

A variational framework for integrating segmentation and registration through active contours

A. Yezzi^a, L. Zöllei^{b,*}, T. Kapur^c

^aSchool of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA

^bArtificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA

^cGE Navigation and Visualization, Lawrence, MA, USA

Abstract

Traditionally, segmentation and registration have been solved as two independent problems, even though it is often the case that the solution to one impacts the solution to the other. In this paper, we introduce a geometric, variational framework that uses active contours to simultaneously segment and register features from multiple images. The key observation is that multiple images may be segmented by evolving a single contour as well as the mappings of that contour into each image.

© 2003 Elsevier Science B.V. All rights reserved.

Keywords: Segmentation-registration; Variational methods; Active contours

1. Introduction

Segmentation and registration have been established as important problems in the field of medical image analysis (Ayache, 1995; Cline et al., 1990; Grimson et al., 1994; Vannier et al., 1985). Traditionally, solutions have been developed for each of these two problems in relative isolation from the other, but with increasing dependence on the existence of a solution for the other. In the rest of this section, we discuss the interdependence of segmentation and registration solutions and introduce our motivation for a method that simultaneously estimates the two.

1.1. Dependence of registration on segmentation

A large class of registration solutions, referred to as ‘feature-based’ methods, require that some features be identified or segmented in the images prior to their registration. These features may be identified using low-level methods such as edge-detection, or segmented using higher level methods that are customized for specific anatomical structures. Contour- and point-based techniques (Tang et al., 2000; Weese et al., 1997a,b; Yaniv, 1998) are

examples of this strategy, as well as registration methods that compare medialness properties of segmented anatomies (Yushkevich et al., 1999). In contrast to feature-based registration methods, a second class of methods, referred to as ‘intensity-based’ segmentation methods, require no a priori segmentation, which makes them an attractive proposition. Some of the most frequently used objective functions in such registration frameworks are: normalized cross-correlation (Lemieux et al., 1994), entropy of the difference image (Buzug et al., 1997), pattern intensity (Weese et al., 1997b), gradient correlation (Brown, 1996) and gradient difference (Penney et al., 1998). Mutual-Information was introduced as a particularly effective intensity-based metric for registration of medical imagery (Collignon et al., 1995; Wells et al., 1995), and its applicability has been repeatedly demonstrated for solving rigid-body (6 degrees of freedom) registration problems. No such consensus, feature-based or intensity-based, seems to have been reached for the domain of non-rigid registration.

1.2. Dependence of segmentation on registration

The dependence of segmentation on registration is somewhat more subtle. A large class of segmentation methods do not depend on explicit registration between multiple data sets. We will refer to these as ‘low-level’

*Corresponding author. Tel.: +1-617-253-2986; fax: +1-617-258-6287.

E-mail address: lzollei@ai.mit.edu (L. Zöllei).

segmentation methods. In these low-level segmentation methods, the algorithm designers typically use information synthesized from their knowledge of several example data sets to set the parameters of their segmentation algorithms, but no explicit process of registering those data sets to a common reference frame is carried out prior to segmentation. These methods may process a single channel input image using image-processing techniques such as thresholding, connectivity analysis, region-growing, morphology, snakes, and Bayesian MAP estimation. Or, they may process multi-channel data in which the channels are naturally registered because they are acquired simultaneously.

While it is easier to get started in segmentation using these methods because there is no need to solve the cumbersome registration problem a priori, efforts in low-level segmentation of medical imagery often conclude that ‘model-based’, higher level information such as the shape, appearance, and relative geometry of anatomy needs to be incorporated into the solution in order to complete the segmentation task (Baillard et al., 2000; Cootes et al., 1994; Kapur et al., 1998; Staib and Duncan, 1992; Szekely et al., 1996). And it is in the building of these models of anatomy that registration plays a key role. Individual data sets need to be registered to a common frame of reference, so that statistics about their shape, appearance, or relative geometry can be gathered.

The work presented in this paper is motivated by the desire to interleave the process of segmentation and registration so that both solutions may be built simultaneously and hence to eliminate the need to completely deliver one solution before being able to start on the other. This challenge has been approached with a min–max entropy-based framework to segment and register portal images to CT (Bansal et al., 1999), and with the ATM SVC algorithm which applies an iterative sequence of elastic warping of the input to an already segmented model in order to automate the classification of normal and abnormal anatomy from medical images (Warfield et al., 2000). A novel extension to level set representations and active contour models by incorporating shape priors (Chen et al., 2001; Paragios and Rousson, 2002; Paragios et al., 2002) have also been recently introduced, which frameworks could potentially be used to address our proposed task.

The focus of this paper is to introduce a geometric, variational, active contour framework that allows us to interleave powerful level-set-based formulations of segmentation with a feature-based registration method.

2. Background on active contours

Active contours have been utilized extensively for problems including image segmentation, visual tracking, and shape analysis (see Blake and Isard, 1998 and refer-

ences therein). A variety of active contour models have been proposed since the introduction of the ‘snake’ methodology in the mid-1980s (Kass et al., 1987). These original models utilized parametric representations of the evolving contour. Shortly thereafter, using the level set methodology of Osher and Sethian (1988), more geometric techniques (such as those presented in (Malladi et al., 1995)) began to arise based upon the theory of curve evolution. An important class of these geometric models was derived via the Calculus of Variations to obtain evolution equations which would minimize energy functionals (or ‘objective functions’) tailored to features of interest in the image data. An in-depth discussion of variational image segmentation methods, as well an extensive list of references, may be found in the book (Morel and Solimini, 1995). The model that will be presented in this paper certainly fits within the context of these geometric variational approaches. However, we will exploit the calculus of variations to address not only the problem of image segmentation, but simultaneously the problem of image registration as well.

Most of the early active contour models for image segmentation, such as (Caselles et al., 1993, 1997; Cohen, 1991; Kass et al., 1987; Malladi et al., 1995; Tek and Kimia, 1995; Yezzi et al., 1997), were designed to capture localized image features, most notably edges. As such, these have come to be known as ‘edge-based’ models. In medical imaging and many other important applications where consistently strong edge information is not always present along the entire boundary of the objects to be segmented, the performance of purely edge-based models is often inadequate. In recent years, a large class of region-based models (such as Chakraborty et al., 1996; Chan and Vese, 1999; Paragios and Deriche, 1999; Paragios et al., 2002; Ronfard, 1994; Samson et al., 1999; Yezzi et al., 1999) have utilized image information not only near the evolving contour but also image statistics inside and outside the contour (in many ways inspired by the ‘Region Competition’ algorithm presented by Zhu and Yuille (1996)) in order to improve the performance.

There are still many cases in which both edge- and region-based active contour models have difficulty yielding correct segmentations of images that present rather subtle information about portions of the object to be captured. Significant improvement may be obtained in such cases by combining information from images of the same object acquired using different modalities (CT and MR, for example). However, to utilize the joint information, the various images must be correctly aligned to each other or ‘registered.’ If this can be done prior to segmenting any of the images, then *registration can assist segmentation*.

It is equally true, on the other hand, that *segmentation can assist registration*. It is typically much easier to align two images if the boundary of a common object or some other set of common point features have been accurately detected in both images beforehand. The images

may then be registered by point feature or contour matching. Furthermore, there may be cases in which registration is *impossible* (at least rigid registration) without some level of segmentation. This is the case when two (or more) images contain multiple common objects which may not be related by a single global mapping between the image domains. For example, an X-ray image of the femur and tibia may not be globally registered to a CT image of the femur and tibia if the knee is bent differently in the two images. Yet it is certainly possible to choose a registration which aligns the two femoral bones or a different registration which aligns the two tibial bones. In either case, though, it is necessary to segment the desired object from both images in order to perform the registration.

Next, we outline a geometric, variational framework for simultaneously segmenting and registering common objects in two or more images (the technical discussion will consider just two images, but the approach is easily adapted to multiple images). While our methodology is quite general and may certainly utilize any number of segmentation energy functionals, we focus our attention around region-based energy functionals; in particular, we will utilize the (Mumford and Shah, 1989) energy presented in (Chan and Vese, 1999).

3. General framework

In this section we outline the general framework for joint registration and segmentation via active contours. In Section 4, we will address rigid registration with scaling as a special case. Our model will be derived first for the two-dimensional case, and then the corresponding three-dimensional active surface model will be presented. We begin by establishing some basic notation.

3.1. Notation and problem statement

Let $I: \Omega \subset \mathcal{R}^2 \rightarrow \mathcal{R}$ and $\hat{I}: \hat{\Omega} \subset \mathcal{R}^2 \rightarrow \mathcal{R}$ denote two images that contain a common object to be registered and segmented, and let $g: \mathcal{R}^2 \rightarrow \mathcal{R}^2$ be an element of a finite dimensional group G with parameters g_1, \dots, g_n . We will denote by $\hat{\mathbf{x}} \in \hat{\Omega}$ the image of a point $\mathbf{x} \in \Omega$ under g (i.e., $\hat{\mathbf{x}} = g(\mathbf{x})$), and we will denote the Jacobian matrix of g by g' and its determinant (which we assume is positive) by $|g'|$.

Our goal may be stated as follows. We wish to find a closed curve $C \subset \Omega$ which captures the boundary of an object in image I , and another closed curve $\hat{C} \subset \hat{\Omega}$ which captures the boundary of the corresponding object in image \hat{I} . If C and \hat{C} were independent, this would simply be two segmentation problems. However, we will relate C and \hat{C} through a mapping $g \in G$.

$$C := g(C).$$

Our problem, then, is to find both a mapping g (which we will refer to from now on as the registration) and a contour C such that C and $\hat{C} = g(C)$ yield desirable segmentations of I and \hat{I} , respectively. In this manner, the segmentation and registration problems become coupled very naturally.

We will make use of the following additional notation. T, N and \hat{T}, \hat{N} will denote the unit tangents and normals to C and \hat{C} , respectively. In the same manner, $d\hat{\mathbf{x}}$ will denote the area measure $d\mathbf{x}$ (of Ω) pushed forward (onto $\hat{\Omega}$) by g , and $d\hat{s}$ will denote the arc length measure ds (of C) pushed forward (onto \hat{C}) by g . The relationships between these measures are given by $d\hat{\mathbf{x}} = \|g'\| d\mathbf{x}$ and by $d\hat{s} = \|g'T\| ds$. Finally, let $C_{\text{in}} \subset \Omega$ and $C_{\text{out}} \subset \Omega$ denote the regions inside and outside the curve C and let $\hat{C}_{\text{in}} \subset \hat{\Omega}$ and $\hat{C}_{\text{out}} \subset \hat{\Omega}$ denote the regions inside and outside the curve \hat{C} .

3.2. Energy functional

If we were charged with the task of segmenting image I and \hat{I} separately (i.e., without enforcing a relationship between C and \hat{C}), then we might choose from any number of geometric energy-based active contour models and would certainly be free to utilize two different models if the characteristics of image I and \hat{I} were sufficiently different. Let us refer to the energy functionals associated with these two models as E_1 and E_2 , respectively.

In order to discuss the problem in more detail, we must choose a specific form for E_1 and E_2 . Because of their wider capture range and greater robustness to noise, we prefer to focus our discussion around region-based energy functionals rather than edge-based energy functions; although, a similar development can be followed for almost any class of geometric active contour energies (even more sophisticated models that incorporate both edge and region measurements, shape priors, anatomical constraints, and other considerations).

A general class of region-based energies exhibit the following form:

$$E_1(C) = \int_{C_{\text{in}}} f_{\text{in}}(\mathbf{x}) d\mathbf{x} + \int_{C_{\text{out}}} f_{\text{out}}(\mathbf{x}) d\mathbf{x}, \quad (1)$$

$$E_2(\hat{C}) = \int_{\hat{C}_{\text{in}}} \hat{f}_{\text{in}}(\mathbf{x}) d\mathbf{x} + \int_{\hat{C}_{\text{out}}} \hat{f}_{\text{out}}(\mathbf{x}) d\mathbf{x}, \quad (2)$$

where the integrands f_{in} and f_{out} depend on I and where the integrands \hat{f}_{in} and \hat{f}_{out} depend on \hat{I} . If we introduce an artificial time variable, we obtain the following gradient evolutions for C and \hat{C} :

$$\frac{\partial C}{\partial t} = (f_{\text{in}} - f_{\text{out}})N \quad \text{and} \quad \frac{\partial \hat{C}}{\partial t} = (\hat{f}_{\text{in}} - \hat{f}_{\text{out}})\hat{N}. \quad (3)$$

For example, the piecewise-constant segmentation model of Chan and Vese (1999), which the authors utilized for the experiments in this paper, favors a curve which yields the least total squared error approximation of the image by

one constant inside the curve and another constant outside the curve. This yields the following particular choices for f_{in} , f_{out} , \hat{f}_{in} and \hat{f}_{out} ,

$$f_{in} = (I - u)^2, \quad f_{out} = (I - v)^2,$$

$$\hat{f}_{in} = (\hat{I} - \hat{u})^2 \quad \text{and} \quad \hat{f}_{out} = (\hat{I} - \hat{v})^2,$$

where u and v denote the mean values of I inside and outside C and where \hat{u} and \hat{v} denote the mean values of \hat{I} inside and outside \hat{C} .

By combining the selected energy functionals and enforcing the relationship $\hat{C} = g(C)$, we may formulate a joint energy that depends on g and C .

$$E(g, C) = E_1(C) + E_2(g(C))$$

$$= \int_{C_{in}} f_{in}(\mathbf{x}) \, d\mathbf{x} + \int_{C_{out}} f_{out}(\mathbf{x}) \, d\mathbf{x}$$

$$+ \int_{\hat{C}_{in}} \hat{f}_{in}(\mathbf{x}) \, d\mathbf{x} + \int_{\hat{C}_{out}} \hat{f}_{out}(\mathbf{x}) \, d\mathbf{x}. \quad (4)$$

We may re-express this energy using integrals only over the space Ω , which contains the contour C , as follows:¹

$$E(g, C) = \int_{C_{in}} (f_{in} + |g'| \hat{f}_{in} \circ g)(\mathbf{x}) \, d\mathbf{x}$$

$$+ \int_{C_{out}} (f_{out} + |g'| \hat{f}_{out} \circ g)(\mathbf{x}) \, d\mathbf{x}. \quad (5)$$

Now that task is to choose g and C in order to minimize (5). In doing so, we simultaneously segment both I and \hat{I} via C and \hat{C} as well as register the detected features (which are guaranteed to have the same detected shape since the contours C and \hat{C} will not be deformed independently) to each other through the mapping g .

Remarks. Obviously a weighted combination of E_1 and E_2 would be more general and useful in the event that one image is easier to segment than the other. However, to keep the development as clean and simple as possible, we will not include such weights. (We will follow a similar convention of ignoring weighting coefficients when we add curvature terms to the upcoming gradient flows.) A more significant point, though, is that (5) does not allow the registration g to be directly influenced by E_1 . This is a result of our arbitrary choice to let the unknown curve C live in the domain Ω of image I . A more symmetric arrangement would involve utilizing a separate domain for C and two mappings $g_1 \in G$ and $g_2 \in G$ to map C into Ω and $\hat{\Omega}$, respectively. Then the actual registration between the Ω and $\hat{\Omega}$ would be given by $g_2 \circ g_1^{-1}$. Once again, we have chosen to keep the presentation as simple as possible

by considering only one mapping g , which requires us to arbitrarily place the unknown curve in one of the two image domains.

3.3. Gradient flows

The most straightforward method for minimizing $E(C, g)$ is to start with an initial guess for both C and g and then evolve the contour C and the registration parameters g_1, g_2, \dots, g_n using a gradient flow.

The gradient evolution for the curve C may be obtained immediately by noticing that (5) has the same form as (1, 2). Thus, its gradient flow has the same form as (3). Simple substitution yields

$$\frac{\partial C}{\partial t} = (f + |g'| \hat{f} \circ g)N,$$

$$\text{where } f = (f_{in} - f_{out}) \quad \text{and} \quad \hat{f} = (\hat{f}_{in} - \hat{f}_{out}).$$

This flow, by itself, however, is not guaranteed to keep the evolving curve smooth. Thus, as is standard in most geometric active contour models, we will add a curvature (κ) term to the gradient flow (which arises if we add an arc length penalty to our energy functional) in order to regularize the curve evolution:

$$\frac{\partial C}{\partial t} = (f + |g'| \hat{f} \circ g)N - \kappa N. \quad (6)$$

The gradient evolutions for the registration parameters g_1, \dots, g_n depend upon the geometry of the curve \hat{C} and are given by

$$\frac{dg_i}{dt} = \frac{\partial E}{\partial g_i} = \int_{\hat{C}} \left\langle \frac{\partial \hat{\mathbf{x}}}{\partial g_i}, \hat{f}(\hat{\mathbf{x}}) \hat{N} \right\rangle ds$$

$$= \int_C \left\langle \frac{\partial}{\partial g_i} g(\mathbf{x}), \hat{f}(g(\mathbf{x})) \hat{N} \right\rangle \|g'T\| ds$$

$$= \int_C \hat{f}(g(\mathbf{x})) \left\langle \frac{\partial}{\partial g_i} g(\mathbf{x}), Jg'J^{-1}N \right\rangle ds \quad (7)$$

$$= \int_C \hat{f}(g(\mathbf{x})) \left\langle \frac{\partial}{\partial g_i} g(\mathbf{x}), \text{Adj}[g']N \right\rangle ds$$

$$= \int_C \hat{f}(g(\mathbf{x})) \left\langle \frac{\partial}{\partial g_i} g(\mathbf{x}), ((g')^{-1}|g'|)^T N \right\rangle ds.$$

(The last few steps use the fact that

$$\hat{N} = J \left(\frac{g'T}{\|g'T\|} \right) = \frac{Jg'J^{-1}}{\|g'T\|} N = \frac{\text{Adj}[g']}{\|g'N\|} N,$$

where J denotes the 90° rotation matrix and where $\text{Adj}[g']$ denotes the adjunct matrix of g' given by $((g')^{-1}|g'|)^T$.

In the 3D case, where S denotes an active surface (in place of the active contour C) and where \hat{S} denotes the transformed surface $\hat{S} = g(S)$, the registration evolution has the following similar form (where dA and $d\hat{A}$ denote the Euclidean area measures of S and \hat{S} and where T_u and T_v

¹The 'circle' operation in this equation stands for the standard symbol of function composition and does not indicate a dot product.

denote orthonormal tangent vectors such that $T_u \times T_v = N$):

$$\begin{aligned} \frac{\partial E}{\partial g_i} &= \int_s \left\langle \frac{\partial \hat{\mathbf{x}}}{\partial g_i}, \hat{f}(\hat{\mathbf{x}}) \hat{N} \right\rangle dA \\ &= \int_s \left\langle \frac{\partial}{\partial g_i} g(\mathbf{x}), \hat{f}(g(\mathbf{x})) \hat{N} \right\rangle \|g' T_u \times g' T_v\| dA \\ &= \int_s \hat{f}(g(\mathbf{x})) \left\langle \frac{\partial}{\partial g_i} g(\mathbf{x}), ((g') | g'|^T N) \right\rangle dA. \end{aligned}$$

The last step uses the fact that

$$\hat{N} = \frac{g' T_u \times g' T_v}{\|g' T_u \times g' T_v\|} = \frac{\text{Adj}[g'](T_u \times T_v)}{\|g' T_u \times g' T_v\|} = \frac{\text{Adj}[g'] N}{\|g' T_u \times g' T_v\|}.$$

3.4. The infinite dimensional (non-rigid/non-affine) case

So far, we have considered finite dimensional registration in the development of this coupled model. We will continue to develop specific finite dimensional cases (namely rigid and affine) in the following sections and demonstrate these cases in our experiments. However, at the helpful suggestion of the reviewers, we wish to take a moment to discuss how our approach may be formulated mathematically in the infinite dimensional case (i.e. neither rigid nor affine). The philosophy remains the same, namely we consider a single underlying contour C and a mapping g which, when applied to the contour C , yields a second contour $\hat{C} = g(C)$. However, if the mapping g is arbitrary, then the coupling between C and \hat{C} is effectively nonexistent. To see this, consider choosing \hat{C} independently of C in order to minimize the second term $E_2(\hat{C})$ in (4) while C is chosen to minimize the first term E_1 . Once the optimal C and \hat{C} are chosen in this manner, the two contours may ‘artificially’ be coupled by choosing g such that $\hat{C} = g(C)$. We therefore see that without imposing some structure on g , we are back to segmenting each image independently.

We may impose a ‘soft’ structure on g by penalizing the variation of g along the curve. For example, suppose we give a completely general form $g(C) = C + T$ where T is a translation vector which varies from point to point along the curve. In this case, the variation of T may be penalized by adding the following regularizing term E_3 to (4):

$$E_3 = \frac{\beta}{2} \int_c \|T_s\|^2 ds. \quad (8)$$

Computing the first variation of E_2 and E_3 with respect to T yields the following gradient flow for T :

$$\frac{\partial T}{\partial t} = \hat{f} \hat{N} + \beta \frac{\partial^2 T}{\partial s^2}. \quad (9)$$

Note that a significant sacrifice is made by replacing the explicit parametric structure of g with this softer penalty term E_3 in that the registration is now only defined

between the two contours C and \hat{C} rather than the between the entire image domains Ω and $\hat{\Omega}$. Thus, while our model philosophically generalizes in this manner, we feel it is practically better suited for rigid, affine, and other final dimensional forms of registration. In the remainder of this paper we will develop and demonstrate the rigid and affine cases.

4. ‘Affine’ registration

Notice that the gradient curve evolution (6) for C and the gradient direction (7) for the vector of registration parameters g_1, \dots, g_n both depend upon the Jacobian, g' of the registration map g . In the special case where G is the group of rigid-body motions following a (possibly nonuniform) scaling operation, then we may represent g by a rotation matrix R , a scaling matrix M and a displacement vector D :

$$g(\mathbf{x}) = R M \mathbf{x} + D. \quad (10)$$

Note, the fully affine case could be obtained by incorporating an additional shearing matrix into the above formulation. In the case described by (10), the Jacobian of g is independent of \mathbf{x} and is simply the product of the rotation and scaling matrices R and M . The determinant of this product equals the determinant of the scaling matrix, $m = |M| = (M_x M_y)$, thereby greatly simplifying both (6) and (7):

$$\frac{\partial C}{\partial t} = (f(\mathbf{x}) + m \hat{f}(g(\mathbf{x}))) N - \kappa N, \quad (11)$$

$$\frac{dg_i}{dt} = \int_c \hat{f}(g(\mathbf{x})) \left\langle \frac{\partial g(\mathbf{x})}{\partial g_i}, m R M^{-1} N \right\rangle ds. \quad (12)$$

4.1. The 2D case

In two dimensions, the rotation matrix R depends upon a single angle θ , the scaling matrix depends upon two scaling factors M_x and M_y , and the displacement vector D depends upon two offsets D_x and D_y in the x and y directions, respectively:

$$R = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix}, \quad M = \begin{bmatrix} M_x & 0 \\ 0 & M_y \end{bmatrix}, \quad D = \begin{bmatrix} D_x \\ D_y \end{bmatrix}.$$

The partial derivatives of $g(\mathbf{x})$ needed in (12) with respect to these five registration parameters are given by

$$\frac{\partial g(\mathbf{x})}{\partial \theta} = \begin{bmatrix} -\sin \theta & \cos \theta \\ -\cos \theta & -\sin \theta \end{bmatrix} \begin{bmatrix} M_x & 0 \\ 0 & M_y \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix},$$

$$\frac{\partial g(\mathbf{x})}{\partial D_x} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad \frac{\partial g(\mathbf{x})}{\partial D_y} = \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

$$\frac{\partial g(\mathbf{x})}{\partial M_x} = R \begin{bmatrix} x \\ 0 \end{bmatrix}, \quad \frac{\partial g(\mathbf{x})}{\partial M_y} = R \begin{bmatrix} 0 \\ y \end{bmatrix}.$$

4.2. The 3D case

In three dimensions, we evolve a surface S rather than a curve C and our registration g now represents a mapping from R^3 to R^3 . However, for the case of our special affine registration, g still has the form of (10). The rotation matrix R can be represented by a product of three separate rotation matrices R_α , R_β and R_γ which cause rotations around the x -, y - and z -axis, respectively. We refer to the corresponding angles as roll (α), pitch (β) and yaw (γ). The scaling matrix M contains three components M_x , M_y and M_z and the translation vector D depends upon three offsets D_x , D_y and D_z in the x , y and z directions, respectively:

$$R = R_\gamma R_\beta R_\alpha,$$

$$R_\alpha = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha & -\sin \alpha \\ 0 & \sin \alpha & \cos \alpha \end{bmatrix},$$

$$R_\beta = \begin{bmatrix} \cos \beta & 0 & \sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & \cos \beta \end{bmatrix},$$

$$R_\gamma = \begin{bmatrix} \cos \gamma & -\sin \gamma & 0 \\ \sin \gamma & \cos \gamma & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

$$M = \begin{bmatrix} M_x & 0 & 0 \\ 0 & M_y & 0 \\ 0 & 0 & M_z \end{bmatrix} \quad \text{and} \quad D = \begin{bmatrix} D_x \\ D_y \\ D_z \end{bmatrix}.$$

The partial derivatives of $g(\mathbf{x})$ with respect to these nine registration parameters are given by

$$\frac{\partial g(\mathbf{x})}{\partial \alpha} = R_\gamma R_\beta R'_\alpha M \mathbf{x} + D,$$

$$\frac{\partial g(\mathbf{x})}{\partial \beta} = R_\gamma R'_\beta R_\alpha M \mathbf{x} + D,$$

$$\frac{\partial g(\mathbf{x})}{\partial \gamma} = R'_\gamma R_\beta R_\alpha M \mathbf{x} + D,$$

$$\frac{\partial g(\mathbf{x})}{\partial D_x} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad \frac{\partial g(\mathbf{x})}{\partial D_y} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad \frac{\partial g(\mathbf{x})}{\partial D_z} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix},$$

$$\frac{\partial g(\mathbf{x})}{\partial M_x} = R \begin{bmatrix} x \\ 0 \\ 0 \end{bmatrix}, \quad \frac{\partial g(\mathbf{x})}{\partial M_y} = R \begin{bmatrix} 0 \\ y \\ 0 \end{bmatrix}, \quad \frac{\partial g(\mathbf{x})}{\partial M_z} = R \begin{bmatrix} 0 \\ 0 \\ z \end{bmatrix}.$$

These derivatives are utilized to update the registration parameters via the coupled flow for the surface S and the registration g according to the following equations (analogous to (11) and (12) for the 2D case):

$$\frac{\partial S}{\partial t} = (f(\mathbf{x}) + m\hat{f}(g(\mathbf{x})))N - HN,$$

$$\frac{dg_i}{dt} = \int_s \hat{f}(g(\mathbf{x})) \left\langle \frac{\partial g(\mathbf{x})}{\partial g_i}, mRM^{-1}N \right\rangle,$$

where $m = |M| = (M_x M_y M_z)$ and H and dA denote the mean curvature and area element of the surface S (N and \hat{N} denote the unit normal of S and $S = g(S)$ just as in the 2D case for curves).

5. Results

In this section, we report segmentation/registration results from three experiments on MRI/CT and one on MR/MR images of the head and the spine. The first experiment was performed in 2D, while the second, third and the fourth ones were performed in 3D. In the 2D experiment, corresponding slices between the MR and the CT were chosen manually, and used as input for our algorithm. In the 3D experiments, a pair of 3D MR and CT or a pair of MR scans was used as input, without any attempt to manually initialize the registration. In all four cases, validation is performed by visual inspection of the results. Quantitative analysis of the algorithm performance is done in the following section.

5.1. 2D MR-CT head experiment

Input: In this experiment, the input consists of two 2D images of the head. The first input image is a single, axial cross-section from a 3D, gradient echo MRI scan (top row of Fig. 1), and the second image is the (manually chosen) corresponding cross-section from a 3D CT scan (bottom row of Fig. 1).

Goal: The goal of the joint segmentation-registration experiment is to simultaneously segment the skin surface and register the two slices.

Initialization: A curve is initialized on the MR image, such that it lies within the head images. This initial curve is shown in the left column of Fig. 1. We choose initial registration parameters such that they map this initial curve inside the head of the CT image but clearly not at the ‘corresponding location’.

Outcome: The final joint segmentation is shown in the right column of Fig. 1. Notice that the contour accurately outlines the skin surface in both the MR and CT images and also that as the contour evolves (left to right) within the MR image (top), its rigid transformation into the CT image (bottom) is also evolving.

5.2. 3D MR-CT spine experiment

Input: In this section, we present results from two similar experiments. In both cases, the input consists of a 3D CT and a 3D MR data volume of the spine. In the first experiment, the acquisition orientation was sagittal while in the second it was axial. The results are shown from a sagittal view in both cases.

Goal: The goal of the joint segmentation-registration experiment is to segment a single vertebra while comput-

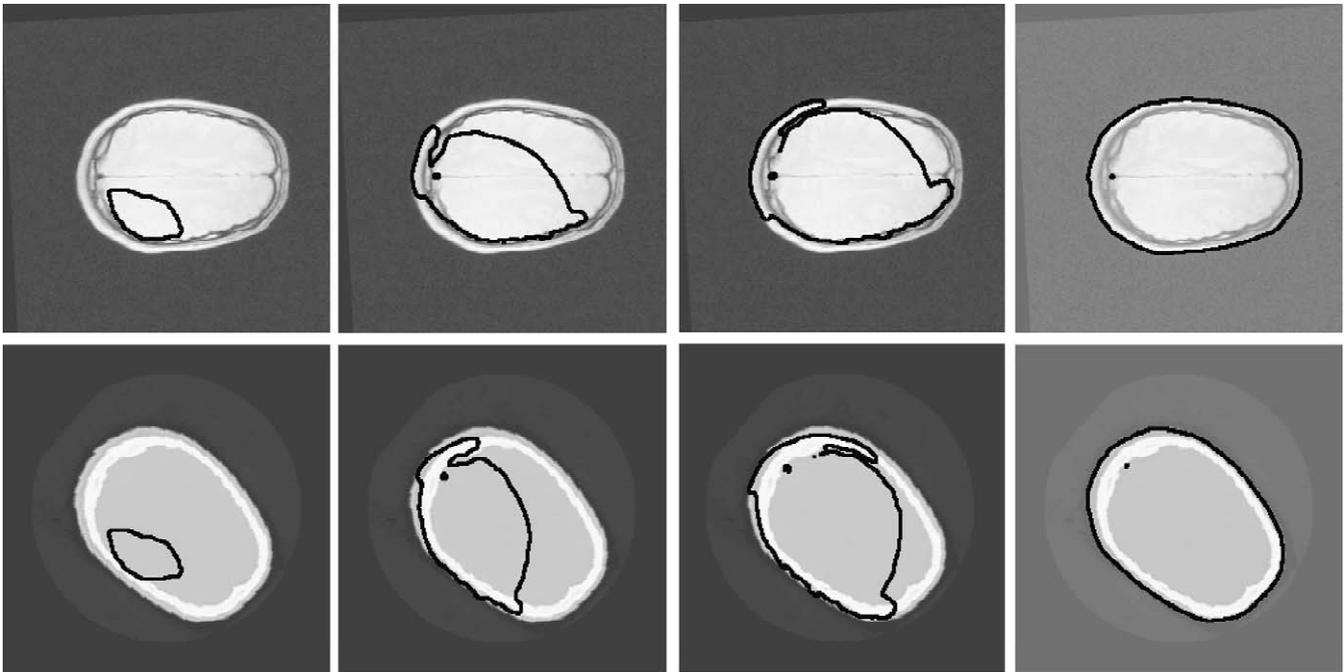


Fig. 1. Registration/segmentation of 2D MR/CT head data: initial (left), intermediate (middle) and final (right) results. The top row shows the evolving contour overlaid on the MR slice, and the bottom row shows it overlaid on the CT slice. Note that the rightmost column shows that the contour has correctly identified the skin boundary in both the CT and the MR.

ing the rigid transform that registers the two corresponding vertebrae in the different modalities. Note that the two spine images could not be registered, as a whole, by a single rigid transform.

Initialization: A surface is initialized within one of the vertebrae in the MR slice (shown in the first row of Fig. 2)

and the top rows of the bottom section of Fig. 3. The initial registration parameters map this initial surface into the neighborhood of the same vertebra in the CT image, but not exactly to the corresponding position.

Outcome: The segmentation component of the result is shown in the last rows of Figs. 2 and 3. Notice, that in both

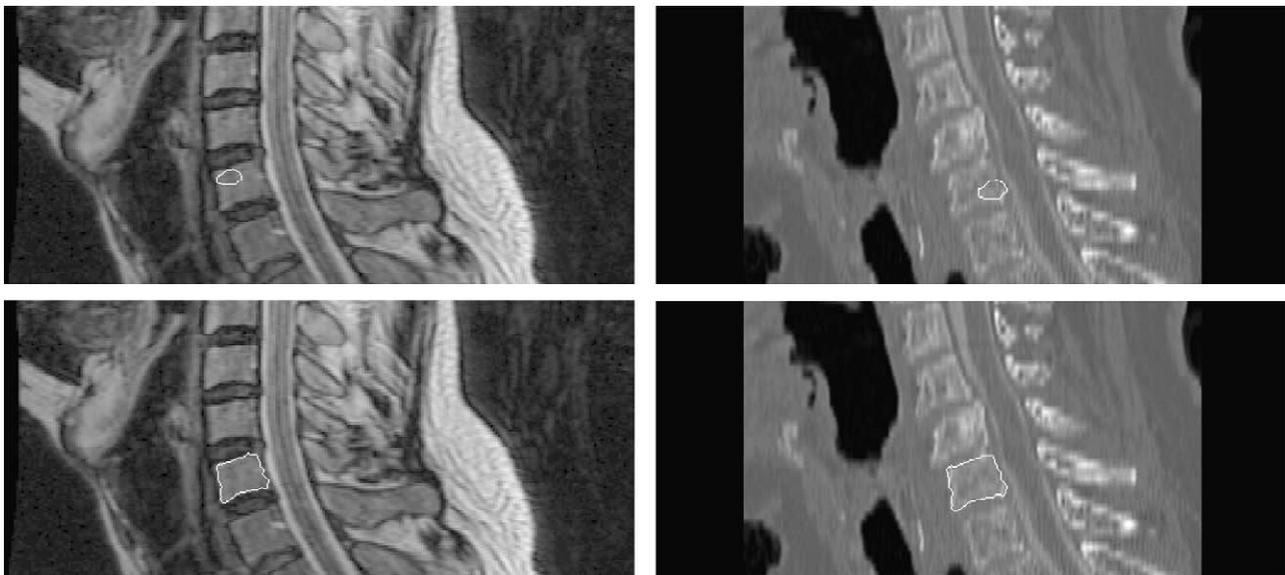


Fig. 2. Registration/segmentation of MR and CT Spine Images: initial (top) and final (bottom) results of Experiment 1. The left column shows MR cross-sections of the spine, and the right column shows CT slices of the same subject. In the top row, the initial contour is located inside the vertebra of interest, overlaid on both the MR and the CT acquisitions. Note the poor contrast around the vertebra of interest in the MR image, as well as the fact that the transform between the two spine volumes is not rigid (the spine curves differently in the two images). The bottom rows show the cross-section of the final surface, which has captured the boundary of the vertebra in the MR and CT scans.

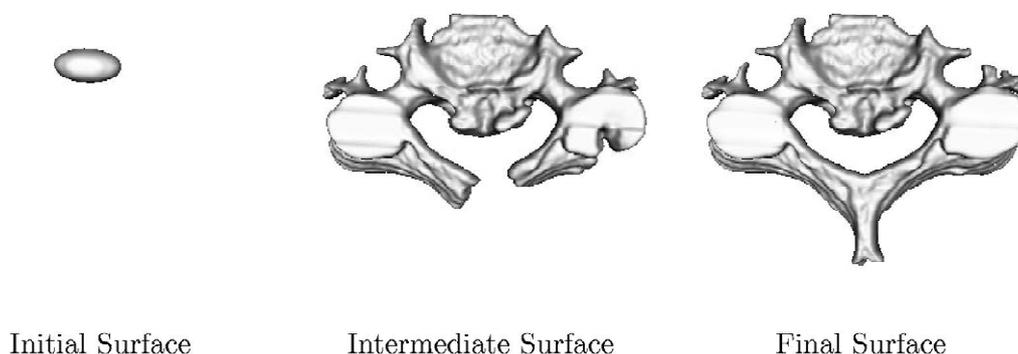


Fig. 3. Registration/segmentation of 3D CT and 3D MR spine images. Experiment 2. The figure shows the evolving 3D surface.

cases, the contour accurately outlines the boundary of the vertebra in both of the MR and the CT slices. By segmenting/registering each vertebra in this manner, the change in the curvature of the spine could be estimated between two scans. A 3D model of the developing surface has also been constructed for the second experiment (first row of Fig. 3). The results of these experiments best illustrate the power of our joint segmentation/registration procedure.

5.3. 3D CT-MR head experiment

Input: In this experiment, the input consists of two 3D data sets of the head, one MR, one CT. Each data set contains 23 slices of size 256×256 . The middle row of Fig. 4 shows an axial slice from the MR data set, and the bottom row shows the corresponding slice from the CT data set.

Goal: The goal of the joint segmentation/registration experiment is to segment the 3D skin surface while registering the two 3D volumes.

Initialization: A surface is initialized such that it lies outside the head in the MR image. This initial surface is shown in the top left corner of Fig. 4.

Outcome: The segmentation component of the result is presented in two different forms in Fig. 4. The first row (images a–c) presents the evolving 3D surface at the initial, at an intermediate and at the final stages of the experiment. The second and third rows demonstrate cross-sections of the surface overlaid on axial MR and CT slices, respectively. In each row, the first column shows the initial, the second column an intermediate, and the third column the final state of the surface model intersected with the MR and CT slices.

5.4. 3D CT-MR ventricle experiment

Input: In this experiment, the input consists of two MR 3D data sets of the same patient, taken 1 month apart. The data volumes are of different dimensionality and their voxels are also of different scales. The second row of Fig.

5 presents a set of two axial MR slices taken from the two different datasets (from different views).

Goal: The goal of the joint segmentation-registration experiment is to segment the surface of the ventricle while registering the two 3D data sets.

Initialization: A surface is initialized such that it contains the ventricle in the MR images. This initial surface is shown in the top left corner of Fig. 5.

Outcome: The segmentation component of the result is presented in two different forms in Fig. 5. The first row presents the evolving 3D surface at the initial, at an intermediate and at the final stages of the experiment. The second row demonstrates cross-sections of the surface overlaid on axial MR slices. In this row, the first column shows the initial, the second column an intermediate, and the third column the final state of the surface model intersected with the MR slices. This experiment is an excellent example for a scenario where it is essential to recover not only a rigid-body movement, but the scaling parameters of the transformation, too.

6. Validation experiments

6.1. Images used for validation

In the following validation experiments, we use a set of synthetic images, which are displayed in Fig. 6. Besides the original binary image ('Original Image'), we created two other images by adding different amounts of Gaussian noise to the former. In both cases the distorting noise is zero-mean, and one has 0.05 variance while the other has 0.5. In the case of Image A, we also applied a rigid transformation to one of the objects present.

6.2. How registration is aided by segmentation

In order to demonstrate in what manner segmentation results aid the registration procedure, we ran two different types of experiments. The first experiment demonstrates the fact that without segmentation some of the above-introduced images cannot be successfully registered. That

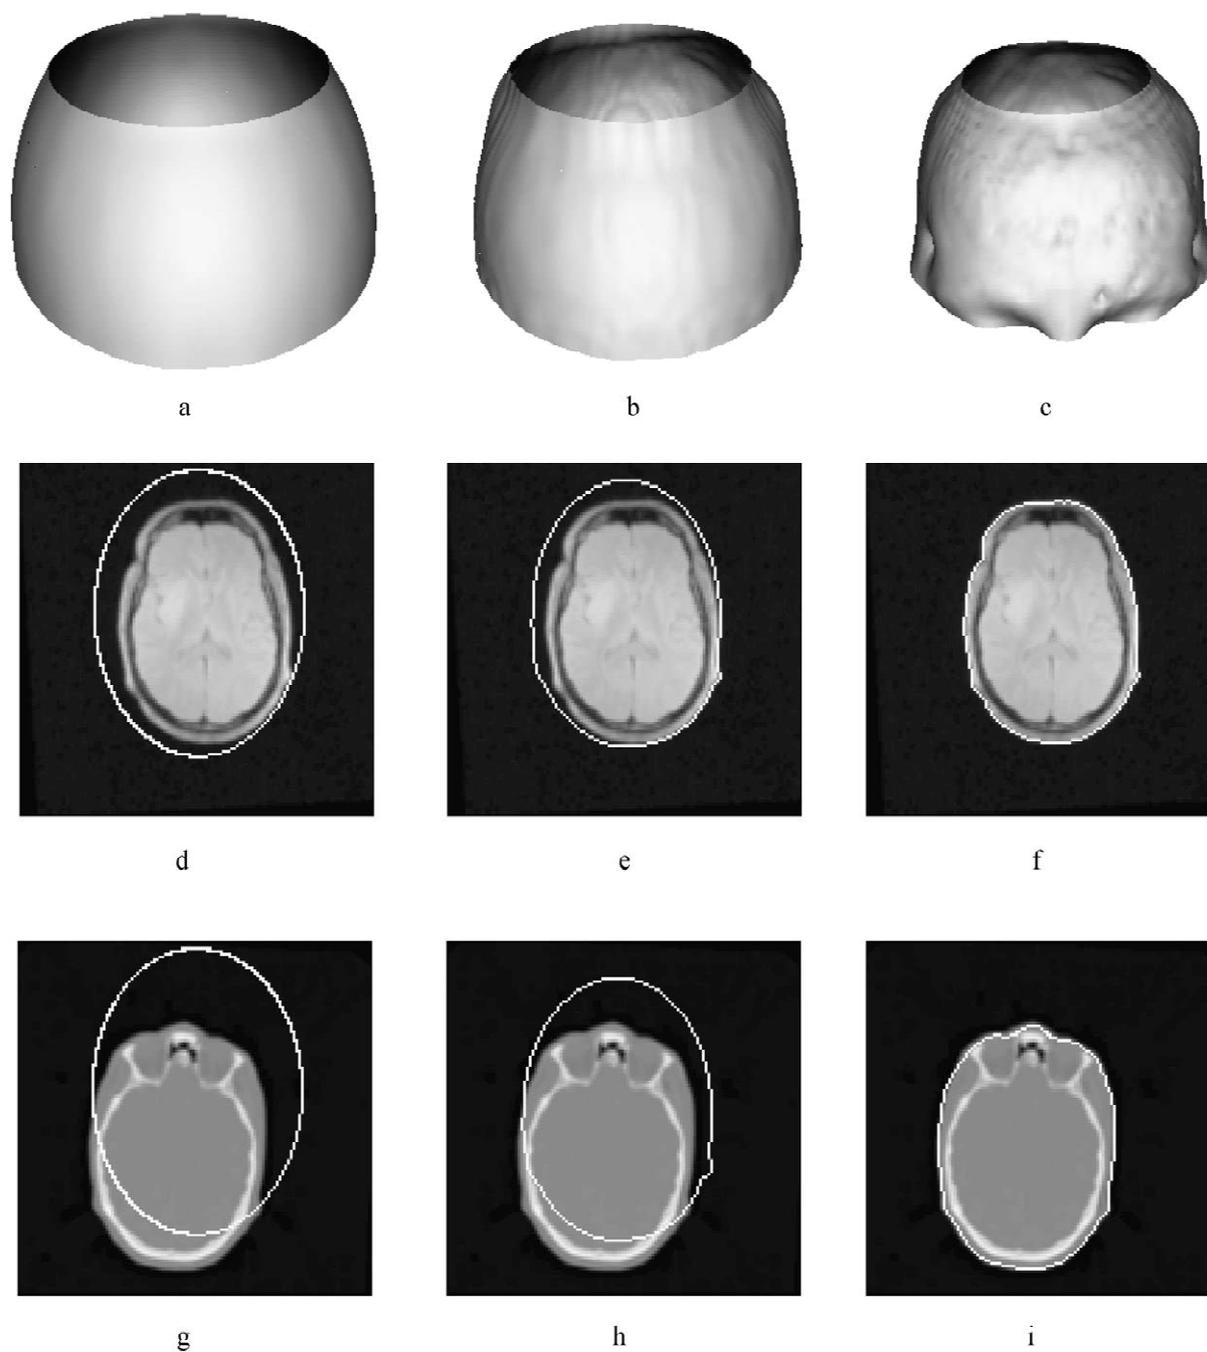


Fig. 4. Registration/segmentation of 3D MR and CT head images: initial (left), intermediate (middle) and final (right) results. The top row (a–c) shows the evolving surface. The bottom two rows shows cross-sections of the evolving surface overlaid on an MR (d–f) and CT slice (g–i), respectively.

is true, for example, in the case of Image A and Image B. Given that the rigid transformation should not be applied to the whole input images, but merely to certain parts of it, the registration results are incorrect. We ran registration experiments implementing a joint entropy minimization objective function with a stochastic gradient descent optimization, but the registration error was always greater than 15 mm.

In the second set of experiments, we used the same input images as above. This time, however, we first ran a

segmentation algorithm on the input images, which was targeted to the objects whose position difference was to be recovered. In addition, we also added a regularization term to the segmentation algorithm. That improved the quality of segmentation results on the noisy inputs by making the resulting boundaries smoother (see the input images on Fig. 7). Using these segmented contours as registration inputs, a matching rigid transformation could more accurately be computed. The average error of the registration experiments was (4.2544, 1.2305) mm, respectively, in the

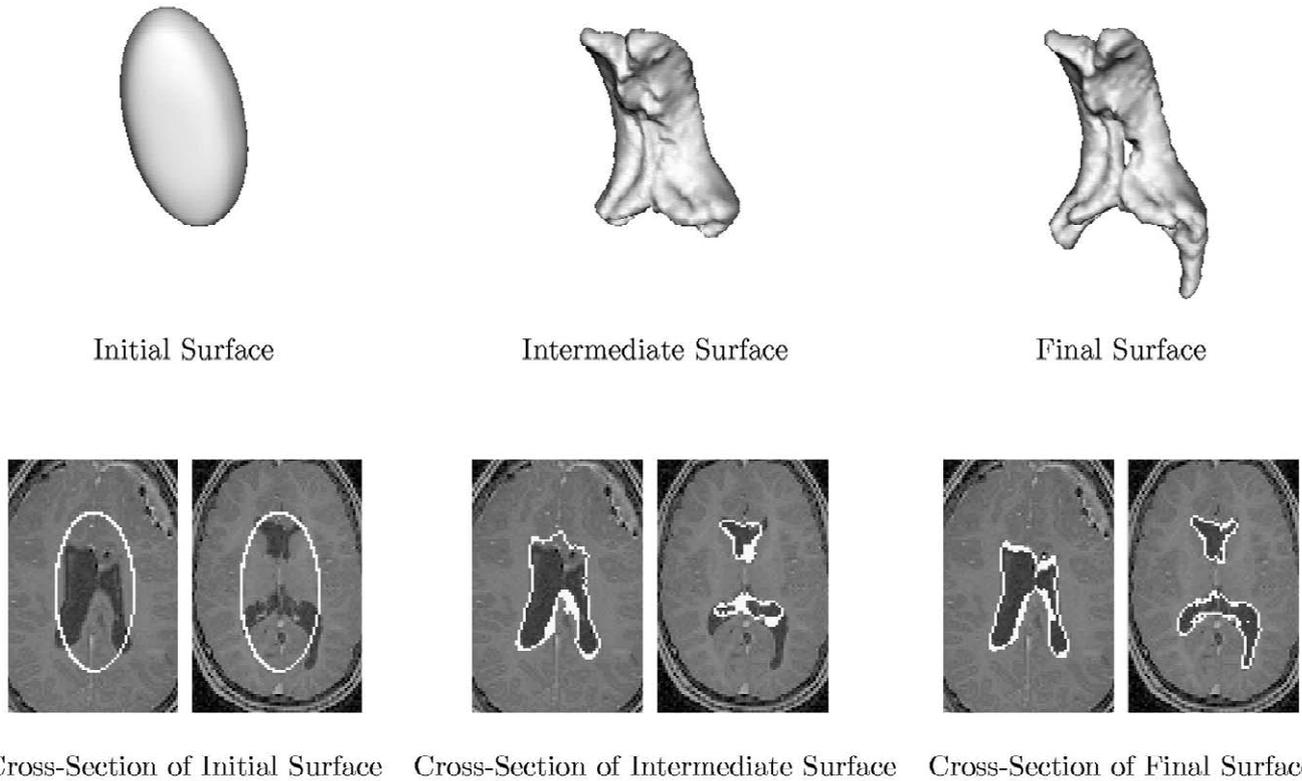


Fig. 5. Registration/segmentation of 3D MR ventricle images: initial (left), intermediate (middle) and final (right) results. The top row shows the evolving surface. The bottom row presents cross-sections of the evolving surface overlaid on the MR slices from the two different volumes.

x and y directions. We should note though that the success of the algorithm is largely affected by the accuracy of the segmentation quality. Although the higher the regularization term we add to the segmentation algorithm, the smoother but less precise the boundaries of the targeted

object become (see Fig. 8). Sharp corners of the targeted objects are often cut off, which can lead to significant degradation of the registration results.

Thus we can conclude that registration algorithms do benefit from segmentation results in certain applications.

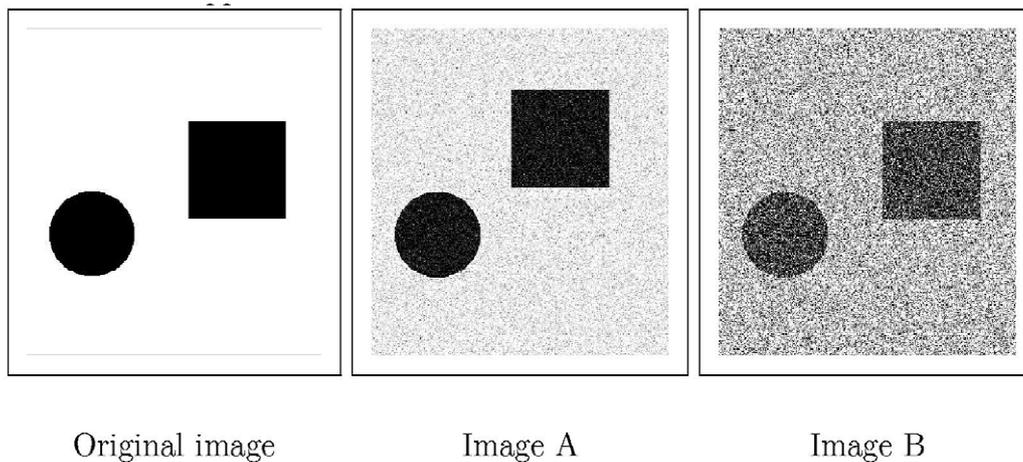


Fig. 6. Synthetic images serving as inputs to the validation experiments. In the case of Image A, the added Gaussian noise is zero-mean with 0.05 variance and in the case of Image B, the noise is zero-mean with 0.5 variance. There is a known rigid transformation applied to the square of Image A.

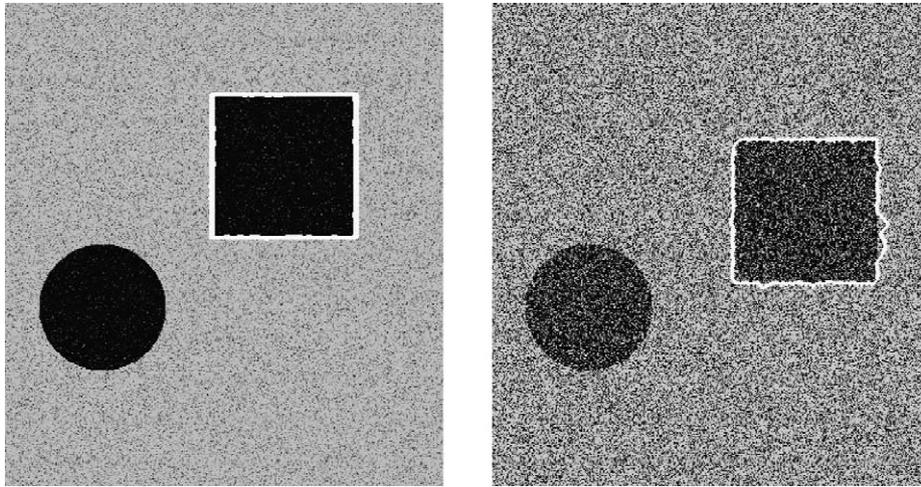


Fig. 7. Input images to the registration validation experiment where the objects of interest were pre-segmented. In the case of Image B (right), a regularization term was added to the segmentation criterion to obtain smoother boundaries.

6.3. How segmentation is aided by registration

We carried out some experiments using the same synthetic datasets in order to demonstrate how segmentation can benefit from the coupled registration process.

In the first round of experiments, we only carried out the segmentation process, separately on the two input images: Image A and Image B. While in the case of Image A the segmentation could be carried out with good results, in the case of Image B the segmentation did not succeed because of the high level of noise. The segmentation boundary remained very rough and numerous short contours were introduced (see Fig. 9). Note, that in these experiments we did not use such a powerful regularization term as demon-

strated in Fig. 8, in order to exactly match later the parameter settings of the segmentation/registration coupled approach.

When, however, these two images were segmented simultaneously, using segmentation coupled with the registration process, proper segmentation results could be accomplished (see Fig. 10). After terminating the experiments, the contours correctly outlined the borders of interest. The registration results were within 0.12 mm accuracy.

6.4. Registration accuracy

For the registration validation experiments, we ran three

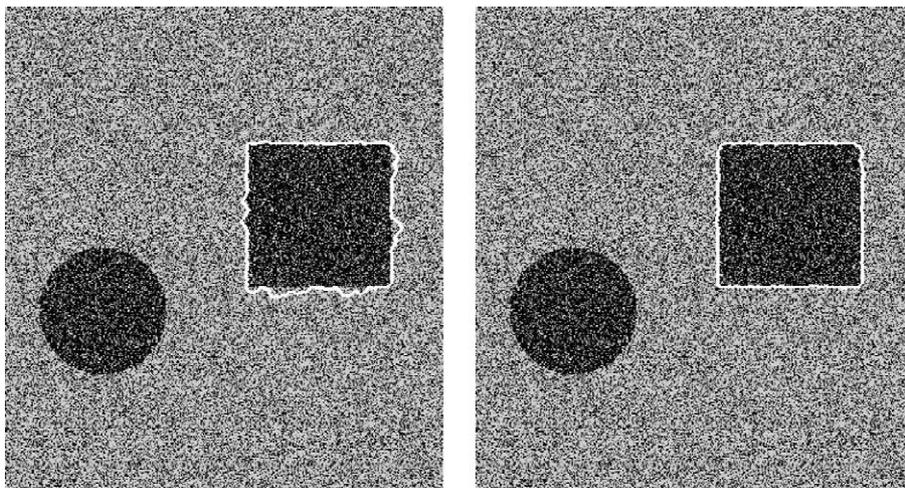


Fig. 8. Object boundaries with (left) a lower and (right) a higher regularization parameter.

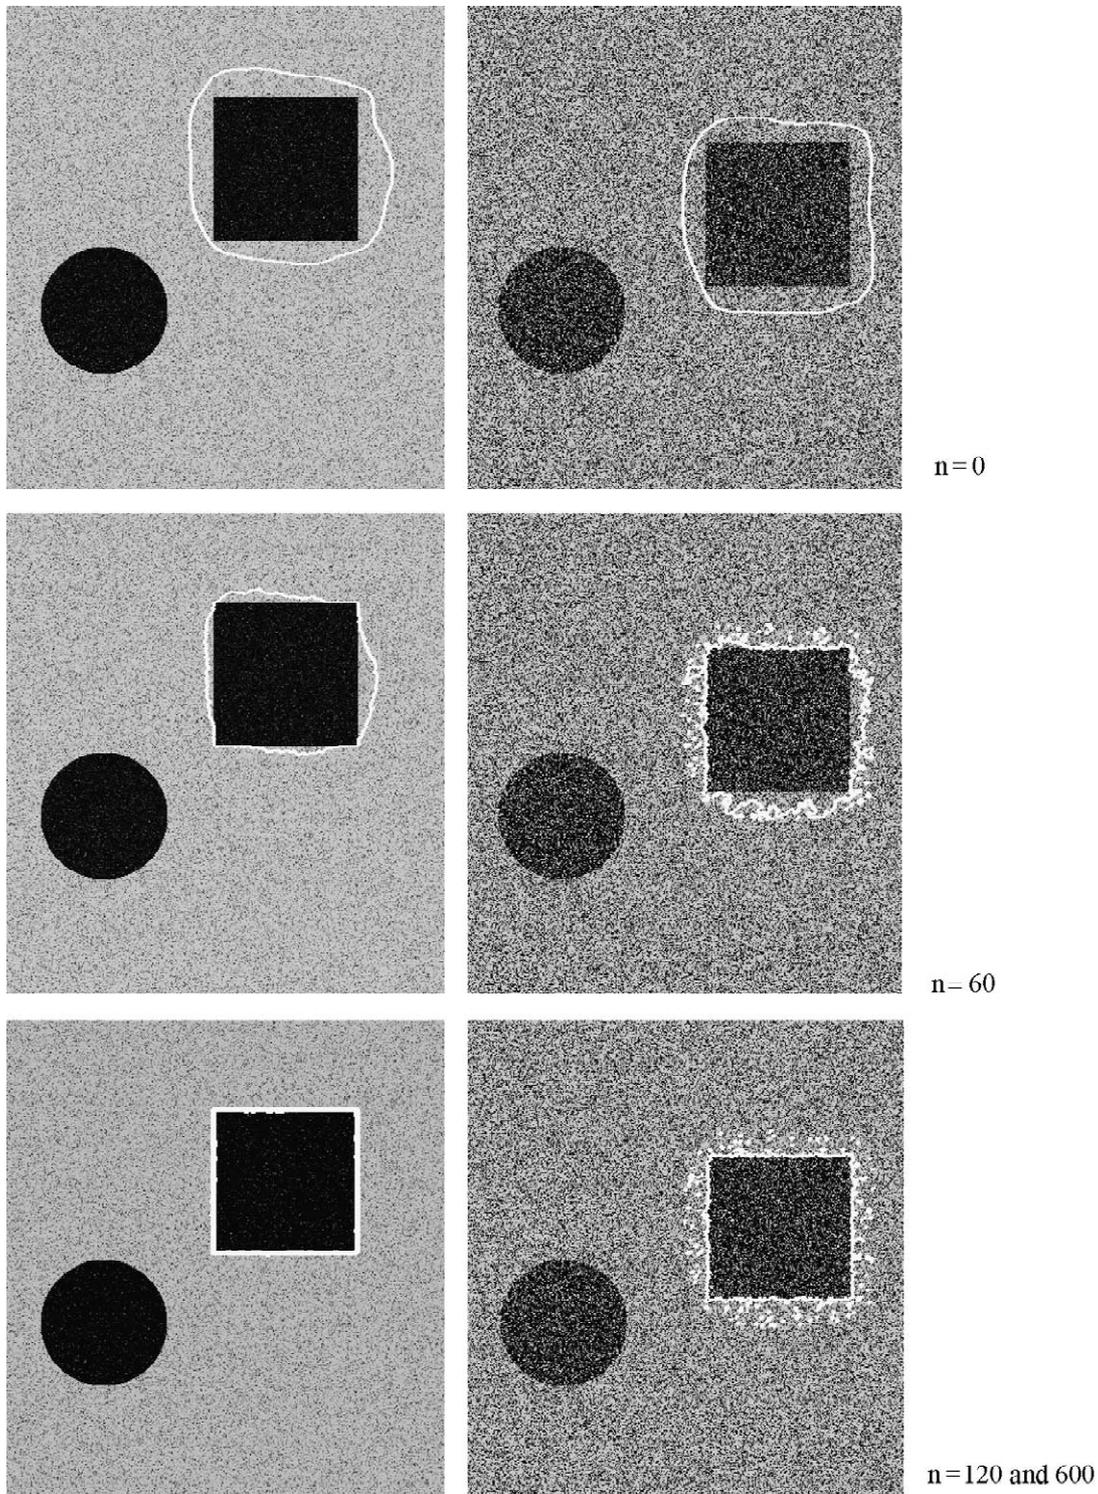


Fig. 9. Individual segmentation experiments on Image A (first column of images) and Image B (second column of images) separately. Note, the same level of regularization was used here as later on in the case of the coupled registration and segmentation method (Fig. 10).

different sets of experiments (Experiment I, II, III). The first two sets used synthetic images as inputs: Image A and Image B. The third set was run on real MR-CT slices of a head (see Fig. 11). In both cases ground truth transforma-

tion parameters were given (see Fig. 11). In each set of experiments, we executed the algorithm several times, each time with a different initial contour which was either within or outside of the object boundary to be segmented.

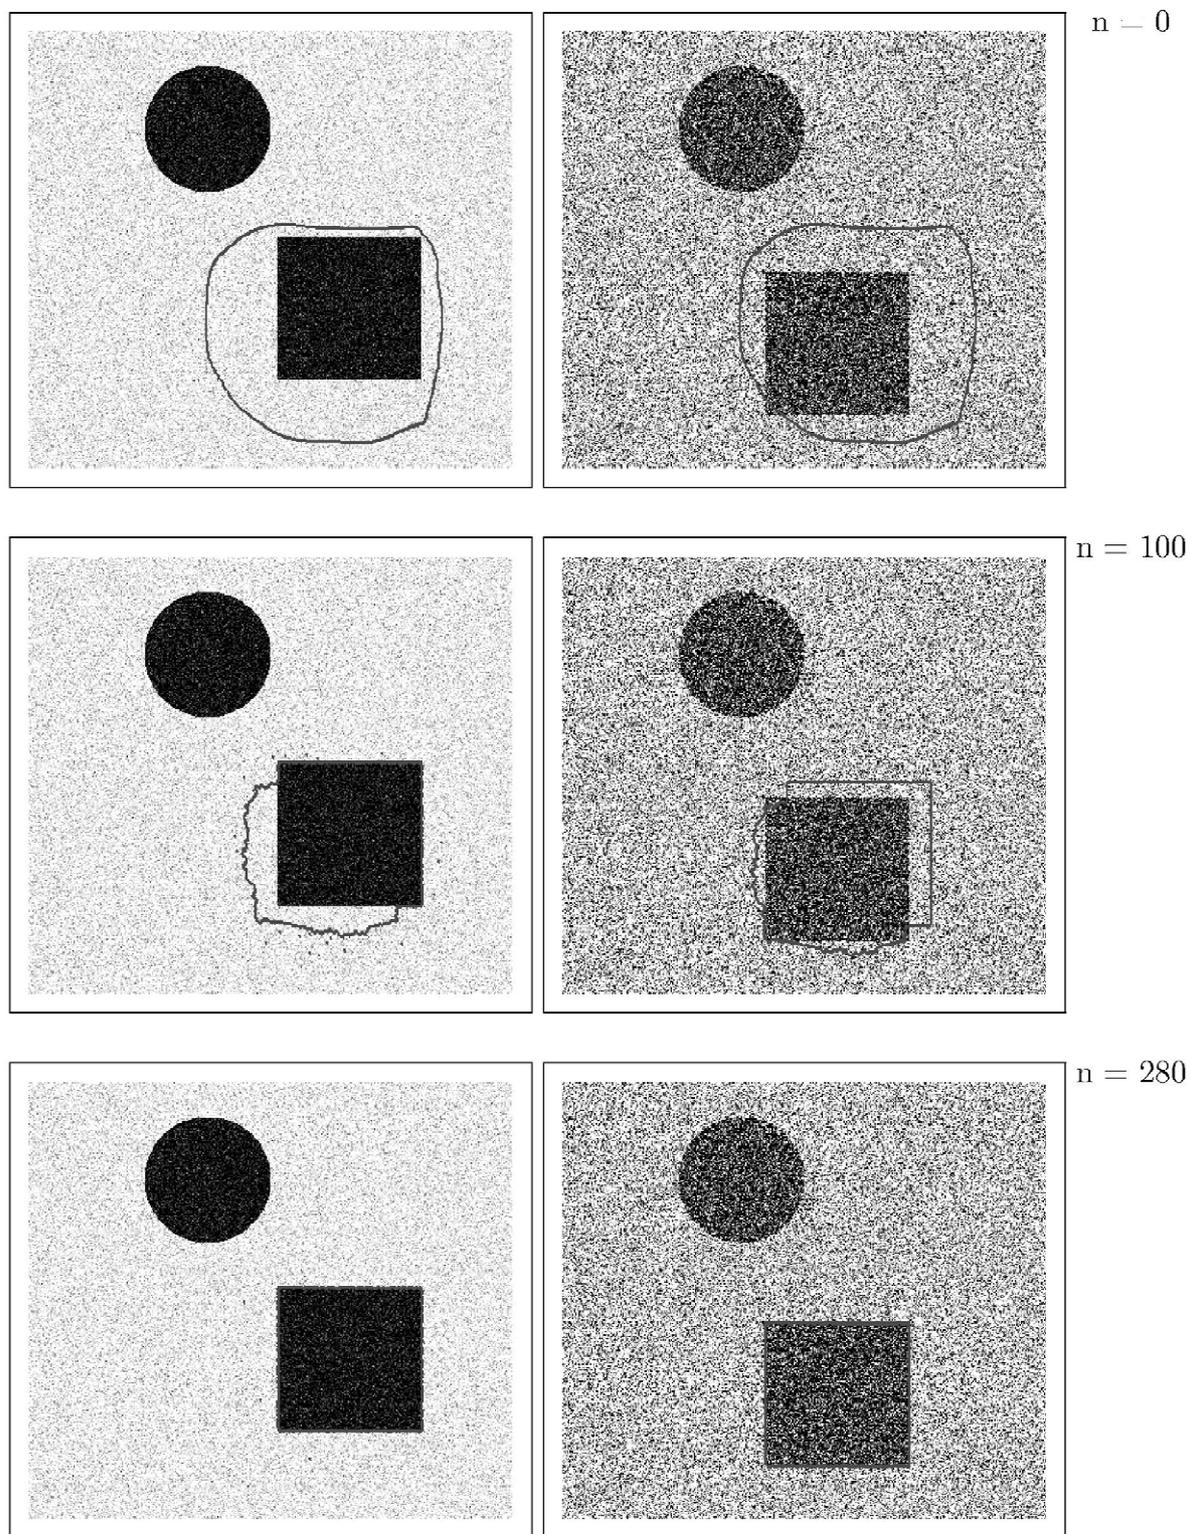


Fig. 10. Registration and segmentation are coupled to obtain the desired task. The images shown in a row were treated together. Note, the same level of regularization was used here as earlier in the case of the individual segmentation procedures (Fig. 9).

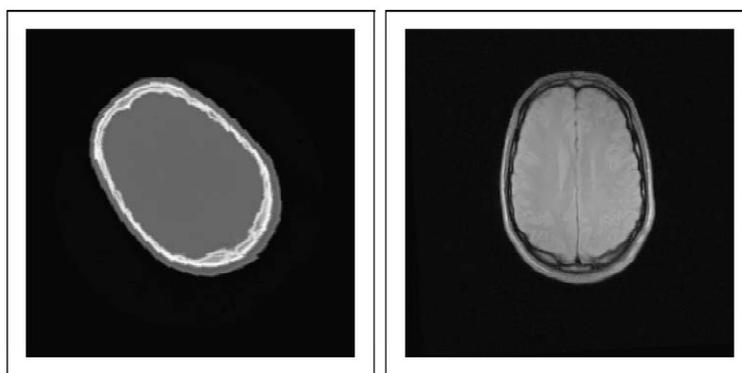


Fig. 11. CT and MR image slices used as input for the registration validation experiments.

Table 1
Registration accuracy results

Experiments	No. of runs	Error measures	Mean (error)	Variance (error)
I	30	Translation (mm)	(0.4970, 0.5873)	(0.0418, 0.0351)
		Rotation (rad)	NA	NA
II	22	Translation (mm)	(0.3432, 0.4586)	(0.0426, 0.0261)
		Rotation (rad)	NA	NA
III	16	Translation (mm)	(−0.0944, 0.0008)	(0.1206, 0.0386)
		Rotation (rad)	0	0

The registration accuracy was very good. Translation errors, in general, were less than 0.5 mm and the rotation errors were insignificant. The results—the mean and variance of the error measures—are summarized in Table 1.

7. Summary and future work

We have presented a variational framework for joint segmentation and registration using active contours. We employ a single contour (or surface in 3D) to segment multiple images. The contour and the registration are both computed to minimize a set of energy functionals, one for each image. The experiments in this paper utilize an intensity-based energy functional, but the framework allows for richer choices that may encode shape priors, textures, or other image statistics, which we are currently exploring.

Acknowledgements

This work was supported by the Whiteman Fellowship and the NSF grant #CCR-0133736.

References

Ayache, N., 1995. *Medical Computer Vision, Virtual Reality and Robotics, Image and Vision Computing*, pp. 295–313.

- Baillard, C., Hellier, P., Barillot, C., 2000. Segmentation of 3D Brain Structures using Level Sets, Publication interne, IRISA, no. 1291.
- Bansal, R., Staib, L.H., Chen, Z., Rangarajan, A., Knisely, J., Nath, R., Duncan, J., 1999. Entropy-based, multiple-portal-to-3D CT registration for prostate radiotherapy using iteratively estimated segmentation. In: *Proceedings of MICCAI*, pp. 567–578.
- Blake, A., Isard, M., 1998. In: *Active Contours*. Springer, London.
- Brown, L.M.G., 1996. Registration of Planar Film Radiographs with Computed Tomography. In: *Proceedings of MMBIA*.
- Buzug, T.M., Weese, J., Fassnacht, C., Lorenz, C., 1997. Image registration: convex weighting functions for histogram-based similarity measures. In: *Proceedings of CVRMed/MRCAS*. Springer, Berlin, Germany, pp. 203–212.
- Caselles, V., Catta, F., Coll, T., Dibos, F., 1993. A geometric model for active contours in image processing. *Numerische Mathematik* 66, 1–31.
- Caselles, V., Kimmel, R., Sapiro, G., 1997. Geodesic active contours. *Int. J. Computer Vision* 22 (1), 61–79.
- Chakraborty, A., Staib, L., Duncan, J., 1996. Deformable boundary finding in medical images by integrating gradient and region information. *IEEE Trans. Medical Imaging* 15 (6), 859–870.
- Chan, T., Vese, L., 1999. An Active Contour Model Without Edges. In: *Int. Conf. Scale-Space Theories in Computer Vision*, pp. 141–151.
- Chen, Y., Thiruvankadam, S., Tagare, H., Huang, F., Wilson, D., Geiser, E.A., 2001. On the Incorporation of Shape Priors into Geometric Active Contours. In: *IEEE Workshop on Variational and Level Set Methods*, pp. 145–152.
- Cline, H., Lorensen, W., Kikinis, R., 1990. Three-dimensional segmentation of MR images of the head using probability and connectivity. *Journal of Computer-Aided Tomography* 14 (6), 1037–1045.
- Cohen, L., 1991. On active contour models and balloons. *CVGIP: Image Understanding* 53, 211–218.
- Collignon, A., Vandermeulen, D., Suetens, P., Marchal, G., 1995. Automated multi-modality image registration based on information theory. In: *Proceedings of IPMI*, pp. 263–274.
- Cootes, T.F., Hill, A., Taylor, C.J., Haslam, J., 1994. Use of active shape models for locating structure in medical images. *Image and Vision Computing* 12 (6), 355–365.

- Grimson, W.E.L., Lozano-Perez, T., Wells, III W.M., Ettinger, G.J., White, S.J., 1994. An automatic registration method for frameless stereotaxy, image, guided surgery and enhanced reality visualization. In: *Proceedings of CVPR*, pp. 430–436.
- Kapur, T., Grimson, W.E.L., Kikinis, R., Wells, III W.M., 1998. Enhanced spatial priors for segmentation of magnetic resonance imagery. In: *Proceedings of MICCAI*, pp. 457–468.
- Kass, M., Witkin, A., Terzopoulos, D., 1987. Snakes: active contour models. *Int. J. Computer Vision* 1, 321–331.
- Lemieux, L., Jagoe, R., Fish, D.R., Kitchen, N.D., Thomas, D.G.T., 1994. A patient-to-computed-tomography image registration method based on digitally reconstructed radiographs. *Med. Phys.* 21 (11), 1749–1760.
- Malladi, R., Sethian, J., Vemuri, B., 1995. Shape modeling with front propagation: a level set approach. *IEEE Trans. PAMI* 17, 158–175.
- Morel, J.M., Solimini, S., 1995. In: *Variational Methods in Image Segmentation*. Birkhauser, Boston, MA.
- Mumford, D., Shah, J., 1989. Optimal approximations by piecewise smooth functions and associated variational problems. *Commun. Pure and Appl. Mathematics* 42 (5), 577–685.
- Osher, S., Sethian, J., 1988. Fronts propagation with curvature dependent speed: algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics* 79, 12–49.
- Paragios, N., Deriche, R., 1999. Geodesic active regions for supervised texture segmentation. *Proceedings of ICCV* 1, 926–932.
- Paragios, N., Rousson, M., 2002. Shape priors for level set representations. *Proceedings of ECCV* 2, 78–92.
- Paragios, N., Rousson, M., Ramesh, V., 2002. Knowledge-based registration and segmentation of the left ventricle: a level set approach, to appear in *IEEE Workshop on Applications in Computer Vision Orlando, FL*.
- Penney, G.P., Weese, J., Little, J.A., Desmedt, P., Hill, D.L.G., Hawkes, D.J., 1998. A comparison of similarity measures for use in 2D–3D medical image registration. *IEEE Trans. Medical Imaging* 17 (4), 586–595.
- Ronfard, R., 1994. Region-based strategies for active contour models. *Int. J. Computer Vision* 13 (2), 229–251.
- Samson, C., Blanc-Feraud, L., Aubert, G., Zerubia, J., 1999. A Level Set Method for Image Classification. In: *Int. Conf. Scale-Space Theories in Computer Vision*, pp. 306–317.
- Staib, L.H., Duncan, J.S., 1992. Boundary finding with parametrically deformable models. *IEEE Trans. PAMI* 14 (11), 1061–1075.
- Szekely, G., Kelemen, A., Brechbuehler, C., Gerig, G., 1996. Segmentation of 3D objects from MRI volume data using constrained elastic deformations of flexible fourier surface models. *Medical Image Analysis* 1 (1), 19–34.
- Tang, T.S.Y., Ellis, R.E., Fichtinger, G., 2000. Fiducial registration from a single X-ray image; a new technique for fluoroscopic guidance and radiotherapy. *Proceedings of MICCAI 1995*, 502–511.
- Tek, H., Kimia, B., 1995. Image segmentation by reaction diffusion bubbles. In: *Proceedings of ICCV*, pp. 156–162.
- Vannier, M., Butterfield, R., Jordan, D., Murphy, W. et al., 1985. Multi-spectral analysis of magnetic resonance images. *Radiology* 154, 221–224.
- Warfield, S.K., Kaus, M., Jolesz, F.A., Kikinis, R., 2000. Adaptive, template moderated, spatially varying statistical classification. *Medical Image Analysis* 4 (1), 43–55.
- Weese, J., Penney, G.P., Desmedt, P., Buzug, T.M., Hill, D.L.G., Hawkes, D.J., 1997a. Voxel-based 2-D/3-D registration of fluoroscopy images and CT scans for image-guided surgery. *IEEE Trans. Information Technology in Biomedicine* 1 (4), 284–293.
- Weese, J., Buzug, T.M., Lorenz, C., Fassnacht, C., 1997b. An approach to 2D/3D registration of a vertebra in 2D X-ray fluoroscopies with 3D CT images. In: *Proceedings of CVRMed-MRCAS*, p. 119.
- Wells, W., Viola, P., Kikinis, R., 1995. Multi-modal volume registration by maximization of mutual information. In: *Proceedings of MRCAS*, pp. 55–62.
- Yaniv, Z.R., 1998. *Fluoroscopic X-ray Image Processing and Registration for Computer-Aided Orthopedic Surgery*, Masters Thesis The Hebrew University of Jerusalem, Israel.
- Yezzi, A., Kichenassamy, S., Kumar, A., Olver, P., Tannenbaum, A., 1997. A geometric snake model for segmentation of medical imagery. *IEEE Trans. Medical Imaging* 16 (2), 199–209.
- Yezzi, A., Tsai, A., Willsky, A., 1999. A statistical approach to image segmentation for bimodal and trimodal imagery. In: *Proceedings of ICCV*, pp. 898–903.
- Yushkevich, P., Fritsch, D., Pizer, S., Chaney, E., 1999. Towards Automatic, Model-Driven Determination of 3D Patient Setup Errors in Conformal Radiotherapy, Technical Report TR99-007, Department of Computer Science, University of North Carolina–Chapel Hill.
- Zhu, S., Yuille, A., 1996. Region competition: unifying snakes, region growing, and bayes/MDL for multiband image segmentation. *IEEE Trans. PAMI* 18 (9), 884–900.