and over-segmentation (related to the problem of choosing the appropriate scale). In contrast to region-based methods, boundary methods primarily use gradient information to locate object boundaries. Low level boundary methods that perform edge detection [7], [25] typically result in false or broken edges. Thus, we are more interested in deformable whole boundary methods [33], [17], [2], [13], [8], [23] which rely on the gradient features at a subset of the spatial positions of an image (near an object boundary). While we have chosen the deformable shape-based parameterization of [33], we believe that our integration method can be used in conjunction with any of the other whole boundary approaches like the classical snakes method of [17] or with the topologically adaptable active contour methods like those of [28], [8], [23] or their variants like the reaction-diffusion approach of [34]. The latter methods are more adaptable to changes in the topology of the object under consideration, but they still base their decisions regarding boundary points on the value of the gradient and thus suffer from its noise sensitivity.

While both the region and boundary methods have their advantages and disadvantages, their problems are not necessarily identical, i.e. they are not affected in the same way by the various problems. While the presence of noise limits the performance of any image processing algorithm, the region-based methods are less affected than gradient-based boundary finding because the gradient is very noise sensitive. Also, if the high frequency information in the image either is missing or is unreliable, boundary finding is more error-prone compared to region-based segmentation. Shape variations, on the other hand, can be better handled

This work was supported in part by NIH-NHLBI grant R01-HL-44803 and NIMH-P01MH grant 49351-01.

A. Chakraborty was with the Departments of Electrical Engineering, and Diagnostic Radiology, Yale University, 333 Cedar Street, New Haven, CT, 06520-8042 USA. Currently he is with Siemens Corporate Research, Princeton, NJ 08540 (E-mail: chakrab@scr.siemens.com). L. Staib and J. Duncan are with the Departments of Diagnostic Radiology and Electrical Engineering, Yale University, 333 Cedar Street, New Haven, CT, 06520-8042 USA.

using a deformable boundary finding framework when we consider such variations to be generally around an average shape and such information can easily be incorporated as priors [33]. Further, since conventional boundary finding relies on changes in the grey level, rather than their actual values, it is less sensitive to changes in the grey scale distributions such as shading artifacts over images. Also, gradient-based methods in general do a better job of edge localization. This brings us to the realization that integrated methods are likely to perform better than either of the methods alone by being able to combine the complementary strengths of these individual methods, as pointed out in [29], [12].

A limited amount of previous work seeking to integrate region and boundary methods has been done. Among the available methods, AI-based techniques have been used where production rules are invoked for conflict removal [29]. Here, region growing is done first followed by a binary edge detection step. There are a few disadvantages to this procedure. First, a region classified image is often over-segmented due to the presence of noise. Thus, one needs a validation scheme to distinguish between true and false edges by looking at high gradient, continuity, etc. Also, such a scheme has no way of differentiating between multiple regions as it deals with the binary edge map obtained from the region segmented image. Further, such a method may suffer from the effects of poor edge localization as is often the case with region-based segmentation. This is also true of other similar efforts [12], [16] where rather than finding complex objects, the intention is in integrating region growing with edge-detection. Probability based approaches like [15], [27], [5], [14] or the nonlinear diffusion methods of [30], [1] achieve integration in the local or dense field sense. Another way of achieving local integration is the reaction-diffusion method of [34]. However, the problem of using such local integration methods is that if any one of the processes makes an error (e.g. a false edge), it is propagated to the final solution. Also, a decision regarding the final object boundary in [34] is made by considering the whole space of reaction-diffusion images and somehow choosing one result, something that can get very complicated. Finally, we note that the recent work of [31], [36] has similar motivations as ours even though the final implementation is similar to the weak membrane method of [5], [27]. Also, the final solution in [36] lacks the notion of shape because integration is primarily carried out in a local pixel-wise sense.

Our ultimate goal is to develop a fully bi-directional framework for integrating boundary finding and region-based segmentation. This would lead to a system where the two modules would operate in parallel so that at each step the outputs of each model get updated using information from the outputs of both the models from the previous iteration. Our initial effort presented in this paper is aimed at using the results of region-based segmentation to assist boundary finding and is one portion of the complete system. Unlike most of the other existing methods, it in-

tegrates region-based segmentation with boundary finding rather than edge detection, thereby making it resistant to the problems of false and broken edges while at the same time having well localized boundaries. The ability to use model-based information such as shape also allows it to compensate for missing and noisy information. This use of shape is particularly advantageous for medical images where the interesting structures vary around an average shape, characteristic of that task or application. Further, since there are no intermediate steps such as line following involved, it is computationally attractive as well.

II. REGION-BASED SEGMENTATION

The main idea here is to classify a particular image into a number of regions or classes. Thus for each pixel in the image it is necessary to decide or estimate which class it belongs to. There are a variety of approaches to region-based segmentation. In the discussion below, we briefly describe the approach we have used which is adapted from the efforts of [24] and [20]. It has the capability of classifying images into distinctive regions based on either their grey level values or for a more general case, based on their texture properties. We consider textured images more general because grey level images can be considered to be the most simple form of texture images [32]. The only assumption is that the variation of the considered feature (either grey level or texture) within a region is of a smaller magnitude than the variation between regions.

The intensity image is modeled as a Gaussian Markov random field (GMRF). This model has been used by many researchers to model texture and other image characteristics [32], [24], [11]. It models the conditional probability of the image intensity given the classification.

Let S denote the $M \times M$ image lattice, i.e. $S = \{(i, j), 1 \leq i, j \leq M\}$. Let $\{L_s, s \in S\}$ and $\{Y_s, s \in S\}$ denote the labels and the zero mean array obtained from the image data respectively. Note that the labels can belong to only a certain number of region classes. This number is prespecified. Let N_s denote the neighborhood of a site s (a first order neighborhood consists of only four neighbors, eight in the case of a second order system and so on). Now, if one further assumes that all the nearest neighbors of s also have the same label as s , one can write the following expression for the conditional density of the intensity at the pixel site s [24], [32]:

$$P(Y_s = y_s | Y_r = y_r, r \in N_s, L_s = l) = \frac{\exp(-U(Y_s = y_s | Y_r = y_r, r \in N_s, L_s = l))}{Z(l | y_r, r \in N_s)} \quad (1)$$

where $Z(l | y_r, r \in N_s)$ is the partition function of the conditional Gibbs distribution, and

$$U(Y_s = y_s | Y_r = y_r, r \in N_s, L_s = l) = \frac{1}{2\sigma_l^2} \left(y_s^2 - 2 \sum_{r \in N_s} \Theta_{s,r}^l y_s y_r \right) \quad (2)$$

In (2), σ_l and Θ^l are the GMRF model parameters of the l^{th} region class. Without going further into the details,

one can see that for a first order GMRF model, if $\hat{\mu}$ is the mean and $\hat{\sigma}^2$ is the variance of the sub-image under consideration, then the feature vector for each region is denoted by (see [32], [24] for further details):

$$F = (\theta_1, \theta_2, \theta_3, \theta_4, \hat{\mu}, \hat{\sigma}^2) \quad (3)$$

where the θ 's correspond to the influence from the four neighbors that make up the first order neighborhood. These parameters for each region are estimated by a least squares estimate (see [32] for details) method using a window around a user specified point, representative of that particular region.

Once the intensity image Y^* has been modeled, the next task is to determine the classification. This is achieved by maximizing the posterior (the Maximum *a posteriori* (MAP) method) distribution of the region labels given the intensity image:

$$P(L|Y^*) = \frac{P(Y^*|L)P(L)}{P(Y^*)} \quad (4)$$

where L corresponds to the classified image with L_s describing the label at the s^{th} pixel. The label field L is modeled as a first or second order MRF, which says that $P(L_s|L_{S/s}) = P(L_s|L_r, r \in \hat{N}_s)$ where $L_{S/s}$ is the whole label field excluding the site s , and \hat{N}_s denotes the appropriate neighborhood of site s . It acts as a prior that emphasizes the property that neighboring pixels of the classified image share the same label (see [20] for details). Maximizing (4) gives an optimal Bayesian estimate. In our results, we used the coordinate-wise descent method of [20], similar to the iterated conditional mode (ICM) algorithm [4]. Here, one starts with an initial labeling, L^0 and then iterates sequentially (or in parallel) through *each pixel, replacing the current label at that pixel with the label that maximizes $P(L_s|Y^*, L_{S/s})$* , where $L_{S/s}$ represents the label set for the whole image except the site s , and the other symbols have the same connotations as before. Without going into the details (see [24], [20]) it can be shown that this is equivalent to maximizing $\forall s \in S$

$$\arg \max_{L_s} P(L_s|Y^*, L_{S/s}) = P(L_s|Y^*, L_r, r \in \hat{N}_s) \quad (5)$$

which is reduced to determining

$$L_s^* = \arg \max_{L_s} P(y_s^*|Y_{S/s}^*, L_s, L_r, r \in \hat{N}_s)P(L_s|L_r, r \in \hat{N}_s) \quad (6)$$

for each of the sites s using the ICM algorithm.

The functional in (6) is optimized at each step where the first term, the conditional probability term is given by (2) and the second term is the prior term given by the MRF assumption. To summarize the algorithm, at each iteration, each pixel is updated to maximize (6) using the data at site s and the neighborhood, and the current estimate of the labels of the neighborhood. This continues so long as the number of changes is above a certain fraction of the image size.

We note here that the computational complexity of this region segmentation step will depend primarily on the number of parameters used to characterize the GMRF model. For most images (except where texture is important), the only parameters of interest are the mean and the variance (thus we don't use the θ 's in (3)). If we use this reduced set of parameters, it greatly speeds up the process.

III. BOUNDARY PARAMETERIZATION

As already mentioned, we are primarily interested in finding a class of objects with smooth boundaries that are continuously deformable. Simple edge detectors are of limited or no use under such circumstances as the output of edge detectors does not necessarily correspond to object boundaries except for very high quality images. Hence we are concerned with a whole boundary method so that a global structure can be imposed on the problem. A variety of deformable contour methods have been proposed (see the introduction for a brief review). In this paper, we use the approach due to [33] which uses a Fourier parameterization for the boundary. It expresses a curve in terms of an orthonormal basis [19], [22]. It makes the estimation problem easier as the cross correlation matrix becomes diagonal. We note that for most practical situations we constrain ourselves to a limited number of harmonics. Thus,

$$v(t) = \begin{pmatrix} x(t) \\ y(t) \end{pmatrix} = \begin{bmatrix} a_0 \\ c_0 \end{bmatrix} + \sum_{k=1}^{\infty} \begin{bmatrix} a_k & b_k \\ c_k & d_k \end{bmatrix} \begin{bmatrix} \cos(kt) \\ \sin(kt) \end{bmatrix} \quad (7)$$

where, $v(t)$ is the contour vector consisting of the x and y coordinates and a_k, b_k, c_k and d_k are the corresponding Fourier coefficients.

The curve is thus represented by

$$\vec{p} = (a_0, c_0, a_1, b_1, c_1, d_1, \dots) \quad (8)$$

The estimation of these parameters to find the boundary is posed as an optimization problem, where, the objective function measures the strength of the boundary given the set of parameters from the image. We shall come to this issue in the following section. We would also like to point out here that we chose to use the Fourier parameterization because in our view, it is suitable for the range of objects that we are interested in and because of its ability to incorporate prior shape information effectively. Further, the version developed in [33] can be made invariant to starting point, scale and 2-D rotation and translation, thus making the whole process view independent. However, we believe that our integration method can be used in conjunction with any of the other whole boundary approaches like the classical snakes method of [17] or with the topologically adaptable active contour methods like those of [28], [8], [23].

IV. INTEGRATION

The basic premise of this work is that the integration of region and boundary information will result in an improved

algorithm for solving the segmentation problem. In this section, we first discuss the motivation behind our approach and then describe it mathematically.

As input to the integration algorithm, we have both the actual image I and the region classified image I_r , which is obtained from I after passing it through the region segmentation step discussed before. This information is introduced as an added prior into the gradient-based deformable boundary finding framework. In its simplest form, this region term enforces the boundary to enclose a single region in I_r . As we shall later see, the assumption of a single region is not strictly necessary, but for the sake of simplicity, we will continue with it and modifications to it will be discussed later. The traditional boundary finding problem does not use the original image directly, but uses instead the gradient image I_g . As in the approach of [33], we shall use the magnitude of the gradient vector at each pixel location. We have also used the gradient direction along with the magnitude (see [35]). The gradient image can be obtained by convolving I with the derivative of a Gaussian. We take the magnitude of the corresponding vector image to get I_g . Hence, the input to the integration system is the gradient image I_g and the region classified image I_r .

We shall pose the boundary estimation problem using gradient and region homogeneity information in the maximum *a posteriori* framework. Our aim is to maximize $P(\vec{p}|I_g, I_r)$, where as described in the previous section, \vec{p} is the vector of parameters used to parameterize the contour. Now,

$$P(\vec{p}|I_g, I_r) = \frac{P(\vec{p}, I_g, I_r)}{P(I_g, I_r)} \quad (9)$$

$$= \frac{P(I_r|I_g, \vec{p})P(\vec{p}, I_g)}{P(I_g, I_r)} \quad (10)$$

$$= \frac{P(I_r|I_g, \vec{p})P(I_g|\vec{p})P(\vec{p})}{P(I_g, I_r)} \quad (11)$$

Thus, ignoring the denominator which does not change with \vec{p} our aim is to determine,

$$\begin{aligned} \vec{p}^* &= \arg \max_{\vec{p}} P(\vec{p}|I_g, I_r) \\ &\propto \arg \max_{\vec{p}} P(I_r|I_g, \vec{p})P(I_g|\vec{p})P(\vec{p}) \end{aligned} \quad (12)$$

or,

$$\begin{aligned} \arg \max_{\vec{p}} P(\vec{p}|I_g, I_r) &\equiv \arg \max_{\vec{p}} [\ln P(\vec{p}) + \ln P(I_g|\vec{p}) \\ &\quad + \ln P(I_r|I_g, \vec{p})] \end{aligned} \quad (13)$$

In the last equation (13), we have just taken the natural logarithm, which is a monotonically increasing function. Knowledge of I_g may be used to calculate I_r , through the use of line processes [15], [5]. However, if we do not use that information, we are effectively discarding information rather than assuming extra information. Thus, finally, the above can be written in the following form:

$$\begin{aligned} \arg \max_{\vec{p}} M(\vec{p}, I_g, I_r) &= \arg \max_{\vec{p}} [M_{prior}(\vec{p}) \\ &\quad + M_{gradient}(I_g, \vec{p}) + M_{region}(I_r, \vec{p})] \end{aligned} \quad (14)$$

We now explain each of the terms in (14).

Prior shape term: The first term in equation (14) corresponds to the prior shape term. This can be obtained from previous experience or from expert knowledge about the natural structures that exist in the image. When it is non-uniform, it biases the model towards a particular range of shapes. Use of prior boundary information can especially be important in clinical applications such as delineating the corpus callosum or the endocardium of the heart, where the nature of the shape does not change a great deal from individual to individual even though the exact reconstruction is different. Prior boundary information under such circumstances can be obtained from previously outlined boundaries. These are first parameterized and then the mean and variance of these parameters are calculated to obtain a multivariate Gaussian prior, which could then be used to constrain the optimization at a later stage for further instances of similar structures (see [33] for more details). Since there might be other objects in the image, we would always need an initial estimate of the boundary to start the optimization process. The prior shape, if available, can be used for this purpose as well. The prior information could be of particular importance if the boundary is ill-defined.

Gradient term: The second term in equation (14) is the likelihood term, which depends on the gradient image. It is a measure of the likelihood of the contour of the described object being the true boundary, once the parameters defining the boundary are given. This is expressed in terms of a contour integral, where the integral is computed over $C_{\vec{p}}$, the curve described by the boundary $(x(\vec{p}, t), y(\vec{p}, t))$. At each point on the contour, the strength of the boundary is evaluated by the magnitude of the gradient at that particular point, given by the gradient image. Thus, the likelihood of the whole contour being the true boundary becomes proportional to the sum of the magnitude of the gradients at all the points that lie on the boundary. If we assume that the noise can be approximated by a zero mean Gaussian, and further assume that the pixels on the boundary are independent, then we may express the above term in the probability expression as the following line integral ([33]) where σ_n^2 is the variance of the underlying noise process (k_1 being a constant):

$$M_{gradient}(I_g, \vec{p}) = \frac{k_1}{\sigma_n^2} \int_{C_{\vec{p}}} I_g(x(\vec{p}, t), y(\vec{p}, t)) dt \quad (15)$$

The derivation of the above equation is straightforward and can be found in [33].

Region term: For every image, the region-based segmentation step is carried out only once using the method described in section 2. The last term (of equation (14)) incorporates this region information into the boundary finding framework. As we have already mentioned, we would like the boundary to enclose within it a homogeneous or slowly varying (w.r.t. grey level or texture) region. In other words, we expect the interior of the boundary to be filled with the region of a single class. This realizes the expectation that the variations within an object are assumed

to be smaller than those between objects. We emphasize here that this similarity measure may be based on image intensity, texture properties or any other image attribute that can be handled by the region process. However, this method can also deal with situations like outlining an annulus, where the central part might belong to a separate class. To describe this term mathematically, let us consider the




Fig. 1. An example demonstrating the use of region information in the boundary finding process

situation in Figure 1 showing a hypothetical region classified image and the boundary at any instance during the optimization process. For simplicity's sake, let us assume that there are only two regions and we want to segment out the object in the center from the surrounding background. As we can see, the boundary does not include the entire central region. At some points it includes the background and at other places it excludes points that belong to the central region. We would like to penalize cases where points from the surrounding regions are included and would like to reward if more and more of the central region is included. A very simple way of doing this would be to do the following. Let us consider for this simple image, all points that lie inside the central region to have value 1.0 and all points that lie outside a value -1.0 . Now, if we sum up the values of all the points that are inside the contour, it is clear that the sum achieves a maximum when the contour is so placed such that it includes all of the points with a value 1.0 and excludes all of the points that have a value of -1.0 , that lie on the surrounding region.

If more than two regions are involved, all pixels of the region that needs to be segmented can be assigned a positive value and the remaining ones negative values, the magnitudes of which reflect how much one expects the target region to be dissociated from the remaining regions. Hence, remote regions are expected to have high negative values, representing a larger penalty for including remote points. This way multiple regions can be handled. For an annulus, e.g. in a transaxial heart image, it does not matter what value the points have in the region within the inner boundary of the annulus. We can also relax the assumption that the interior of a region needs to be a connected region. If we know that the target object consists of more than one region, then all those regions are assigned positive values and those that lie outside are assigned negative ones. Thus, the only requirement is that the points immediately inside

and outside of where we expect the boundary to be located should be different.

For a mathematical description of the term, assume that the prior conditional distribution of $I_r|\vec{p}$ be given by an exponential of the form $P(I_r|\vec{p}) = \frac{1}{Z} e^{\int_{A_{\vec{p}}} I_r dA}$ where Z is a normalizing constant given by $Z = \int_{\vec{p} \in \mathcal{P}} e^{\int_{A_{\vec{p}}} I_r dA} d\vec{p}$, \mathcal{P} is the range of p and $A_{\vec{p}}$ is the area bounded by the contour described by the parameter vector \vec{p} . Thus,

$$M_{region}(I_r, \vec{p}) \propto \int \int_{A_{\vec{p}}} I_r(x, y) dA \quad (16)$$

Simplified Integration using Green's theorem: The objective function involving all the three above terms can be expressed as:

$$\begin{aligned} \arg \max_{\vec{p}} M(\vec{p}, I_g, I_r) &= \arg \max_{\vec{p}} [M_{prior}(\vec{p}) \\ &+ M_{gradient}(I_g, \vec{p}) + M_{region}(I_r, \vec{p})] \\ &\equiv \arg \max_{\vec{p}} \left[M_{prior}(\vec{p}) + K_1 \int_{C_{\vec{p}}} I_g(x(\vec{p}, t), y(\vec{p}, t)) dt \right. \\ &\left. + K_2 \int \int_{A_{\vec{p}}} I_r(x, y) dA \right] \quad (17) \end{aligned}$$

where K_1 and K_2 are the weighting constants which signifies the relative importance of the two terms in the above equation. We will discuss this issue later.

Of the last two terms in (17), one is a line integral and the other is an area integral. In general, computing a line integral is much less expensive compared to an area integral. Thus we would save a lot of computation, especially when we carry out an iterative optimization procedure, if we could convert the area integral to a line integral. We already compute a line integral in the second term which is present in the original boundary finding method. The above conversion can be obtained using Green's theorem [3] as follows. Let,

$$M_r(x, y) = \int_0^x I_r(z, y) dz; \quad N_r(x, y) = - \int_0^y I_r(x, z) dz \quad (18)$$

Hence, using Green's Theorem,

$$\int \int_{A_{\vec{p}}} I_r(x, y) dA = \frac{1}{2} \int_{C_{\vec{p}}} [N_r(x, y) \frac{\partial x}{\partial t} + M_r(x, y) \frac{\partial y}{\partial t}] dt \quad (19)$$

Thus, finally we have,

$$\begin{aligned} \arg \max_{\vec{p}} M(\vec{p}, I_g, I_r) &= \arg \max_{\vec{p}} [M_{prior}(\vec{p}) + \int_{C_{\vec{p}}} (K_1 I_g(x, y) \\ &+ K_2 \{N_r(x, y) \frac{\partial x}{\partial t} + M_r(x, y) \frac{\partial y}{\partial t}\}) dt] \quad (20) \end{aligned}$$

In this section we have presented a deformable boundary finding procedure that introduces a prior term which incorporates information that we obtain from region-based segmentation. Further, using Green's theorem, we reduced

the whole problem to computing line integrals only rather than both line and area integrals. Since, M_r and N_r are evaluated only once for an image, and is not repeated at every iteration, the computational speedup due to the use of the Green's theorem formalism is from $O(\text{target area})$ to $O(\text{target perimeter})$. For most objects this difference is quite significant because the number of pixels within the target object is generally much larger than the number of sample points on the contour.

It is essential to point out here that in practice, the functions M_r and N_r are evaluated by simply summing the pixel values of I_r in either the x or y direction. Thus, $M_r(i, j) = \sum_{k=0}^i I_r(k, j)$ and similarly along the other coordinate for $N_r()$. Hence, even though I_r is discontinuous, $M_r()$ and $N_r()$ are continuous in the x and y direction respectively. Further, the derivatives of these functions which are necessary for the gradient calculation yield I_r . Thus, for most of the computations, these differentiations are not being carried out numerically due to the way the objective function is constructed. We emphasize here that at no stage are we taking derivatives of I_r , which is discontinuous. Optimization is achieved using the conjugate gradient method. The process is terminated when the change in the objective function falls below a pre-determined threshold.

V. RESULTS

Experiments were carried out both with synthetic and natural images to verify the performance of the above method.

Synthetic Images: The output of all these experiments are object boundaries. We have set up experiments using synthetic images to quantitatively evaluate the method.

For our purposes, we created a synthetic image that has one target object in the center surrounded by a background as shown in Figure 2(a). The target object has both convex and concave parts to it. Further, it also has some high curvature points that make the boundary finding process non-trivial. To make it even more difficult, we have smoothed the image using a Gaussian kernel so that the edges become fuzzy. On top of that, white noise was added which would again affect the boundaries the most as they were smoothed. The result is a considerably lower SNR at the boundaries that at any other location in the image. Thus, this procedure represents many of the difficulties associated with structure segmentation in biomedical images. No prior information was used for these images. The region-based segmentation was carried out using only the mean and the variances of the two regions under consideration.

For evaluation we used the following procedure. First, the true boundary was evenly sampled into 256 points. The boundary finding process was initialized with a boundary that is spatially distant from the true boundary and results in an output boundary. For all the examples using synthetic images, we used 8 harmonics (see Discussion section also). This result was then evenly sampled into the same number of points. To find out how closely the output boundary approximated the true boundary, we need to cal-

culate the distance between them. To solve the problem of pointwise correspondence, we keep the true boundary fixed, and vary the starting point of the other boundary point by point, calculating at each step the total distance as a sum of the distances between each corresponding points. The minimum of these is considered to be the reported value of the distance. The comparisons were done using three versions of the objective function. When only the gradient-based term in the objective function was used, we have the traditional gradient-based boundary finder. The second method only uses the region-based term, where information only from the region classified image is used. Finally, the proposed method uses both of the above terms.

First, we compare the performance of the three methods under varying amounts of noise. For each noise level, the three methods were executed using the same initial boundary placement. The three methods were allowed to run for the same number of iterations or until the change between successive iterations fell below a certain predetermined level.




Fig. 2. Performance of the three methods for a noisy image (SNR=2.0). In the last three images, the more gray contours represent the initial boundaries and the brighter ones the final derived boundaries. (a)Top,Left: Original image; (b)Top,Right: Region classified image; (c)Middle,Left: Output of Boundary finding using gradient information only; (d)Middle,Right: Output of Boundary finding using region information only; (e)Bottom: Output of Boundary finding using the integrated approach.

Figure 3(a) shows a comparison of the three methods for

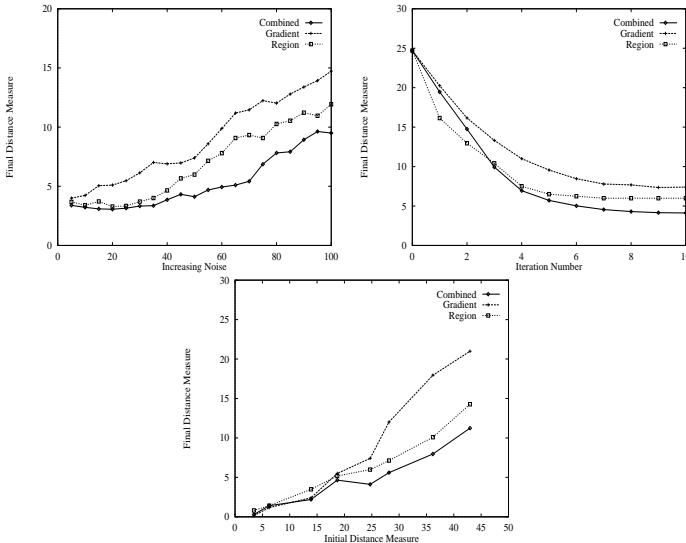


Fig. 3. Performance plots (a)Top,Left: Performance under noise using the three methods. The SNR varies from 4.0 to 1.0; (b)Top,Right: Convergence speed using the three methods (SNR=2.0); (c)Bottom: Performance under different Initializations (SNR=2.0)

this example. The Y-axis gives a measure of the distance between the approximated contour and the true one. (This is a scaled version of the square of the distance between the contours calculated as described above.) In this and all the following plots, 'Gradient' represents the gradient-based method, 'Region' corresponds to the boundary found based only on the region classified image, and 'Combined' refers to the proposed integrated method. As we might expect, the performance of all the three methods worsen as we increase the amount of noise. However, as can be seen, the combined method is the least sensitive to noise. Also, it performs much better compared to the traditional gradient-based method. Since the gradient-based method relies upon the first derivative, it is more susceptible to noise than the region-based method, which uses homogeneity within the image. On the other hand, the combined method seems to give uniformly better results.

In Figure 2, we show the result for a particular value of the noise given by SNR=2.0. The initial boundary is considerably displaced and disfigured when compared to the actual boundary. We see that the combined method produces the best result.

The next part of the experiment is to see how fast the methods converge. Figure 3(b) shows the situation when the image used has a noise level given by SNR=2.0 and has the same starting position shown in Figure 2. Clearly, the combined method performs better than the other two, especially compared to the gradient-based boundary finding. Due to the global nature of the region term, initially the convergence is faster for the region-only method. But once it comes close to the actual boundary, due to the better localization property of the integrated method (which comes from the gradient term), the integrated method outperforms the region-only method. As for the gradient-only

method, the effect of noise and fuzzy boundaries resist it from getting the best match. Thus not only the results are better (as we concluded from the previous experiment), but also the convergence is faster. The next plot, Figure 3(c) shows the situation under different starting positions. The initial contours vary from each other not only by translation and rotation, but also in terms of their exact shapes. We can observe that when the initial contour is close to the actual one, there is very little difference among the three methods. However, when the initialization is far enough, the integrated method does considerably better than the gradient-only method. The region-only method comes close, the difference being due to the better localization property of the integrated approach. It must be noted that for the results shown in Figures 3, the experiments were repeated ten times under exactly the same settings, the only variation being that the noise sample is different even though the noise distribution is the same. The reported results were obtained by taking the average of the above which filtered out minor variations that might have arisen. One needs to do this because even though the noise distribution remains the same, the pixel values of the noise could be different.

Now, we consider a situation where one or both of the methods might produce unrealistic results. To investigate such a situation, we devised an experiment with a synthetic occlusion. In Figure 4(a), in a part of the image the intensity is lowered to a level slightly above the background. However, as it touches the background there is a further fall in the intensity level. For this part, the SNR=1.5, and for the rest of the image, the SNR=3.0, where the SNR is calculated in a similar fashion as in the previous case. In addition, as in the previous case, we smooth the boundaries and add noise to it. This situation may happen in a medical X-ray image where an object is just behind part of another object of interest. Partial volume effects might also result in such spatial variations. Now, if we perform region-based segmentation under the assumption that there are two regions, then that part of the central region beyond the artificial occlusion boundary will be classified with the background as can be seen in the figure. Thus, as far as the region-based segmentation is concerned, it sees an occlusion. Now, if we base the boundary finding on this region classified image, it would tend to pull its boundary to the occlusion boundary. For those parts where there is no occlusion, it would still work. For those parts where there is occlusion, it would make gross mistakes resulting in a huge overall error. The gradient-based method doesn't see much of a difference as it searches for a local gradient maxima, which is still there. Since the initialization is closer to the actual boundary than the occlusion boundary, the gradient-based method ignores the occlusion boundary. Thus, even if one of the methods fail partly, the integrated approach still gives reasonable results as it uses information from both the methods. This method can potentially be used in conjunction with methods to handle occlusions (see [26] for example, which preserves depth information)




Fig. 4. Performance of the three methods for an occluded image. In the last three images, the more gray contour represents the initial boundary placing and the brighter one the final derived boundary. (a)Top,Left: Original image; (b)Top,Right: Region segmented image; (c)Middle,Left: Output of Boundary finding using gradient information only; (d)Middle,Right: Output of Boundary finding using region information only; (e)Bottom: Output of Boundary finding using an integrated approach.

for even better results. If the initialization were close to the occlusion boundary, of course, none of the methods would work and for all the cases, including the integrated method, we would get grossly wrong results. It is precisely in situations like this that the importance of the prior boundary information can be realized.

Real Images: In this section we apply the algorithm to real clinical images. Figure 5a shows a short axis Magnetic Resonance (MR) image of a dog heart. The aim is to outline the endocardium. A simple region-based segmentation using only the means and the variances was done and the output is shown in Figure 5b. Thus, once again, the user selected three points (by mouse clicks), one inside the endocardium, one on the epicardium and another on the outer background. The means and variances of these regions were calculated by using as least squares method on the window of pixels surrounding these user-selected points. Also, as with the synthetic images, we used 8 harmonics to describe the boundary. The initial boundary is shown in Figure 5c. Figure 5d shows the probable boundary as outlined by an expert. We display this only for the sake of comparison. Once again no prior information




Fig. 5. Example with an MR heart image. (a)Top,Left: Original image; (b)Top,Right: Region classified image; (c)Middle,Left: Original image with the initialization as marked by a user; (d)Middle,Right: Original image with the contour of the endocardium drawn by an expert; (e)Bottom,Left: Output of Boundary finding using gradient information only; (f)Bottom,Right: Output of Boundary finding using an integrated (gradient + region) approach using the same initialization.

was used. In all the experiments, we used the same initialization and the same number of harmonics to describe the boundaries for both the methods. As we can see, the image quality is very poor and the edges seem to be very fuzzy. Thus the gradient information is weak. If we apply gradient-based boundary finding, due to a lack of strong edge information, the boundary will diverge after a few iterations, as shown in Figure 5e. On the other hand, as we can see the region information in itself is also not very good. Figure 5f shows the results of the integrated method, which though not perfect, is much better compared to the other method. The main reason for this improved performance is that neither region-based segmentation nor gradient-based boundary finding will actually fail as there is some information in both the gradient and the region classified image. By themselves, they have the limitations previously described and thus neither method alone produces desirable results. Once we combine them, the output improves due to the information fusion, which relieves some of the limitations found when using the algorithms separately.

Figure 6 shows a similar sequence using a mid-sagittal MR brain image, where the task is segmenting the cor-




Fig. 6. Example with an MR brain image. (a)Top,Left: Original image; (b)Top,Right: Region classified image; (c)Middle,Left: Original image with the initialization as marked by a user; (d)Middle,Right: Original image with the contour of the corpus callosum drawn by an expert; (e)Bottom,Left: Output of Boundary finding using gradient information only; (f)Bottom,Right: Output of Boundary finding using an integrated (gradient + region) approach using the same initialization.

pus callosum. Here, we used 10 harmonics to describe the boundary. Once again, the expert drawing is used only for the sake of comparison. Region-based segmentation was done for three regions: the corpus callosum, the grey matter and the background (that includes the CSF) using only the means and the variances (in the same fashion as before).

In Figure 7, the task is to outline the epicardium on an MR image of a dog heart. Since the epicardium, is considerably smoother than the endocardium we used only 5 harmonics to describe the boundary. As in the example for the endocardium, region-based segmentation was done for three regions: endocardium, epicardium and the background using only the means and the variances, using a similar procedure as described for the other images (see Figure 7(b) for the region-segmented result). This is a particularly difficult image because the right ventricle is not at all conspicuous and thus only an expert eye can point out the epicardial boundary as shown in Figure 7c. Without any constraints, the output of a normal gradient-based boundary finding algorithm is as shown in Figure 7e. One way of introducing constraints, where otherwise




Fig. 7. Example with an MR heart image to outline the epicardium; (a)Top,Left: Original image; (b)Top,Middle: Region classified image; (c)Top,Right: Initial boundary; (d)Bottom,Left: Original image with the contour of the epicardium drawn by an expert; (e)Bottom,Middle: Output of Boundary finding using gradient information only without any constraints; (f)Bottom,Right: Output of Boundary finding using an integrated approach using region-based constraints and under the same initialization.

there are no image features, is to mark out areas beyond which the boundary should not go. This can simply be done in the present framework by negating the pixels of the classified image in these forbidden areas. The forbidden region removes the right ventricle (parts of which can be differentiated in the region classified image, and essentially, the forbidden region just connects them). In this example, a region was roughly pointed out as forbidden for the epicardium by the user depending on a rough estimate of where the epicardial wall separates the right ventricle. Using this procedure, the output of the integrated approach is as shown in Figure 7f.

Next, we consider the case of an ultrasound image of the heart as shown in Figure 8(a), where the task is to outline the left ventricle. This is an ideal example where the texture and not just the pixel intensity is important. Region classification was done using texture properties and is shown in Figure 8(b). Here, three regions were chosen where the user selected a point, one inside the left ventricle, one in the surrounding myocardium, and one more in the background. As usual, a window was selected around the user-defined points, and then the mean, variance and




Fig. 8. Example with an ultrasound image of the heart. (a)Top,Left: Original image; (b)Top,Right: Region classified image; (c)Middle,Left: Initial boundary as marked by a user; (d)Middle,Right: Output of a gradient-based boundary finder. That it diverges is clear; (e)Bottom: Output of an integrated (gradient + region) boundary finder which looks quite reasonable.

four other GMRF parameters were estimated using a least squares procedure as mentioned in section 2, and described in detail in [24]. A gradient-based method by itself is not likely to produce good results. Thus, as shown in Figure 8(d), using a gradient-based boundary finder leads us nowhere. On the other hand, the boundary finder in the integrated case as shown in Figure 8(e) does not diverge and the result looks reasonable.

Reproducibility: Since there is still some human operator interaction required to use our proposed algorithm, we present results aimed at testing the reproducibility of the boundaries generated by the algorithm. We test the algorithm’s reproducibility against results generated by a human operator (i.e. manual tracing). We hypothesize that our approach will have a smaller variance, implying better reproducibility.

First, a bank of ten images were selected. These MR images constitute a canine heart from the apex to the base. For each one of these images, the algorithm was executed ten different times to find out the boundary of the endocardium from different initial settings as provided by an human operator. Note, that for the initialization, the op-

erator quickly clicked a few points around the boundary. However, for manual tracings, the human operator carefully traced the boundary. The output of the algorithm consisted of the boundary of the endocardium. The boundaries were uniformly sampled to always have the same number of points. The mean and the variance in the contour position was calculated for each one of these images. The variance was then divided by the number of points on the contour.

For the manual tracings by a human operator, a domain expert traced carefully the boundaries for each of these images ten times and the same procedure as above was repeated. To minimize the effect of memorizing, the order of the images were randomized and the tracings were considerably spaced out in time. The results are shown in Table 1. Figure 9 shows the mean contours (human and algorithm




Fig. 9. Four of the images used in the *Reproducibility* experiment with the mean contour as hand traced by an operator, and as found by the algorithm overlaid. The broken line represents the hand-traced contour and the solid line the algorithmic procedure. (Please see the text for explanations.)

generated) overlaid on four of the images, where the broken contour represents the human drawn and the solid one the algorithmic version. As we can see from the table, the variance is always at least twice for the human tracings when compared to the algorithmic results. This constitutes a significant deviation in reproducibility between the two methods. That it is indeed the case, is borne out by a standard pairwise *T-test* on the variance data which shows that the difference between the two cases is highly significant ($p < 0.001$).

We note again that this experiment clearly shows two important things. First, as demonstrated in Figure 9, the results from the algorithm and the human tracings are in close agreement. Second, the result of the *T-test* shows that the variability in the algorithm is significantly smaller when compared to human tracings. Thus, the algorithm produces results that reasonably agree with a human ex-

pert’s desired result, but provides a much more stable (i.e. less variable) estimate of the boundary location.

Image Number	Algorithm variance ($pixel^2$)	Human variance ($pixel^2$)
1	0.631	1.47
2	0.5	1.50
3	0.576	1.371
4	0.67	1.7
5	0.53	1.85
6	0.72	1.81
7	0.75	1.71
8	0.83	1.99
9	0.24	1.4
10	0.576	2.1

TABLE I
VARIATIONS USING HUMAN HAND DRAWN AND AUTOMATED TRACINGS.

VI. DISCUSSION

While the algorithm works under a wide range of parameter settings, to get an optimal performance, it is important to choose them in a rational way.

Choice of the number of Harmonics: The performance of the algorithm depends considerably on the proper choice of the number of harmonics, especially if the data are noisy, because this choice constrains the range of shapes that the algorithm can allow, thereby acting as a regularizer. A large number of harmonics allows the algorithm to produce boundaries representing a large range of shapes, but makes the optimization procedure computationally long, and in some cases might even result in divergence. However, most of the biological forms that we are interested in are relatively smooth and un-convoluted. Hence, they can be represented by a limited number of harmonics. We may choose to have a larger number of harmonics for more complicated structures at higher computational cost. A principled selection of the number of harmonics necessary, is made by calculating the modeling error for the type of objects under investigation and choosing the least number (of harmonics) that produces an error smaller than a pre-determined limit (see also [33]). Here, in this paper, we used 8 harmonics for the synthetic images and the endocardium example. For the corpus callosum, we used 10 harmonics. However, for the smoother epicardium, only 5 harmonics were used.

Choice of the number of regions in the image: The region-based segmentation method discussed here doesn’t automatically determine the number of regions the image pixels need to be classified into. Instead, we use an supervised segmentation procedure, where the user determines the relevant number of classes. We used two considerations to make an appropriate choice. On one hand, we wanted to choose a minimum number of classes that would provide a reasonable region segmentation. Also, we had in consideration the fact that our final outcome was the targeted structure. Thus, we chose the number of classes in

such a way that the targeted structure is well differentiated from the neighboring tissue. As long as this criteria is met, even if the segmentation of the background is sub-optimal, the final outcome is not effected. Some knowledge regarding the anatomy of the imaged region can be of considerable help in making an educated guess. From our experience, as we have noted in the experiments, three regions for the heart images (endocardium, epicardium and the background) and three for the brain (grey matter, white matter and the CSF) images were adequate.

Choice of the relative weighting between the gradient and the region term: We performed some theoretical and practical analysis to help us make this selection [10]. From that study, our observation was that a reasonable weighting scheme would be to use $K_1 = 1$ and $K_2 = 1/SNR$ in equation (20) i.e. give more weight to the region-based term as the noise increases. For practical images, one can make this choice by first making a rough noise analysis of the given image. This can easily be done using the region properties which are being calculated during the region segmentation step.

Computational costs: The appropriate use of Green’s theorem assures us that the increase in the final optimization time that produces the final boundary is minimal. However, we note that the integrated algorithm is still computationally more involved than the original gradient-based boundary finding. This is so because after all, the region-based segmentation needs to be done as an additional step. But, as mentioned before and as we saw in the examples, for most images, it can be simplified by only using the means and the variances of the regions involved rather than the whole set of parameters. On an average, it takes about a minute to obtain the region classification on a Sun Sparc10 workstation. Another 1.0-1.5 minutes are necessary for the boundary optimization. However, this depends on the number of regions involved, the complexity of the boundary, etc. Also, a better optimization of the code is likely to make the procedure faster.

VII. CONCLUSIONS

We have presented in this paper a new technique for integrating gradient and region information within the deformable boundary finding framework. We note again that the region homogeneity constraint does not necessarily mean regions of constant value. All it requires is that the intra-region variation is smaller than the inter-region variation. Further, it is also not constrained just to the use of intensity value. Other region features like texture can also be used within this framework as our example with the ultrasound image demonstrates. The method is posed within a Bayesian framework of maximization of the *a posteriori* probability. Appropriate use of Green’s theorem makes the algorithm computationally attractive. As the examples show, the integrated approach is more robust to both increased amounts of noise, as well as increasingly displaced initialization of the initial boundary. Almost uniformly there is an improvement over the con-

ventional gradient-based boundary finding. To prove this, we have devised a variety of experiments and the results from all of them are favorable.

Application of this method on real medical images results in noticeable improvement as has been shown. We are using it for clinical research purposes for outlining the endocardial and epicardial boundaries of the heart and the results are much better than what we had achieved using the purely boundary-based method of [33].

However, there remain areas of potential improvement. We note that some of the probabilistic assumptions may not hold if the parameterized boundary is too far away from the true boundary at any instant. However, the resulting objective function that is generated from the probabilistic assumptions can also be interpreted as finding the boundary that best matches (agrees) with the gradient and region image. This interpretation continues to be valid. Thus in some sense, the method transcends its probabilistic limitations albeit in an ad hoc way. This can be made mathematically more complete. Also, the method can be extended to find multiple objects simultaneously. Generating priors, maybe from a multiscale image representation, could also be an area of potential investigation. Currently, we are in the process of extending the whole integration framework to three-dimensional images [9].

REFERENCES

- [1] L. Alvarez, P.L. Lions, and J.M. Morel. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journal of Numerical Analysis*, 29:845–866, 1992.
- [2] A.A. Amini, T.E. Weymouth, and R.C. Jain. Using dynamic programming for solving variational problems in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:855–867, September 1990.
- [3] P. Baxandall and Hans Liebeck. *Vector Calculus*. Oxford University Press, 1986.
- [4] J. Besag. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society*, 48:259–302, 1986.
- [5] A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press, Cambridge, MA, 1987.
- [6] P.J. Burt, T.H. Hong, and A. Rosenfeld. Segmentation and estimation of region properties through co-operative hierarchical computation. *IEEE Transactions on System, Man and Cybernetics*, 11:802–809, 1981.
- [7] J.F. Canny. A computational approach to edge-detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:679–698, 1986.
- [8] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal of Computer Vision*, to appear.
- [9] A. Chakraborty, L.H. Staib, and J.S. Duncan. An integrated approach to surface finding for medical images. *Proc. IEEE Workshop on Mathematical Models in Biomedical Image Analysis*, pages 253–262, 1996.
- [10] A. Chakraborty, M. Worring, and J.S. Duncan. On multi-feature integration for deformable boundary finding. *Proceedings of the International Conference on Computer Vision*, pages 846–851, 1995.
- [11] R. Chellappa. *Two dimensional discrete Gaussian Markov random field models for image processing*. Progress in Pattern Recognition 2, L.N. Kanal and A. Rosenfeld Ed., 1985.
- [12] C.C. Chu and J.K. Agarwal. The integration of image segmentation maps using region and edge information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15:1241–1252, 1993.
- [13] T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham. Active shape models - Their training and application. *Computer Vision, Graphics and Image Processing: Image Understanding*, 61:38–59, 1995.
- [14] D. Geiger and A. Yuille. A common framework for image segmentation. *International Journal of Computer Vision*, 6:227–243, 1991.
- [15] D. Geman and S. Geman. Stochastic relaxation, Gibbs distribution and Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.
- [16] J.F. Haddon and J.F. Boyce. Image segmentation by unifying region and boundary information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:929–948, 1990.
- [17] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:312–331, 1988.
- [18] J. Kittler and J. Illingworth. On threshold selection using clustering criteria. *IEEE Transactions on System, Man and Cybernetics*, 15:652–655, 1985.
- [19] F.P. Kuhl and C.R. Giardina. Elliptic Fourier features of a closed contour. *Computer Graphics and Image Processing*, 18:236–258, 1982.
- [20] R. Leahy, T. Herbet, and R. Lee. Application of Markov random fields in medical imaging. *Information Processing in Medical Imaging*, pages 1–14, 1989.
- [21] Y.G. Leclerc. Constructing simple stable descriptions for image partitioning. *International Journal of Computer Vision*, 3:73–102, 1989.
- [22] C. S. Lin and C. L. Hwang. New forms of shape invariants from elliptic Fourier descriptors. *Pattern Recognition*, 20(5):535–545, 1987.
- [23] R. Malladi, J.A. Sethian, and B.C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17:158–175, 1995.
- [24] B.S. Manjunath and R. Chellappa. Unsupervised texture segmentation using markov random field models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13:478–482, 1991.
- [25] D. Marr and E. Hildreth. Theory of edge detection. *Proc. Royal Society, London*, B207:187–217, 1980.
- [26] D. Mumford. Bayesian rationale for the variational formulation. In B. H. Romeny, editor, *Geometry Driven Diffusion in Computer Vision*, pages 135–146. Kluwer, 1994.
- [27] D. Mumford and J. Shah. Boundary detection by minimizing functionals. *Proceedings, IEEE Conf. Computer Vision and Pattern Recognition*, page 22, 1985.
- [28] S. Osher and J.A. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79:12–49, 1988.
- [29] T. Pavlidis and Y. Liow. Integrating region growing and edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:225–233, 1990.
- [30] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:629–639, 1990.
- [31] R. Ronfard. Region-based strategies for active contour models. *International Journal of Computer Vision*, 13:229–251, 1994.
- [32] B.S. Manjunath T. Simchony and R. Chellappa. Stochastic and deterministic networks for texture segmentation. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 38:1039–1049, 1990.
- [33] L.H. Staib and J.S. Duncan. Boundary finding with parametrically deformable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:161–175, 1992.
- [34] H. Tek and B.B. Kimia. Image segmentation by reaction diffusion bubbles. *Proceedings of the International Conference on Computer Vision*, pages 156–162, 1995.
- [35] M. Worring, A.W.M. Smeulders, L.H. Staib, and J.S. Duncan. Parametrized feasible boundaries in gradient vector fields. *Information Processing in Medical Imaging*, pages 48–61, 1993.
- [36] S.C. Zhu, T.S. Lee, and A.L. Yuille. Region competition: Unifying snakes, region growing and Bayes/MDL for multi-band image segmentation. *Proceedings of the International Conference on Computer Vision*, pages 416–423, 1995.