



A Coupled Minimization Problem for Medical Image Segmentation with Priors

YUNMEI CHEN

Department of Mathematics, University of Florida, Gainesville, FL 32611

yun@math.ufl.edu

FENG HUANG

Invivo Corporation, Gainesville, FL 32608

fhuang@invivocorp.com

HEMANT D. TAGARE

*Department of Diagnostic Radiology, Department of Electrical Engineering, Yale University,
New Haven, CT 06520*

hemant.tagare@yale.edu

MURALI RAO

Department of Mathematics, University of Florida, Gainesville, FL 32611

rao@math.ufl.edu

Received August 11, 2004; Revised November 2, 2005; Accepted March 14, 2006

First online version published in June, 2006

Abstract. We present a coupled minimization problem for image segmentation using prior shape and intensity profile. One part of the model minimizes a shape related energy and the energy of geometric active contour with a parameter that balances the influence from these two. The minimizer corresponding to a fixed parameter in this minimization gives a segmentation and an alignment between the segmentation and prior shape. The second part of this model optimizes the selection of the parameter by maximizing the mutual information of image geometry between the prior and the aligned novel image over all the alignments corresponding to different parameters in the first part. By this coupling the segmentation arrives at higher image gradient, forms a shape similar to the prior, and captures the prior intensity profile. We also propose using mutual information of image geometry to generate intensity model from a set of training images. Experimental results on cardiac ultrasound images are presented. These results indicate that the proposed model provides close agreement with expert traced borders, and the parameter determined in this model for one image can be used for images with similar properties.

Keywords: segmentation, registration, shape prior, intensity prior, mutual information of image geometry, variational method

1. Introduction

In numerous medical imaging modalities, the boundaries of anatomical structures cannot be detected by algorithms that only use edge or region information. Reasons for this are significant signal loss, noise, and non-uniformity of regional intensities. These problems are ever present for images acquired in cardiac ultrasound, where the boundary detection problem is further complicated by the presence of confusing anatomical structures such as the mitral valve and papillary muscles. In some cases, an image sequence may even have portions of the myocardium that lie outside of the sector scan so that some segments of its boundary may not be visible at all.

In an effort to overcome these difficulties, various techniques have been developed to incorporate prior information into the segmentation process. In Cootes et al. (1995) and Wang and Staib (1998), a statistical shape model was constructed from a set of corresponding points across the training images. This information was used in a Bayesian formulation to find the object boundary. In Cootes et al. (1999) a Gaussian model was fit to a training set of corresponding feature points. In Cootes and Taylor (1999) mixed models were used to fit to the data for specific applications where the distributions are non-Gaussian. In an alternate approach in Staib and Duncan (1992), Staib and Duncan specified the shape of the curve by creating statistical priors on the Fourier coefficients of the contour. This prior was incorporated into segmentation processing in a Bayesian framework. In Szekely et al. (1996) also developed Fourier parameterized shape models. In their model, an elastic fit of the shape model in the subspace of eigenmodes restricts the deformations. Another approach using shape templates was presented in Tagare (1997), where an energy function was created whose minimum corresponded to a boundary in the image that is similar in shape to the template. In Yuille et al. (1992) a method of deformable templates was proposed for feature extraction from faces. The features of the interest were described by a parameterized template which interacts dynamically with the image to minimize the energy function. More deformable models in medical image segmentation can be found in McInerney and Terzopoulous (1996).

Recently, statistical shape knowledge has been incorporated in edge based or region based active contours. In Leventon et al. (2000b), Leventon et al. extended geometric active contours developed in Caselles et al.

(1997) and Kichenassamy et al. (1995) by incorporating shape information into the evolution process. A principal component analysis was used to form a statistical shape model from a training set represented by using signed distance functions. The evolution of the interface was first driven by a force depending on image gradient and curvature, and then a correction was made by maximizing a posterior estimate of shape and pose. In Chen et al. (2002) incorporated the shape information into geodesic active contours by using variational methods. They created a shape term in the energy of the geodesic active contour, so that the deformation of the active contour stops when it arrives at high image gradients and forms a shape similar to the prior. The shape term also recovers a similarity transformation that aligns the interface to the prior shape better. Shape priors have also been used in region based segmentation schemes. In Cremers et al. (2002) incorporated statistical shape knowledge into the Mumford-Shah segmentation scheme (Mumford and Shah, 1989) by minimizing a functional that includes the shape energy and the Mumford-Shah energy. In Tsai et al. (2001) build a parametric shape model, and the parameters adjusted to minimize a region based objective function which provides the segmentation.

Besides using prior shape in segmentation, in Leventon et al. (2000a) incorporated prior intensity and curvature profiles of the features extracted from a training set into the segmentation process by an approach similar to the one developed in Leventon et al. (2000b).

While experimental results have shown the effectiveness of using image and prior information in segmentation in numerous medical applications, many problems remain including the complexity and variability of the images, the accuracy of the measurements obtained, and the rapid computation times required by the user. One practical problem that remains is how to determine the parameter that balances the influences from image information and priors. If the evolution of an active contour is mainly governed by a force defined by image gradients, it may be sensitive to the initial step or may leak through the boundary where the edge feature is not salient. Conversely, if the force defined by the shape prior is the dominating term, the active contour may not arrive at the boundary of the object of interest even though it has a shape similar to the prior.

In this paper we propose to use prior intensity profile to determine the parameter that balances the influence from image information and shape prior. The basic idea is to find the optimal balance of the forces that govern

the evolution of an active contour, so that the active contour arrives at high gradients, forms a shape similar to the prior, and captures prior intensity profile. Our approach is solving a coupled optimization problem. The first problem minimizes a weighted sum of two energies depending on image gradient and prior shape, respectively. The solution of this problem with fixed weight provides a segmentation and a transformation that aligns the interface to the prior shape. The second problem maximizes the mutual information of image geometry (see definition in Section 2) between the prior image and the aligned novel image. The alignment is the solution of the first problem. By maximizing the energy of the second problem over all the solutions of the first problem corresponding to different weights we can get the ‘best’ estimate for the parameter used in the first problem, and hence obtain a desirable segmentation. A similar idea was proposed by Ma in his unpublished thesis (Ma, 1997).

The proposed model indeed performs segmentation and registration simultaneously. The registration in our model combines a rigid transformation and a local deformation. The rigid transformation is determined explicitly by shape matching, while the local deformation is determined implicitly by image gradient and prior intensity profile. This is somewhat similar to the idea used in Soatto and Yezzi (2002), Paragios et al. (2002), Rousson and Paragios (2002) for matching nonequivalent shapes. In those works they consider a general deformation as the composition of a finite dimensional group action (rigid or affine transformation) and a local deformation.

Another reason to use prior intensity profiles to assist segmentation is that in some cases information about the expected shape alone may not be sufficient to guide the active contour arriving at the boundary of the object of interest. For instance, in some 2-chamber cardiac ultrasound images, the shape of the epicardium is nonequivalent to the prior while simultaneously having low contrast. In these images the models incorporating prior shape in active contours may not be able to give an accurate segmentation, as shown in Fig. 3 below. Using prior intensity profile across the boundary of the object can help the active contour to arrive at the right location.

Moreover, in this paper, we present an alternative method for generating *average* intensity profiles from a set of training images. To generate an intensity model in Leventon et al. (2000a) a Gaussian model was used to compute the joint distribution of the intensity values

and signed distances to boundary from a set of training images and segmentations. In ultrasound images the signal is partially in the form of speckle (Tao et al., 2002), and the statistics of speckle are non-Gaussian. Therefore, modeling the randomness of ultrasound images using a Gaussian model is not appropriate. Our method of generating an intensity model is model free and based on maximizing mutual information of image geometry between the intensity model and aligned training images.

We report experimental results on ultrasound images. We show that with appropriate prior shapes and intensity profiles, our technique is capable of finding boundaries in images, that are complicated by significant signal loss, poor signal to noise ratio, and non-uniformity of intensities. Moreover, our experiments indicate that the proposed model is not too sensitive to the initial step, and the optimized parameter for one image can be used for a group of images with similar properties. The existence problem for our model is also discussed.

2. Model Description

The literature on active contours is vast and we do not review all of it here. Our aim is to briefly discuss the formulation of active contours and then review work that directly deals with the incorporation of shape in active contours.

2.1. Shape Model

The notion of shape in our model is assumed independent of translation, rotation, and scaling. The shape model C^* used in our algorithm was obtained by averaging the aligned training contours with similar shapes. The alignment was made as follows. Fix C_1 , align C_j ($j = 2, \dots, n$) to C_1 by finding a scale μ_j , a rotation matrix R_j and a translation vector T_j such that the overlap area of C_1 and $\mu_j R_j C_j + T_j$ is maximized. i.e.

$$a(C_1, C_j) = \text{area of } (A_1 \cup A_j - A_1 \cap A_j) \quad (2.1)$$

is minimized, where A_1 and A_j are the inside areas of C_1 and $\mu_j R_j C_j + T_j$ respectively. Then the average shape C^* is given by $C^* = (C_1 + \sum_2^n \mu_j R_j C_j + T_j)/n$.

In case the sample curves have large shape variations, we use a clustering process to create the priors. The procedure is as follows. Fix a curve from

the training set (say C_1), and align the rest of the curves to C_1 as above to get $\mu_j R_j C_j + T_j$, and correspondingly, $a(C_1, C_j)$ ($j = 1, \dots, n-1$). Define $\sigma = \sum_{i=2}^n a(C_1, C_i)/(n-1)$. Then the first cluster consists of all the curves C_k , such that $a(C_1, C_k) < \sigma/3$. The remaining curves then are grouped into several clusters by repeating the process. The details are available in Chen et al. (2001). The drawback of this method is not symmetrical in the sense that the average shape may depend on the choice of C_1 .

An alternative method we used to create the prior shapes was the self-organizing maps algorithm. This method requires to determine in advance the number of the clusters containing all the training shapes. If we want to group the n training contours C_i $i = 1, \dots, n$ into k clusters (say $k = 3$), we first take three arbitrary contours as the initial contours, denoted by $m_j(0)$ ($j = 1, 2, 3$). At $t + 1$ iteration, randomly select a contour denoted by $X(t + 1)$ from the training set, and compare the disparity in shape between $X(t + 1)$ and each of $m_j(t)$ ($j = 1, 2, 3$). To do this comparison we first align $X(t + 1)$ to each $m_j(t)$, and denote the aligned $X(t + 1)$ by $\tilde{X}_j(t + 1) = \mu_j R_j X(t + 1) + T_j$. Then we compute $A_j = a(\tilde{X}_j(t + 1), m_j)$ defined in (2.1). Suppose A_1 is the smallest number in A_j $j = 1, 2, 3$. We keep $m_2(t)$ and $m_3(t)$ unchanged, and update $m_1(t)$ by

$$m_1(t + 1) = m_1(t) + \alpha(t)[\tilde{X}_1(t + 1) - m_1(t)],$$

where $\alpha(t)$ is a smooth function of t , and decreases to zero as $t \rightarrow \infty$. After a large number, say N , of iterations, three *average* shapes $m_j(N)$ ($j = 1, 2, 3$) are generated. Then three clusters are formed by the curves that are closest to the *average* shapes. The closeness is again measured by the measurement in (2.1).

2.2. Intensity Model

In this section we introduce our method for generating an intensity model (*average* intensity profile across the *average* shape) from a set of training images.

Let $C_i(p)$ ($i = 1, \dots, m$) be a set of training segmentations in a cluster, and I_i ($i = 1, \dots, m$) be the set of images associated with $C_i(p)$. Let also $C^*(p)$ $p \in [0, 1]$ be the *average* shape in this cluster. Our task is to generate an *average* intensity profile I^* across C^* from the training images I_i .

To do this, we first align each contour C_i ($i = 1, \dots, m$) to C^* by a similarity transformation

(μ_i, R_i, T_i) that minimizes $a(C^*, C_i)$ defined in (2.1). As a consequence, $I_i(\mu_i^{-1} R_i^{-1}(x - T_i))$ is aligned with $I^*(x)$. Then, we examine the disparity in the training intensity profiles I_i across the training segmentation C_i as follows. Let

$$V_{\epsilon_0} = \{x \in \Omega \mid d(C^*, x) < \epsilon_0\}, \quad (2.2)$$

be a ϵ_0 neighborhood of C^* . If two images I_i and I_j are related approximately by

$$I_i(\mu_i^{-1} R_i^{-1}(x - T_i)) \simeq a I_j(\mu_j^{-1} R_j^{-1}(x - T_j)) + b \quad (2.3)$$

for some constants a and b , we may define the disparity measurement in the intensity profiles near the segmentations C_i and C_j by

$$\min_{a,b} \int_{V_{\epsilon_0}} |I_i(\mu_i^{-1} R_i^{-1}(x - T_i)) - a I_j(\mu_j^{-1} R_j^{-1}(x - T_j)) - b|^2 dx.$$

The integral is over V_{ϵ_0} rather than the entire image domain, since the intensity profiles near the segmentations are more meaningful. If the relation (2.3) is not valid, maximizing mutual information has been proven to be effective in solving matching problems, in particular in matching multi-modality images (see e.g. Collignon et al., 1995; Hermosillo et al., 2002; Thevenaz and Unser, 2000; Viola and Wells, 1997; Wells et al., 1996 and the references there). One of the advantages of using mutual information is that it does not require an explicit function that relates two images, but only assumes that aligned images explain each other better than when they are not aligned.

Mutual information between two random vectors X and Y is defined as

$$MI(X, Y) = H(X) + H(Y) - H(X, Y), \quad (2.4)$$

where

$$H(Z) = - \int_{R^N} p_Z(z) \log p_Z(z) dz, \quad (2.5)$$

is the Shannon entropy of a random N -vector Z with density $p_Z(z)$, and

$$H(X, Y) = - \int_{R^N} \int_{R^N} p_{X,Y}(x, y) \log p_{X,Y}(x, y) dx dy, \quad (2.6)$$

is the joint entropy of X and Y , and $p_{X,Y}(x, y)$ is the joint density function. Consider intensity at each pixel as the realization of a random variable, and all the random variables in an image have the same distribution. By using (2.4)–(2.6) the *common* mutual information (MI) of two images $I_i(x)$ and $I_j(x)$ on V_{e_0} can be computed as follows: Let

$$\begin{aligned} f(x) &= I_i(\mu_i^{-1}R_i^{-1}(x - T_i)), \\ g(x) &= I_j(\mu_j^{-1}R_j^{-1}(x - T_j)). \end{aligned} \quad (2.7)$$

$$MI_{V_{e_0}}(I_i, I_j) =: \int_{R^2} p_{f,g}(i_1, i_2) \log \frac{p_{f,g}(i_1, i_2)}{p_f(i_1)p_g(i_2)} di_1 di_2. \quad (2.8)$$

where $p_{f,g}(i_1, i_2)$, $p_f(i_1)$ and $p_g(i_2)$ are computed over V_{e_0} . From this definition we can see that the mutual information of two random variables f and g is the same as the Kullback-Leibler distance of the random variables with the probability density functions $p_{f,g}(i_1, i_2)$ and $p_f(i_1)p_g(i_2)$. Therefore, this is a measure of mutual dependence of f and g .

Note that if the locations of two points in the image are switched, the intensity profile is changed, but the density function remains the same. In order to match intensity profiles of two images, we propose to maximize the mutual information of image geometry (MIIG) rather than mutual information (MI).

By MIIG of I_i and I_j on V_{e_0} we mean that

$$\begin{aligned} \text{MIIG}_{V_{e_0}}(I_i, I_j) \\ = \int_{R^{10}} p_{F,G}(m, n) \log \frac{p_{F,G}(m, n)}{p_F(m)p_G(n)} dmdn \end{aligned} \quad (2.9)$$

where $p_F(m)$ and $p_G(n)$ are the the probability density function of

$$\begin{aligned} F(x) = \langle f(x - 2h), f(x - h), f(x), f(x + h), \\ f(x + 2h) \rangle, \end{aligned}$$

and

$$\begin{aligned} G(x) = \langle g(x - 2h), g(x - h), g(x), g(x + h), \\ g(x + 2h) \rangle, \end{aligned}$$

over V_{e_0} with f and g are defined in (2.8), respectively, $p_{F,G}(m, n)$ is the joint density function of $F(x)$ and $G(x)$ over V_{e_0} , $m = \langle i_1, i_2, i_3, i_4, i_5 \rangle$,

$n = \langle i_6, i_7, i_8, i_9, i_{10} \rangle$, $dm = di_1 di_2 di_3 di_4 di_5$, and $dn = di_6 di_7 di_8 di_9 di_{10}$. h is a vector, that can be chosen according to user's interest. The principle goal of h is to better use neighborhood intensity information. For instance, in order to segment cardiac borders in ultrasound images we set $x + h$ to be the first right neighboring point of x in the horizontal direction, since the intensity profiles across the cardiac boundaries vary more in this direction than other directions. Since MIIG uses neighborhood information, MIIG of two images gives better description of the closeness of the intensity profiles of these two images than MI.

The method was tested against a group of 85 apical 2-chamber end diastolic (ED) echocardiographic images acquired retrospectively from 61 normal patients. The expert traced epicardial borders were given in these images as the training segmentations. The images were grouped into three clusters by applying the grouping method described in Section 2.1 to these training segmentations. The 20 images from one of the clusters with expert borders superimposed are displayed in Fig. 1.

The *average* shape C^* for this cluster was created by the method mentioned in Section 2.1. We used both methods explained there, and got similar results. Then, the training images I_i in this cluster were aligned by aligning the corresponding training segmentation C_i to C^* using a similarity transformation (μ_i, R_i, T_i) that minimizes $a(C^*, C_i)$ defined in (2.1). Next, we define the neighborhood V_{e_0} as a three pixel neighborhood of C^* to compute the $\text{MIIG}_{V_{e_0}}(I_i, I_j)$. If the images in this cluster have relatively larger $\text{MIIG}_{V_{e_0}}(I_i, I_j)$ after some outliers are dropped, then we don't do any subgrouping. Otherwise, we may group images in this group into several subclusters by using $\text{MIIG}_{V_{e_0}}(I_i, I_j)$ as the distance function in the k-means clustering algorithm, so that the images in each subcluster has similar intensity profiles. In our experiment we grouped this 20 images into two subclusters. One of them has 11 images, which are displayed as the first 11 images in Fig. 1.

Denoting the subgroup of 11 images by I_1, \dots, I_{11} , we now describe the construction of the *average* intensity profile I^* across C^* for this subgroup of 11 images. Ideally, we would like to generate I^* by

$$\max_{I^*} \sum_{j=1}^{11} \text{MIIG}_{V_{e_0}}(I^*, I_j), \quad (2.10)$$

where $\text{MIIG}_{V_{e_0}}(I^*, I_j)$ is defined as in (2.9) with $f = I^*(x)$ and $g = I_j(\mu_j^{-1}R_j^{-1}(x - T_j))$. Since this formulation is computationally intensive, we restricted




Figure 1. 20 training images with the segmented epicardia used as training images and shapes.

ourselves to the cases where

$$I^* = \sum_{j=1}^{11} a_j I_j(\mu_j^{-1} R_j^{-1}(x - T_j)). \quad (2.11)$$

The weights a_j were determined by using (2.10). We applied this method to this subgroup of 11 training images, and obtained the *average* intensity profile I^* shown in Fig. 2(b). Moreover, we computed the sum in (2.10) for the a_j determined in (2.10) and $a_j = 1/11$ ($j = 1, \dots, 11$). The results were 196.73 and 184.08 respectively, confirming our suspicion that the weighted average is better. The *average* intensity profile of the other subgroup can be generated by the same method.

2.3. Segmentation with Priors

Now we are going to present our variational approach for segmentation using prior shape and intensity profile. The key point of our model is to propagate a curve/surface by a velocity that depends on




Figure 2. The image $I^* = \sum_{j=1}^{11} a_j I_j(\mu_j^{-1} R_j^{-1}(x - T_j))$ (a). $a_j = 1/11$; (b). a_j is determined by (2.10). In (a) and (b) the superimposed contours are the *average* shapes.

the image gradients, prior shape and intensity profile, such that the propagation stops when the active contour/surface forms a shape similar to the shape prior, arrives at high gradients, and captures the prior intensity profile.

To begin the description of the proposed model, we first briefly review the active contour with a shape prior in Chen et al. (2002). Let $C^*(p)$ ($p \in [0, 1]$) be a curve, representing the shape prior. The model in Chen et al. (2002) minimizes the energy functional $E(C, \mu, R, T)$ defined as

$$\int_0^1 \{g(|\nabla I|)(C(p)) + \frac{\lambda}{2}d^2(\mu RC(p) + T)\}|C'(p)|dp, \tag{2.12}$$

where (μ, R, T) are similarity transformation parameters, $d((x, y))$ is the distance of the point (x, y) from C^* , and $g(|\nabla I|) = \frac{1}{1+\beta|\nabla G_{\sigma^*}I|^2}$, with a parameter $\beta > 0$, and $G_{\sigma}(x) = \frac{1}{\sigma}e^{-\frac{|x|^2}{4\sigma^2}}$. The first term in (2.12) is the energy functional of geodesic active contours (Caselles et al., 1997; Kichenassamy et al., 1995), that measures the amount of high gradient under the trace of the curve. The second term is the shape related energy, that measures the disparity in shape between the interface and the prior. The constant $\lambda > 0$ is a parameter, that balances the influence from the image gradient and prior shape. The curve C and the transformation parameters μ, R and T evolve to minimize $E(C, \mu, R, T)$. At the stationary point, the contour C lies over points of high gradient in the image and forms a shape close to C^* . This model also provides accurate estimates for a similarity transformation that aligns the interface to prior shape. A geometric active contour often “leaks” through “gaps” in the boundary, which have low gradients, since it does not have any information about how the gaps are to be bridged. However, model (2.12) incorporates the information about the expected overall shape of the boundary in geometric active contours so that the active contour can compare its shape with the expected shape and bridge the gaps in a meaningful way. The experimental results in Chen et al. (2002) showed model (2.12) is able to get a satisfactory segmentation in the presence of gaps, if the boundary of the object of interest has a shape similar to the prior. However, there are still some problems in using this model. First, when the gaps are a substantial fraction of the overall boundary, and at the gaps the shape of boundary has relatively large geometric distortion from the prior, model (2.12) can not provide an accurate segmentation, since the knowledge about expected shape alone is not sufficient to provide the correct information regarding how the gaps should be bridged. Another problem is how to determine the parameter λ in (2.12). Inappropriate value of λ will result in a bad segmentation.

To have a better solution to these problems, we propose to incorporate both prior shape and intensity profile in segmentation. One solution for doing this could be the creation of a intensity profile related term in the energy functional (2.12). As discussed in the previous section, we propose to use the MIIG of the prior image I^* and aligned novel image I in a neighborhood of C^* as the measurement of the disparity in intensity profiles across the interface and prior shape. Then, we may incorporate this intensity related energy into (2.12) to form the following optimization problem:

$$\begin{aligned} \min_{C, \mu, R, T} & \int_0^1 \{g(|\nabla I|)(C(p)) \\ & + \frac{\lambda_1}{2}d^2(\mu RC(p) + T)\}|C'(p)| dp \\ & - \lambda_2 \text{MIIG}_{V_{\epsilon_0}}(I^*(x), I(\mu^{-1}R^{-1}(x - T))), \end{aligned} \tag{2.13}$$

where V_{ϵ_0} is defined in (2.2), and $\lambda_i > 0$ ($i = 1, 2$) are parameters balancing the influences from the image gradient, prior shape and prior intensity profile. By adding the third term in (2.13), the active contour governed by (2.12) is forced to arrive at high gradient, form a shape similar to the prior, and capture the prior intensity profile near the feature. However, the problem of determining the parameters $\lambda_i > 0$ ($i = 1, 2$) is not trivial.

In this paper we present an alternative approach, that is not only able to incorporate both shape and intensity information in segmentation, but also able to provide better estimate for the parameter used in the model. Our model is a coupled optimization problem, which consists of a minimization and a maximization problems below:

$$C_{\lambda}, \mu_{\lambda}, R_{\lambda}, T_{\lambda} = \arg \min_{C, \mu, R, T} E_{\lambda}(C, \mu, R, T), \tag{2.14}$$

and

$$\lambda^* = \arg \max_{\lambda} F(\lambda), \tag{2.15}$$

where

$$\begin{aligned} E_{\lambda}(C, \mu, R, T) = & \int_0^1 \{g(|\nabla I|)(C(p)) \\ & + \frac{\lambda}{2}d^2(\mu RC(p) + T)\}|C'(p)|dp, \end{aligned}$$

and

$$F(\lambda) = F(\mu_\lambda, R_\lambda, T_\lambda) = \text{MIIG}_{V_{\epsilon_0}}(I^*(x), I(\mu_\lambda^{-1} R_\lambda^{-1}(x - T_\lambda))).$$

The basic idea of this model is to optimize the selection of the parameter, that balances the forces from deformation from shape prior and edge information, by maximizing the MIIG between the novel image and intensity prior in the neighborhoods of the shape prior. For a given λ the solution of (2.14) provides an optimal segmentation C_λ and a transformation $(\mu_\lambda, R_\lambda, T_\lambda)$. By maximizing the energy functional in (2.15) over all the possible solutions $(\mu_\lambda, R_\lambda, T_\lambda)$ of (2.14) corresponding to λ , we can get an optimal estimate λ^* for λ , and hence a better segmentation C_{λ^*} corresponding to λ^* .

Therefore, although the intensity prior was not explicitly included in the energy functional (2.14) to get a segmentation, the segmentation result in fact was greatly influenced by the intensity prior from its role in the selection of the parameter in (2.14). By the coupling of (2.14) and (2.15) the segmentation arrives at higher image gradient, forms a shape similar to the prior, and captures the prior intensity profile.

2.4. Level Set form and EL Equations

Level set methods (Osher and Sethian, 1988) have been used extensively in active contour models because they allow for cusp, corners, and automatic topological changes. In this part we will present the level set formulation of the proposed model (2.14)–(2.15).

Represent a contour C by the zero level set of a Lipschitz function u such that $\{x \mid u(x) > 0\}$ is the set inside C . Let $H(z)$ be the Heaviside function, that is $H(z) = 1$ if $z \geq 0$, and $H(z) = 0$ if $z < 0$, and $\delta = H'(z)$ (in the sense of distribution) be the Dirac measure. Then, the length of the zero level set of u in the conformal metric $ds = g|C'(p)|dp$ can be computed by $\int_\Omega g|\nabla H(u)| = \int_\Omega \delta(u)g|Du|$. The similarity of the shapes between the zero level set of u and C^* can be evaluated by $\int_\Omega \delta(u)d^2(\mu Rx + T)dx$. Then, the level set formulation of (2.14) is given by

$$\min_{u, \mu, R, T} \int_\Omega \delta(u) \left\{ g(|\nabla I|)(x) + \frac{\lambda}{2} d^2(\mu Rx + T) \right\} |\nabla u|. \tag{2.16}$$

The evolution equations associated with the Euler-Lagrange equations of (2.16) are

$$\frac{\partial u}{\partial t} = \delta(u) \text{div} \left\{ \left(g + \frac{\lambda}{2} d^2 \right) \frac{\nabla u}{|\nabla u|} \right\}$$

$$\frac{\partial u}{\partial n} = 0, \text{ on } \partial\Omega, t > 0, \tag{2.17}$$

$$\frac{\partial \mu}{\partial t} = -\lambda \int_\Omega \delta(u) d \nabla d \cdot (Rx) |\nabla u| dx \tag{2.18}$$

$$\frac{\partial \theta}{\partial t} = -\lambda \int_\Omega \delta(u) \mu d \nabla d \cdot \left(\frac{dR}{d\theta} x \right) |\nabla u| dx, \tag{2.19}$$

$$\frac{\partial T}{\partial t} = -\lambda \int_\Omega \delta(u) d \nabla d |\nabla u| dx, \tag{2.20}$$

where R is the rotation matrix in terms of the angle θ , d is evaluated at $\mu Rx + T$.

3. Numerical Method and Experimental Results

We used the level set form (2.16) in our experiments. (2.16) was solved by finding a steady state solution of (2.17)–(2.20). (2.17) was implemented by the following iteration scheme:

$$\frac{u_{i,j}^{n+1} - u_{i,j}^n}{\Delta t} = \frac{1}{h^2} \delta_\epsilon(u_{i,j}^n) \times \left\{ \Delta_x^- \left(\frac{A_{ij}^n \Delta_x^+ u_{i,j}^{n+1}}{\sqrt{(\Delta_x^+ u_{i,j}^n)^2/h^2 + (u_{i,j+1}^n - u_{i,j-1}^n)^2/(2h)^2}} \right) + \Delta_y^- \left(\frac{A_{ij}^n \Delta_y^+ u_{i,j}^{n+1}}{\sqrt{(\Delta_y^+ u_{i,j}^n)^2/h^2 + (u_{i+1,j}^n - u_{i-1,j}^n)^2/(2h)^2}} \right) \right\}$$

where $u_{i,j}^n = u(n\Delta t, ih, jh)$, and

$$A_{ij}^n = g_{i,j}^n + \frac{\lambda}{2} d_{i,j}^{n,2}.$$

The function δ_ϵ is a smooth version of the Dirac δ function, we used

$$\delta_\epsilon(z) = \begin{cases} 0 & \text{if } |z| > \epsilon \\ \frac{1}{2\epsilon} [1 + \cos(\frac{\pi z}{\epsilon})] & \text{if } |z| \leq \epsilon \end{cases}$$

To prevent the level set function to become too flat near the front we applied a reinitialization technique to reinitialize u to the signed distance function to its zero level curve, as in Sussman et al. (1994) and Zhao et al. (1996). This procedure is made by using a new

function $v(x)$, which is the steady state solution to the equation

$$\frac{\partial v}{\partial s} = \text{sign}(u(\cdot, t))(1 - |\nabla v|), v(\cdot, 0) = u(\cdot, t),$$

as $u(\cdot, t)$ for the next iteration $t + \Delta t$.

However, how to solve (2.15) efficiently still remains an open question. In our experiments we first solved (2.16) to get a sequence of solutions $(C_{\lambda_i}, \mu_{\lambda_i}, R_{\lambda_i}, T_{\lambda_i})$ corresponding to a sequence of $\lambda_i, i = 1, \dots, k$. Then we computed $F(\mu_{\lambda_i}, R_{\lambda_i}, T_{\lambda_i})$ in (2.15) for each $1 \leq i \leq k$. In this computation we partitioned the intensities into 16 bins, and used the discrete form of Shannon entropy to calculate the MIIG. If $F(\mu_{\lambda_j}, R_{\lambda_j}, T_{\lambda_j})$ is the largest one among them, then we take $C_{\lambda_j}, \mu_{\lambda_j}, R_{\lambda_j}, T_{\lambda_j}$ as our model solutions. In this way we could get a better estimate for λ , and hence, a better segmentation, that captures high gradients, shape prior, and intensity profile. However, it may not be the absolute optimal, since the comparison of the energy values in (2.15) was only for finitely many λ 's.

We applied this algorithm on 2-chamber cardiac ultrasound images. The epicardia and the endocardia in these images were not completely imaged, and our task was to find and complete these boundaries using prior shape and intensity profile.

To create the prior shape, epicardial boundaries were outlined by an expert echocardiographer on 85 images acquired at ED from 61 patients. Using the method described in Section 2.1, the boundaries were grouped into three clusters and the *average* shape of each cluster was computed. Using the method described in Section 2.2, the *average* intensity profiles for each *average* cluster (or subcluster) can be generated. The *average* shape C^* and the associated *average* intensity profile I^* for one of the subclusters is displayed in Fig. 2(b). (The images in this subcluster are the first 11 images displayed in Fig. 1.)

To segment the epicardial border in a novel image displayed in Fig. 3(b), we have to select an *average* contour and the associated intensity profile of one cluster from all the clusters as the priors in our model. Making a good selection is very important and needs some input from experts or experienced users. This is a problem that needs to be further studied. In our experiment suggested by a cardiologist we used the *average* contour and intensity profile near the contour in Fig. 2(b) as the priors.

The active contour was initialized with the ellipse displayed in Fig. 3(a). Evolving the active contour




Figure 3. (a) An ellipse used as the initial contour in our experiments for three images in the following three rows. (b)–(g). Each row presents the segmentations (solid, red), expert’s borders (dotted, black), and average shape (green) in an image. The segmentations in the left column and right column are the solutions of (2.16) corresponding to $\lambda = 0.04$ and $\lambda = 20$ respectively.

using equations (2.17)–(2.20) with a fixed λ , we obtained a segmentation C_λ together with a similarity transformation $(\mu_\lambda, R_\lambda, T_\lambda)$, that are the solutions of (2.16). By varying λ we generated a sequence of solutions of (2.16). We chose the optimal value λ^* of λ to be the one maximizing (2.15). Finally, the solutions for the model (2.16) coupled with (2.15) were chosen as the solutions of (2.16) corresponding to λ^* .

The first column of the table displays 8 different values of λ . By the procedure described above we obtained $C_{\lambda_i}, \mu_{\lambda_i}, R_{\lambda_i}, T_{\lambda_i}, (i = 1, \dots, 8)$. The third and the fourth columns present the $MI_{V_{\epsilon_0}}(I^*(x), I(\mu_{\lambda_i}^{-1} R_{\lambda_i}^{-1}(x - T_{\lambda_i})))$, and $MIIG_{V_{\epsilon_0}}(I^*(x), I(\mu_{\lambda_i}^{-1} R_{\lambda_i}^{-1}(x - T_{\lambda_i})))$, respectively. Since the 4th column of this table is largest when $\lambda = 0.04$, we selected the solutions of (2.11) to correspond to this choice of $\lambda = 0.04$. The segmentation (solid) corresponding to this λ is shown in Fig. 3(b) together with the expert traced border (dotted). The distance between the expert and algorithm generated borders are tabulated in column 2 of the table. It is defined as $\sum_{i=1}^N d_{C^{ep}}(C(p_i))/N$, where $d_{C^{ep}}$ is the distance function of the experts traced border C^{ep} , and $C(p)$ ($0 \leq p \leq N, P_1 = 0, p_N = 1$) is our segmentation. The units of the distance are the numbers of the pixels, and the pixel size is $0.62 \text{ mm} \times 0.62 \text{ mm}$. From this table we see that the segmentation corresponding to $\lambda = 0.04$ is the one having smallest distance from expert's contour and largest value of MIIG. This statement is not true for MI.

λ	dist	MI	MIIG
0.04	3.1142	1.8894	7.9971
0.20	3.2903	1.9192	7.9794
0.40	3.3552	1.9207	7.9698
2	3.3962	1.9579	7.9623
20	4.7160	1.9207	7.8914
40	6.6697	2.0133	7.7524
0.02	25.4498	0	0
0.004	26.2190	0	0

The second row in Fig. 3 displays this experimental result. The segmentations (solid) in Fig. 3(b) and (c) are the solutions of (2.16) with $\lambda = 0.04$ and $\lambda = 20$, respectively. Comparing the segmentation results with the expert's borders (dotted) this figure provides visual confirmation of the result presented in the table.

To further test the method, we also used the same initial contour and the 'optimal' value $\lambda = 0.04$ determined in the previous test in two additional images. The segmentation results (solid) together with expert's border's (dotted) are presented in Figs. 3(d)–(g). The segmentations in the left and right columns are the solutions of (2.16) corresponding to $\lambda = 0.04$ and $\lambda = 20$, respectively. Comparing the results in the left column of Fig. 3 with those in the right column, we observe that $\lambda = 0.04$ also provides good segmentations in these two new images indicating that the 'optimal' estimate of λ from one image can possibly be used for other members of the cluster. Of course, the shapes of the object boundaries and their intensity profiles must be similar.

Moreover, the proposed model is less sensitive to the initial step than the edge-based active contours, since the deformation of the contour is influenced by intensity profiles in addition to image gradients.

Next we present an experiment result on synthetic data to show more clearly the relationship between maximization of MIIG and optimality of the parameter for getting a better segmentation. Figure 4(a) shows a synthetic image with dimension 128×128 and the expected segmentation (the red contour). Our task is to find a contour in this image, that is close to the expected one. Fig. 4(b) presents the prior shape (the red contour) and intensity profiles. The image in Fig. 1(b) was generated by rotating and then blurring the image in Fig. 1(a) The rotation was $\pi/6$ degree, and the blurring was made by convolution with a Gaussian kernel.

We applied model (2.16) with the prior shape and intensity profiles presented in Fig. 1(b) to the image in Fig. 1(a). For each fixed λ , by solving (2.16) we obtained a segmentation C_λ and a similarity transform $(\mu_\lambda, R_\lambda, T_\lambda)$. By varying λ we got a sequence of solutions corresponding to the different λ 's. The table below and the Fig. 4(c)–(e) show how the optimal λ and its corresponding segmentation are related to the MIIG of the image and prior in the neighborhood of the prior shape.

In Fig. 4(c)–(e) the white ellipse was used as the initial contour to solve the EL equation associated with (2.16), while the red solid contours were the solutions of (2.16) corresponding to $\lambda = 0.04, 1, \text{ and } 0.45$, respectively. When $\lambda = 0.04$, the gradient term weighted too much so that the contour in Fig. 4(c) stopped at high gradient. While $\lambda = 1$, the shape term weighted too heavy so that the contour in Fig. 4(d) stopped when it formed a shape similar to the prior shape, and the




Figure 4. (a) A synthetic image with the expected segmentation (red contour) (b) Shape (red contour) and intensity priors (c)–(e). Segmentations (solid, red) as the solution of (2.16) corresponding to $\lambda = 1, 0.04$ and 0.45 , respectively. The ellipse in each figure is the initial contour.

small circle in the expected shape (the red contour in Fig. 4(a)) was missed. The contour in Fig. 4(e) corresponding to $\lambda = 0.45$ was close to the expected segmentation. It arrived at high gradient, and captured the shape and intensity information.

This is further illustrated in the following table. The first column of the table displays these 3 different values of λ . The second column shows the distance between the expected contour and our algorithm generated contours. The third and the fourth columns present the MI and MIIG between the intensity prior in Fig. 4(b) and the novel image in Fig. 4(a) corresponding to different λ 's, respectively. The definition of the distance is the same as that in the previous experiment of finding the epicardial border. This table indicates that the segmentation corresponding to $\lambda = 0.45$ is the one having smallest distance from expected contour and largest value of MIIG. Again, this statement is not true for MI.

λ	dist	MI	MIIG
0.45	1.41	0.1591	0.3019
1	0.8487	0.1050	0.2828
0.04	2.84	0.1209	0.1814

We also applied our algorithm to segment the endocardia in ultrasound images. The shape model of endocardium and the prior intensity profiles across the shape model shown in Fig. 5(b) were generated from a set of 12 training shapes and their associated images shown

in Fig. 5(a) by using our proposed method explained above.

To segment the endocardial border in a novel image displayed in Fig. 5(b), we used the *average* contour and intensity profile near the contour in Fig. 5(b) as the priors. The active contour was initialized with the ellipse displayed in Fig. 5(a). Evolving the active contour using equations (2.17)–(2.20) with a fixed λ , we obtained a segmentation C_λ together with a similarity transformation $(\mu_\lambda, R_\lambda, T_\lambda)$. By varying λ we generated a sequence of solutions of (2.16). The optimal value of λ was chosen as the one maximizing (2.15). Finally, the solutions of (2.16) corresponding to this optimal λ is our model solution.

The table below shows a list of the values of λ 's used in (2.16) together with the distance between the expert's borders and the segmentations, the values of MI and MIIG on the neighborhoods of the segmentations, which are the solutions of (2.16) corresponding to these λ 's,

λ	dist	MI	MIIG
0.1	1.97	1.3711	6.8060
0.2	2.14	1.4390	6.7197
0.01	2.36	1.1653	6.6282
0.001	2.78	1.0783	6.4244
1.0	8.18	1.1719	6.4050

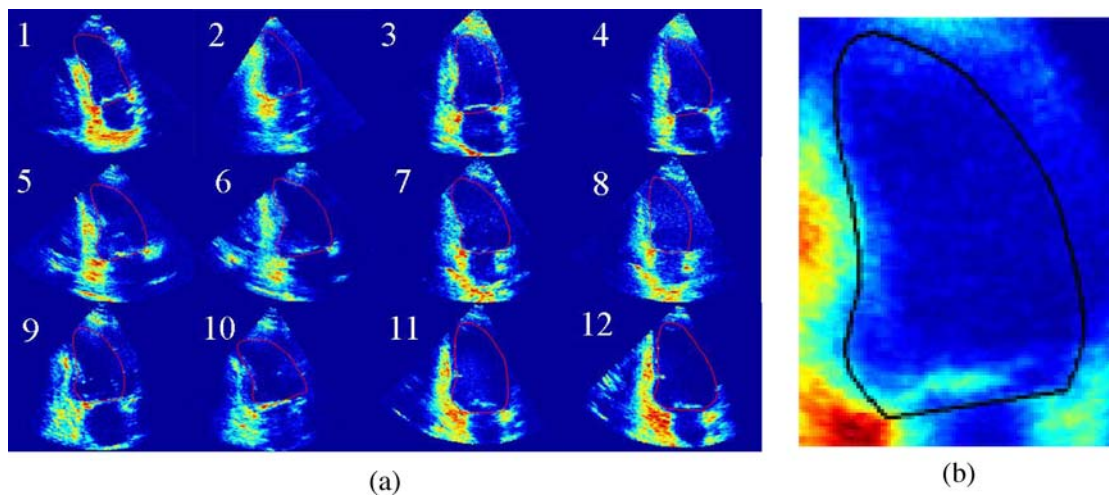


Figure 5. (a) 12 training images with the segmented endocardia used as training images and shapes, (b) Average endocardium with associated intensity profiles generated from the training shapes and images in Fig. 4(a) by the proposed algorithms.

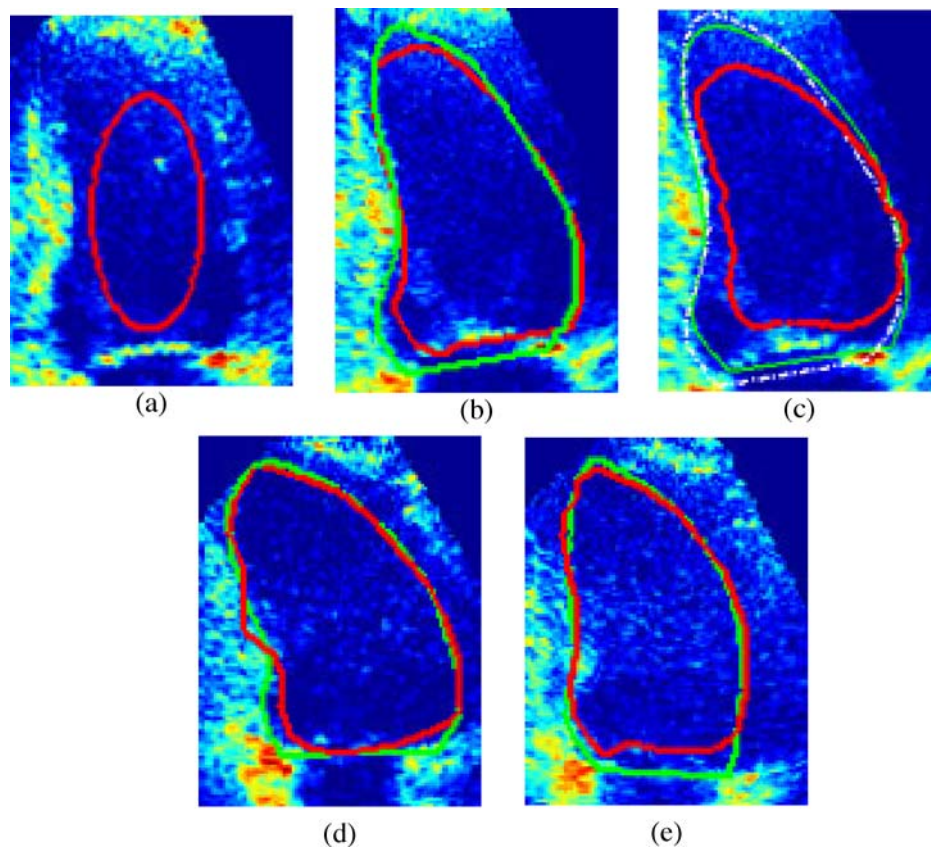


Figure 6. (a) An ellipse used as the initial contour in our experiments for three images. (b), (d), and (e). Segmentations as the solutions of (2.16) (solid, red) corresponding to $\lambda = 0.1$ comparing with the expert's borders (dotted, green) in three different images. (c). Segmentation as the solution of (2.16) (solid, red) corresponding to $\lambda = 1$ comparing with the expert's borders (dotted, green) in the same image as in (b).

We observed the same phenomena as that in the previous two tables that the larger the MIIG, the smaller the distance of the segmentation from the expert's contour. This statement is not true for MI. From this table we see that $\lambda = 0.1$ is the optimal value of λ , and the solution corresponding to this λ is the model solution of (2.16).

Figure 6 presents the results of the segmentations of endocardium. Figure 6(a) shows the initial contour in a novel image. Figure 6(b) shows the solution C to the model (2.16) with the optimal $\lambda = 0.1$, while the segmentation in Fig. 6(c) is the solution of (2.16) with $\lambda = 1$. With this λ the model over weights the shape energy so that the contour stopped when a shape similar to the prior was formed. However, the contour has not captured either the high gradient or the image intensity profiles. Fig. 6(d)–6(e) display the segmentation results in two more novel images. They are the solution of (2.16) with $\lambda = 0.1$. In these figures, the segmentations and the expert's contours are represented by solid (red) and dotted (green) contours, respectively. This experiment results indicates again that the proposed model provides close agreement with expert traced borders, and the parameter determined in this model for one image can be used for images with similar properties.

4. Conclusion

We proposed a new method for segmentation that incorporates the shape and intensity priors in active contours by solving a coupled optimization problem. The energy function in one part of the model is a weighted sum of the energy of the geometric active contour and a shape related energy. The minimizer of this energy function gives the segmentation and the transformation that aligns it to the prior shape. The second part of the model provides an optimal estimate of the weight used in the first energy function by maximizing the MIIG of the intensity prior and the aligned novel image near the feature over all the alignments that are the solutions of the first part corresponding to different weights.

The improvements of this model over existing active contour algorithms are in two aspects. First, we used maximizing MIIG rather than MI to match intensity profiles of two images. The reason for doing this is that the MIIG takes neighborhood intensity distribution into account, and hence, gives better description of intensity profile than MI. Secondly, the parameter in the model of active contour with shape can be selected by model itself, and the segmentation with the optimized parameter arrives at higher image gradients, forms a shape

similar to the prior, and captures the prior intensity profile. Moreover, the coupling idea can be used for any models where the solution is influenced by three forces.

We applied our model to the problem of cardiac boundary determination in ultrasound images, for which the methods using edge or region information only can not give a good result. Even the active contours with shape prior struggle with such data. The proposed model was tested against a database of epicardial borders traced by an expert on echocardiographic images acquired from the apical 2-chamber view. The preliminary results were encouraging. The existence of the solution to the proposed model was proved.

However, much work needs to be done. Our method for generating shape model and intensity model has two steps: first creating the shape model, and then building a intensity model in the neighborhood of the shape model. The clustering process also involves two steps. We first group all of the training shapes into several clusters, then each cluster may divided into sub-clusters according to the intensity similarity of the images in the cluster. Finally in each sub-cluster the training shapes and their associated images were very close. As suggested by the reviewer, we will try to generate intensity and shape models simultaneously by using a disparity measurement depending on both the shape and intensity near the shape. We will also further study the approach presented in alez Ballester et al. (2003, 2004), where a generalized image model is presented. They present images as sets of four-dimensional (4D) sites combining position and intensity information, as well as their associated uncertainty and joint variation. This model allows for the representation of both images and statistical models as well as other representations. such as landmarks or meshes.

Moreover, our current model is for 2-d images. The extension of this model to 3-d cases will have additional computational cost for computing MIIG. Therefore, better numerical methods for implementing this model need to be further studied. We also aim to test the proposed method against the entire database of normal images and the images acquired from patients with some observable abnormality.

Acknowledgments

We thank Prof. David Wilson, Department of Mathematics, University of Florida, and Dr. Edward A. Geiser, Division of Cardiology, Department of Medicine, University of Florida for providing the

ultrasound images and expert's contours. We also thank Jung-ha An, Department of Mathematics, University of Florida, for the experiment of generating the shape prior using self-organizing map.

Chen was partially supported by NIH P50-DC03888 and R01 NS42075. Rao was partially supported by R01 NS42075. Tagare was partially supported by R01-LM06911-05 from the National Library of Medicine.

References

- Cootes, T., Beeston, C., Edwards, G., and Taylor, C. 1999. Unified framework for atlas matching using active appearance models. In *Int'l Conf. Inf. Proc. in Med. Imaging*, Springer-Verlag, pp. 322–333.
- Caselles, V., Kimmel, R., and Sapiro, G. 1997. Geodesic active contours. *IJCV*, 22(1):61–79.
- Collignon, A., Maes, F., Vandermeulen, D., Suetens, P., and Marchal, G. 1995. Automated multimodality image registration using information theory. In *Info. Processing in medical Images*, Brest, France, pp. 263–274.
- Cootes, T. and Taylor, C. 1999. Mixture model for representing shape variation. *Image and Vision Computing*, 17(8):567–574.
- Cootes, T., Taylor, C., Cooper, D., and Graham, J. 1995. Active shape model—their training and application. *Computer Vision and Image Understanding*, 61:38–59.
- Chen, Y., Tagare, H., Thiruvankadam, S.R., Huang, F., Wilson, D., Geiser, A., Gopinath K., and Briggs, R. 2002. Using prior shapes in geometric active contours in a variational framework. *IJCV*, 50(3):315–328.
- Cremers, D., Tischhäuser, F., Weickert, J., and Schnörr, C. 2002. Diffusion-snakes: Introducing statistical shape knowledge into the Mumford-Shah functional. *International Journal of Computer Vision*, 50(3):295–315.
- Chen, Y., Wilson D., and Huang, F. 2001. A new procrustes methods for generating geometric models. In *Proceedings of World Multiconference on Systems, Cybernetics and Informatics*, Orlando, pp. 227–232.
- alez Ballester, M.A.G., Pennee X., and Ayache, N. 2003. Generalized image models and their application as statistical models of images. In *Proceedings of MICCAI'03, Montreal, Lecture Notes in Computer Science*, Canada, 2879:150–157.
- alez Ballester, M.A.G., Pennee, X., Linguraru, M.G., and Ayache, N. 2004. Generalized image models and their application as statistical models of images. *Medical Image Analysis*, 8(3):361–369.
- Hermosillo, G., Chefid'Hotel, C., and Faugeras, O. 2002. A variational methods for multimodal image matching. *IJCV*, 50(3):329–345.
- Kichenassamy, S., Kumar, A., Olver, P., Tannenbaum, A., and Yezzi, A.J. 1995. Gradient flows and geometric active contour models. In *Proc. ICCV'95*, pp. 810–815.
- Leventon, M.E., Faugeras, O., Grimson, E., and Wells, W. 2000a. Level set based segmentation with intensity and curvature priors. *Mathematical Methods in Biomedical Image Analysis*.
- Leventon, M.E., Grimson, E., and Faugeras, O. 2000b. Statistical shape influence in geodesic active contours. In *Proc. IEEE Conf. CVPR*, pp. 316–323.
- Ma, T. 1997. Active contour models: Consistency, stability, and parameter estimation. Ph.D. Thesis, *Dept. of Computer Science*, Yale University.
- Mumford, D. and Shah, J. 1989. Optimal approximation by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42:557–685.
- McInerney, T. and Terzopoulos, D. 1996. Deformable models in medical image analysis: A survey. *Medical Image Analysis*, 1(2):91–108.
- Osher, S. and Sethian, J.A. 1988. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulation. *J. Comp. Physics*, 79:12–49.
- Rousson, M. and Paragios, N. 2002. Shape priors for level set representations. In *Proceedings of the 7th European Conference on Computer Vision*, Copenhagen, Denmark, pp. II:78–93.
- Paragios, N., Rousson, M., and Ramesh, V. 2002. Marching distance functions: A shape-to-area variational approach for global-to-local registration. In *Proceedings of the 7th European Conference on Computer Vision*, Copenhagen, Denmark, pp. 775–789.
- Staib, L. and Duncan, J. 1992. Boundary finding with parametrically deformable contour methods. *IEEE Trans. Patt. Anal. and Mach. Intell.*, 14(11):1061–1075.
- Szekely, G., Kelemen, A., Brechbuhler, C., and Gerig, G. Segmentation of 2D and 3D objects from MRI volume data using constrained elastic deformations of flexible Fourier contour and surface models. *Medical Image Analysis*, 1(1):199–234.
- Sussman, M., Smereka P., and Osher, S. 1994. A level set approach for computing solutions to incompressible two phase flow. *J. Comput. Phys.*, 119:146–159.
- Soatto S. and Yezzi, A. 2002. Deformation: Deforming motion, shape average and joint registration and segmentation of images. In *Proceedings of the 7th European Conference on Computer Vision*, Copenhagen, Denmark.
- Tagare, H.D. 1997. Deformable 2-D Template Matching Using orthogonal curves. *IEEE Trans. on Medical Imaging*, 16(1):108–117.
- Tao, Z., Beatty, J., Jaffe, C.C., and Tagare, H.D. 2002. Gary level models for segmenting myocardium and blood in cardiac ultrasound images. *Intl. Symp. on Biomedical Imaging*, Washington D.C.
- Tsai, A., Yezzi, Jr. A., and Willsky, A.S. 2001. Curve evolution implementation of the Mumford-Shah functional for image segmentation, denoising, interpolation, and magnification. *IEEE Trans. on Image Proc.*, 10(8):1169–1186.
- Thevenaz, P. and Unser, M. 2000. Optimization of mutual information for multiresolution image registration. *IEEE Trans. on Image Proc.*, 9(12):2083–2099.
- Viola P. and Wells III, W.M. 1997. Alignment by maximization of mutual information. *IJCV*, 24(2):137–154.
- Wang Y. and Staih, L. 1998. Boundary funding with corresponding using statistical shape models. In *Proc. IEEE Conf. CVPR*, pp. 338–345.
- Wells III, W.M., Viola, P., Atsumi, H., Nakajima, S., and Kikinis, R. 1996. Multi-Modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1(1):35–51.
- Yuille, A., Hallinan, P.W., and Cohen, D.S. 1992. Feature extration from faces using deformable templates. *IJCV*, 8:99–111.
- Zhao, H.K., Chan, T., Merriman B., and Osher, S. 1996. A variational level set approach to multiphase motion. *J. Comput. Phys.*, 127:179–195