

Joint Segmentation via Patient-Specific Latent Anatomy Model

T. Riklin Raviv¹, B.H. Menze^{1,5}, K. Van-Leemput^{1,2,3}, B. Stieltjes⁶,
M.A. Weber^{6,7}, N. Ayache⁵, W. M. Wells III^{1,4} and P. Golland¹

¹ Computer Science and Artificial Intelligence Laboratory, MIT, USA

² Department of Information and Computer Science, Helsinki University of
Technology, Finland

³ Department of Radiology, MGH, Harvard Medical School, USA

⁴ Brigham and Womens Hospital, Harvard Medical School, USA

⁵ Asclepios Research Project, INRIA Sophia Antipolis, France

⁶ German Cancer Research Center (DKFZ), Heidelberg, Germany

⁷ Diagnostic Radiology, University Hospital, Heidelberg, Germany

Abstract. We present a generative approach for joint 3D segmentation of patient-specific MR scans across different modalities or time points. The latent anatomy, in the form of spatial parameters, is inferred simultaneously with the evolution of the segmentations. The individual segmentation of each scan supports the segmentation of the group by sharing common information. The joint segmentation problem is solved via a statistically driven level-set framework. We illustrate the method on an example application of multimodal and longitudinal brain tumor segmentation, reporting promising segmentation results.

Key words: patient-specific latent anatomy, spatial parameters, tumor segmentation, level-set framework

1 Introduction

Modeling patient-specific anatomy is essential in longitudinal studies and pathology detection. We present a generative approach for joint segmentation of MR scans of a specific subject, where the latent anatomy, in the form of spatial parameters is inferred concurrently with the segmentation. While the methodology can be applied to a variety of applications, here we focus on segmentation of pathological tissues. Specifically, we demonstrate our algorithm on a problem of multimodal segmentation of brain tumors in longitudinal studies. Patient-specific datasets acquired through different modalities at a particular time point are segmented simultaneously, yet individually, based on the specific parameters of their intensity distributions. The spatial parameters that are shared among the scans facilitate the segmentation of the group.

While probabilistic atlases are commonly used as priors for segmentation of MR scans of normal tissues or structures [1, 8, 20–22, 31] the standard methods fail to handle pathologies or to detect subtle anatomical deformation, for example, due to aging. A few methods generate patient-specific atlases by iteratively refining the normal template model [13, 17, 24]. Other methods detect tumors

from differences in images acquired at different time points [26, 29]. Both approaches rely heavily on priors such as tumor shape, intensities, growth and expected evolution [5, 13, 16, 17, 24, 26, 29, 34]. Discriminative approaches [4, 9–11, 25, 32, 33] segment lesions by constructing feature distributions that characterize healthy subjects, so that the pathology can be specified as outliers. However, the variability of normal brain scans and the effects some tumors have on their surrounding healthy tissues lead to a high false positive detection rate. Moreover, mild anomalies can be wrongly classified as normal.

Here we propose and demonstrate a fully automatic groupwise segmentation method. No prior knowledge or external information is required but a couple of mouse clicks at approximately the center and the boundary of a single tumor slice (out of the few dozen volumes to segment) that are used to initialize the segmentations of the images acquired at the first time point. All model parameters, spatial and intensity, are inferred from the patient scans alone. Tumor segmentations at a given time point are used to initialize the segmentations at the next time point for scans of corresponding modalities. The output of the algorithm consist of individual segmentations for each modality and time point. This is in contrast to many discriminative methods, e.g., [32], that use multi-modal datasets for multivariate feature extraction, assuming spatial coherence of the tumor outlines in different image modalities. Here we relax this assumption and search for systematic, structural differences of the visible tumor volume acquired by different imaging protocols.

Our latent anatomy segmentation model is based on probabilistic principles but is solved using partial differential equations (PDEs) and energy minimization criteria. We describe a statistically-driven level-set algorithm that expresses segmentation uncertainty via the logistic function of the associated level-set values, similar to [23]. We relate the image likelihood term to the region based constraint that relaxes the piecewise smoothness assumption of [18], in the spirit of [3, 19, 35]. We also draw the connection between a Markov random field (MRF) prior on the individual segmentations and two continuous-form energy terms: the commonly used smoothness constraint, originally proposed in [12] and the spatial constraint, associated with the latent anatomy parameters. We developed this approach in [27] and validated the algorithm on joint segmentation of cortical and subcortical structures in a population. Here we investigate its application to a patient-specific tumor data set.

The paper is organized as follows. Section 2 defines the problem of latent-anatomy segmentation. In Section 3 we derive our level-set framework for fitting the probabilistic model to the image data. The alternating minimization algorithm is presented in Section 4. Section 5 reports the experimental results followed by a discussion in Section 6.

2 Problem definition and probabilistic model

This section summarizes the formulation of [27] for the joint segmentation of N aligned images. The images can, for example, represent N scans of a specific

patient acquired via different imaging protocols. Our objective is to segment a particular region of interest, a brain lesion for example, that may appear slightly differently across the images. Let $I_n: \Omega \rightarrow \mathbb{R}^+$, be a gray level image with V voxels, defined on $\Omega \subset \mathbb{R}^3$ and let $\Gamma_n: \Omega \rightarrow \{0, 1\}$ be the unknown segmentation of the image I_n , $n = 1, \dots, N$. We assume that each segmentation Γ_n is generated iid from a probability distribution $p(\Gamma; \theta_\Gamma)$ where θ_Γ is the set of the unknown spatial parameters. We also assume that Γ_n generates the observed image I_n , independently of all other image-segmentation pairs, with probability $p(I_n | \Gamma_n; \theta_{I,n})$ where $\theta_{I,n}$ are the parameters corresponding to image I_n . Since the images are acquired by different imaging protocols we assign a different set of intensity parameters to each of them.

Let $\{I_1 \dots I_N\}$ be the given set of aligned images that form the observed variable in our problem and let $\Gamma = \{\Gamma_1, \dots, \Gamma_N\}$ be the corresponding unknown segmentations. The joint distribution $p(I_1 \dots I_N, \Gamma_1 \dots \Gamma_N; \Theta)$ is governed by the composite set of parameters $\Theta = \{\theta_\Gamma, \theta_{I,1} \dots \theta_{I,N}\}$. Our goal is to estimate the segmentations Γ . This, however, cannot be accomplished in a straightforward manner since the model parameters are also unknown. We therefore jointly optimize Γ and Θ :

$$\{\hat{\Theta}, \hat{\Gamma}\} = \arg \max_{\{\Theta, \Gamma\}} \log p(I_1 \dots I_N, \Gamma_1 \dots \Gamma_N; \Theta) \quad (1)$$

$$= \arg \max_{\{\Theta, \Gamma\}} \sum_{n=1}^N [\log p(I_n | \Gamma_n; \theta_{I,n}) + \log p(\Gamma_n; \theta_\Gamma)]. \quad (2)$$

We alternate between estimating the maximum a posteriori (MAP) segmentations and updating the model parameters. For a given setting of the model parameters $\hat{\Theta}$, Eq. (2) implies that the segmentations can be estimated by solving N separate MAP problems:

$$\hat{\Gamma}_n = \arg \max_{\Gamma_n} [\log p(I_n | \Gamma_n; \theta_{I,n}) + \log p(\Gamma_n; \theta_\Gamma)]. \quad (3)$$

We then fix $\hat{\Gamma}$ and estimate the model parameters $\Theta = \{\theta_\Gamma, \theta_{I,1}, \dots, \theta_{I,N}\}$ by solving two ML problems:

$$\hat{\theta}_{I,n} = \arg \max_{\theta_{I,n}} \log p(I_n; \Gamma_n, \theta_{I,n}), \quad (4)$$

$$\hat{\theta}_\Gamma = \arg \max_{\theta_\Gamma} \sum_{n=1}^N \log p(\Gamma_n; \theta_\Gamma). \quad (5)$$

3 Probabilistic view of the level-set framework

Now we draw the connection between the probabilistic model presented above and a level-set framework for segmentation. Let $\phi_n: \Omega \rightarrow \mathbb{R}$ be the level-set function associated with image I_n . The zero level $C_n = \{\mathbf{x} \in \Omega | \phi_n(\mathbf{x}) = 0\}$

defines the interface that partitions the image space of I_n into two disjoint regions ω and $\Omega \setminus \omega$. Similar to [20, 23] we define the level-set function ϕ_n using the log-odds formulation instead of the conventional signed distance function:

$$\phi_n(\mathbf{x}) \triangleq \epsilon \operatorname{logit}(p) = \epsilon \log \frac{p(\mathbf{x} \in \omega)}{1 - p(\mathbf{x} \in \omega)} = \epsilon \log \frac{p(\mathbf{x} \in \omega)}{p(\mathbf{x} \in \Omega \setminus \omega)}, \quad (6)$$

where $p(\mathbf{x} \in \omega)$ can be viewed as the probability that the voxel in location \mathbf{x} belongs to the foreground region. The constant ϵ determines the scaling of the level-set function ϕ_n with respect to the ratio of the probabilities. The inverse of the logit function for $\epsilon = 1$ is the logistic function:

$$H_\epsilon(\phi_n) = \frac{1}{2} \left(1 + \tanh \left(\frac{\phi_n}{2\epsilon} \right) \right) = \frac{1}{1 + e^{-\phi_n/\epsilon}}. \quad (7)$$

Note, that $H_\epsilon(\phi_n)$ is similar, though not identical, to the regularized Heaviside function introduced by Chan and Vese [3]. We use this form of Heaviside function and its derivative with respect to ϕ in the proposed level-set formulation. To simplify the notation, we omit the subscript ϵ in the rest of the paper.

3.1 Cost functional for segmentation

The joint estimation problem of the hidden variables Γ , or ϕ_n (using the level-set notation) and the unknown model parameters Θ can be solved as an energy minimization problem, where

$$E(\phi_n) = -\log p(I_n | \Gamma_n; \theta_{I,n}) - \log p(\Gamma_n; \theta_\Gamma).$$

As in [27], we establish the correspondence between the log probability and the level-set energy terms. Let $E_I(\phi_n, \Theta)$ denote the term corresponding to the image likelihood in Eq. (3). Then

$$E_I(\phi_n, \Theta) = - \int_{\Omega} [\log p_{\text{in}}(I_n; \theta_{I,n}) H(\phi_n(\mathbf{x})) + \log p_{\text{out}}(I_n; \theta_{I,n}) (1 - H(\phi_n(\mathbf{x})))] d\mathbf{x}, \quad (8)$$

where, p_{in} and p_{out} denote the probability distributions of the foreground and background intensities of a particular image I_n , respectively. If we use, for example, Gaussian densities for p_{in} and p_{out} we get the familiar minimal variance term [3, 19, 35]. Here, we use a Gaussian mixture to model the background, as described later in the paper.

Let us now consider the prior probability $p(\Gamma_n; \theta_\Gamma)$ in Eq. (2) and its corresponding energy terms. Specifically, we construct an MRF prior for segmentations:

$$\begin{aligned} \log p(\Gamma_n; \theta_\Gamma) &= \sum_{v=1}^V [\Gamma_n^v \log(\theta_\Gamma^v) + (1 - \Gamma_n^v) \log(1 - \theta_\Gamma^v)] \\ &\quad - \sum_{v=1}^V f(\Gamma_n^v, \Gamma_n^{\mathcal{N}(v)}) - \log Z(\theta_\Gamma), \end{aligned} \quad (9)$$

where $Z(\theta_\Gamma)$ is the partition function and $\mathcal{N}(v)$ is the set of the closest neighbors of voxel v . The function

$$f(\Gamma_n^v, \Gamma_n^{\mathcal{N}(v)}) = \sum_{v' \in \mathcal{N}(v)} w_{(v,v')} (\Gamma_n^v - \Gamma_n^{v'})^2,$$

accounts for the interactions between neighboring voxels. It can be configured, by setting the values of $w_{(v,v')}$, to act as a finite difference operator approximating the gradient of Γ_n at the voxel v [15]. This approximation allows us to represent the discrete term $\sum_{v=1}^V f(\Gamma_n^v, \Gamma_n^{\mathcal{N}(v)})$ as an approximation of the continuous term

$$E_{\text{LEN}}(\phi_n) = \int_{\Omega} |\nabla H(\phi_n(\mathbf{x}))| d\mathbf{x}, \quad (10)$$

which is the commonly used length term. Note that if we omit the pairwise term in Eq. (9), the prior on segmentations $p(\Gamma_n | \theta_\Gamma)$ reduces to a Bernoulli distribution, where the parameters θ_Γ represent the probability map for the structure of interest. The introduction of the pairwise clique potentials complicates the model but encourages smoother labeling configurations.

We define the spatial energy term E_S based on the singleton term in Eq. (9). Using the level-set formulation we obtain:

$$E_S(\phi_n, \Theta) = - \int_{\Omega} [\log \theta_\Gamma(\mathbf{x}) H(\phi_n(\mathbf{x})) + \log(1 - \theta_\Gamma(\mathbf{x})) (1 - H(\phi_n(\mathbf{x})))] d\mathbf{x}. \quad (11)$$

Note, that ignoring the partition function in the equations that follow Eq. (9) has no effect on the estimation of Eq. (3), but it changes Eq. (5) to be maximum pseudo likelihood [2], rather than maximum likelihood.

We construct the cost functional for $\phi_1 \dots \phi_N$ and the parameters Θ by combining Eq. (8), (10) and (11):

$$E(\phi_1 \dots \phi_N, \Theta) = \gamma E_{\text{LEN}} + \beta E_I + \alpha E_S \quad (12)$$

where $\alpha = 1 - \beta - \gamma$. As in [28], we adaptively tune the weights such that the contributions of the energy terms E_{LEN} , E_I and E_S to the overall cost are balanced.

4 Gradient descent and parameter estimation

We optimize Eq. (12) by a set of alternating steps. For fixed model parameters Θ , the evolution of each level-set function ϕ_n is determined by the following gradient descent equation:

$$\phi_n(\mathbf{x}, t + \Delta t) = \phi_n(\mathbf{x}, t) + \frac{\partial \phi_n}{\partial t} \Delta t, \quad (13)$$

where $\frac{\partial \phi_n}{\partial t}$ is obtained from the first variation of $E(\phi_n, \Theta)$. Using the Euler-Lagrange equations we get:

$$\begin{aligned} \frac{\partial \phi_n}{\partial t} = \delta(\phi_n) \left\{ \gamma \operatorname{div} \left(\frac{\nabla \phi_n}{|\nabla \phi_n|} \right) + \beta [\log p_{\text{in}}(I_n(\mathbf{x}); \theta_{I,n}) - \log p_{\text{out}}(I_n(\mathbf{x}); \theta_{I,n})] \right. \\ \left. + \alpha [\log \theta_\Gamma - \log(1 - \theta_\Gamma)] \right\}, \end{aligned} \quad (14)$$

where $\delta(\phi_n)$ is the derivative of $H(\phi_n)$ with respect to ϕ_n :

$$\delta_\epsilon(\phi_n) = \frac{1}{2\epsilon} \operatorname{sech}\left(\frac{\phi_n}{2\epsilon}\right) = \frac{1}{\epsilon \cosh\left(\frac{\phi_n}{\epsilon}\right)}.$$

For fixed segmentations ϕ_n , the model parameters are recovered by differentiating the cost functional in Eq. (12) with respect to each parameter.

4.1 Intensity parameters

We assume that the intensities of the structure of interest are drawn from a normal distribution, i.e., $p_{\text{in}}(I_n; \theta_{I,n}) = \mathcal{N}(I_n; \mu_n, \sigma_n^2)$. The intensities of the background tissues are modeled as a K-Gaussian mixture:

$$p_{\text{out}}(I_n; \theta_{I,n}) = \text{GMM}(\mu_n^1 \cdots \mu_n^K, \sigma_n^1 \cdots \sigma_n^K, \lambda_n^1 \cdots \lambda_n^K),$$

where λ_n^k is the mixing proportion of component k in the mixture. We estimate the Gaussian mixture model parameters using the expectation maximization (EM) method [6].

4.2 Spatial parameters

We estimate the spatial function $\theta_\Gamma(\mathbf{x})$, which represents a dynamically evolving latent atlas, by optimizing the sum of the energy terms that depend on θ_Γ :

$$\hat{\theta}_\Gamma = \arg \max_{\theta_\Gamma} \sum_{n=1}^N \int_{\Omega} [\tilde{H}(\phi_n(\mathbf{x})) \log(\theta_\Gamma(\mathbf{x})) + (1 - \tilde{H}(\phi_n(\mathbf{x}))) \log(1 - \theta_\Gamma(\mathbf{x}))] d\mathbf{x},$$

yielding

$$\hat{\theta}_\Gamma(\mathbf{x}) = \frac{1}{N} \sum_{n=1}^N \tilde{H}(\phi_n(\mathbf{x})). \quad (15)$$

4.3 Algorithm

We summarize the proposed latent-anatomy segmentation algorithm assuming the following setup. The input consist of N aligned volumes $\{I_{\tau,m}\}$, where $I_{\tau,m}$ is a volume acquired at time τ and modality m .

Initialization The user selects one of the volumes acquired at the first time point and identifies a single sagittal, axial or coronal slice where the tumor or the structure of interest is clearly seen. The user marks with a couple of mouse clicks the approximate location of the tumor center and one of its boundary points. This input determines a sphere that is used to initialize the segmentations of all the volumes acquired at the first time point. We denote their number by M_1 . In our implementation M_1 identical level-set functions that are defined by the signed distance function of this sphere are used for initialization.




Fig. 1. Axial slice of the tumor volumes and the automatic 3D segmentations (red outlines) across 6 modalities and 10 time points. Not all the modalities were acquired at each time point.

1. Set $\tau = 1$. Note that τ is the index of the actual time in which the scans were acquired. It should not be confused with t in Eq. (13) that denotes the number of gradient descent iterations.
2. Calculate the background and foreground intensity parameters, for each of the volumes acquired at time τ based on the current estimates of their corresponding level-set functions ϕ_n , according to subsection 4.1.
3. Calculate the latent anatomy parameters θ_I based on the current estimate of the level-set functions ϕ_n , corresponding to the image volumes acquired at time τ (Eq. (15)).
4. Use Eq. (14) (gradient descent) to evolve the level-set functions associated with the volumes acquired at time τ based on the current estimates of the respective intensity parameters θ_I and the spatial parameters θ_G .
5. Repeat steps 2-4 until convergence.
6. Use the final state of the level-set functions ϕ_n associated with the volumes acquired by modality m at time τ to initialize the corresponding level-set associated with the volumes acquired at time $\tau + 1$. Set $\tau := \tau + 1$
7. Repeat steps 2-6 sequentially, for all the time points.

5 Experimental Results

We applied the proposed method to a set of 44 image volumes of a patient with histologically confirmed low-grade glioma, acquired at 10 different time points at the German Cancer Research Center (Heidelberg, Germany) using 1.5T Siemens Magnetom and 3T Siemens TRIO MR scanners. The volumes were acquired via six imaging protocols: T1, T2, FLAIR, DTI, and contrast-enhanced T1 sequences (T1gd). We note that not all acquisition modalities were used at each time point, as illustrated in Fig. 1. We aligned the images using the MedINRIA registration software [30]. We calculated fractional anisotropy (FA) and apparent diffusion coefficient (ADC) maps from the diffusion tensor images (DTI) using the same software. To enable quantitative evaluation, three manual segmentations of three orthogonal slices that pass through the center of the tumor were provided for each volume.

Fig. 1 presents axial slices of the available image volumes together with the boundaries of the automatic 3D segmentation. Fig. 2 shows the manual segmentations for three lateral slices through the tumor, together with the automatic segmentation.

Table 1 provides quantitative evaluation of the overlap between the automatic and the manual segmentations as measured by the Dice coefficients [7]. We compared the automatic segmentations with the corresponding triplets of manually segmented slices. The first number in each cell in Table 1 reports the mean and the standard deviation of these nine Dice scores. The second number in each cell in Table 1 reports the mean and the standard deviation of the Dice scores obtained by comparing one of the manual segmentations with the other two in the three slices. We marked with asterisks cells that show similar Dice scores. The overall average Dice score for the automatic segmentation is

Time of acquisition	T1	T1gd	T2	Flair	DA	ADC
Nov. 2005	.71±.11*	.49±.14*	.87±.01	.91±.03	.84±.02	.73±.08
	.80±.07	.52±.15	.95±.01	.94±.02	.94±.01	.91±.07
March 2006	.78±.04	.77±.09*	.93±.02*	.84±.02		
	.92±.02	.85±.06	.95±.02	.96±.01		
Sep. 2006	.87±.03*	.85±.08*	.84±.02	.94±.01*	.88±.03	.81±.04
	.90±.04	.82±.12	.95±.02	.94±.02	.94±.02	.93±.03
Dec. 2006	.87±.02	.89±.03*	.93±.01	.86±.02		
	.91±.03	.90±.04	.95±.01	.94±.01		
March 2007	.84±.02	.82±.11*	.93±.02*	.84±.02		
	.91±.04	.87±.11	.94±.03	.94±.01		
June 2007	.81±.09*	.85±.09*	.93±.02*			
	.86±.11	.81±.11	.93±.03			
Sep. 2007	.87±.04	.84±.08*	.87±.02		.86±.03	.87±.02
	.92±.03	.86±.07	.94±.02		.93±.01	.94±.02
Jan. 2008	.88±.02*	.89±.02*	.91±.03*	.83±.04		
	.90±.01	.90±.02	.91±.01	.94±.02		
July 2008	.87±.03	.85±.02	.91±.03	.86±.03		
	.93±.02	.93±.02	.94±.01	.94±.02		
Oct. 2008	.88±.03	.85±.03	.91±.03*	.84±.04		
	.93±.03	.93±.01	.94±.01	.93±.02		

Table 1. Dice coefficients for 44 volumes in the study. The first number in each cell reports the mean and the standard deviation of the Dice scores of the automatic segmentation with respect to three manual segmentations. The second number in each cell reports the mean and the standard of the Dice scores calculated between one of the manual segmentations and the average of the other two (see text for detail). Automatic segmentations that did not differ significantly from the manual ones are marked by the asterisk.

above 0.85 while the average Dice scores obtained for the manual segmentations is 0.91. The top plot in Fig. 3 presents the average Dice score over all modalities at a given time point obtained by our method (red) and the Dice scores of the multivariate tissue classification [32] (green). The plot shows that the Dices scores obtained via the latent anatomy method are consistently higher. The bottom plot in Fig. 3 compares the overlap among the manual segmentations for each individual modality ('intra-modal') with the overlap among the manual segmentations for all the modalities together ('inter-modal'). We define overlap as the mean Dice score among the three manual segmentations, as described above. This plot suggests that even the manual segmentations vary significantly across modalities for the same time point, justifying our approach of generating separate tumor segmentation for each volume.




Fig. 2. Manual segmentations (red, green, blue) and automatic segmentation (black) for lateral T1, T2 and DTI-FA images acquired at the same time point. The fourth image shows the corresponding section of the average of the associated 3D level-set functions. The black line indicates the zero level. Gray dashed lines indicate the tumor boundaries of all the modalities available for that time point.




Fig. 3. Top: A comparison of the average Dice scores of the proposed latent anatomy method (red) and the Dice scores of the multivariate EM for lesion segmentation of [32] (green). Note that the segmentation results obtained by the proposed latent anatomy method are consistently better. Bottom: comparison of the correspondence between the manual segmentations for each individual modality (red) with the correspondence between the manual segmentations for all the modalities together (green). Correspondence, here, is defined as the mean Dice score between the three manual segmentations (see text for detail).

6 Discussion and future directions

We presented a statistically driven level-set approach for joint segmentation of subject-specific MR scans. The latent patient anatomy, which is represented by a set of spatial parameters is inferred from the data simultaneously with the segmentation through an alternating minimization procedure. Segmentation of each of the channels, or modalities, is therefore supported by the common information shared by the group. The method is demonstrated by addressing the problem of multi-modal brain tumor segmentation across 5 – 10 time points. Promising segmentation results were obtained. An on-going research is now conducted to construct a functional model of the tumor growth based on the image sequences [14] that will be used to predict the evolution of the tumor outlines.

Acknowledgments. This work was supported in part by the Leopoldina Fellowship Programme (LPDS 2009-10), NIH NIBIB NAMD U54-EB005149, NIH NCRN P41-RR13218, NIH NINDS R01-NS051826 grants and NSF CAREER Award 0642971.

References

1. J. Ashburner and K. Friston. Unified segmentation. *NeuroImage*, 26:839–851, 2005.
2. J. Besag. Statistical analysis of non-lattice data. *The Statistician*, 24(3):179–195, 1975.
3. T.F. Chan and L.A. Vese. Active contours without edges. *IEEE TIP*, 10(2):266–277, 2001.
4. Cobzas *et al.* 3D variational brain tumor segmentation using a high dimensional feature set. In *ICCV*, 1–8, 2007.
5. Cuadra *et al.* Atlas-based segmentation of pathological brain MR images. *IEEE TMI*, 23:1301–1314, 2004.
6. A. Dempster, N. Laird, and D. Rubin. Maximal likelihood from incomplete data via the EM algorithm. *Proceedings of the Royal Statistical Society*, 39:1 – 38, 1977.
7. L. Dice. Measure of the amount of ecological association between species. *Ecology*, 26(3):297–302, 1945.
8. Fischl *et al.* Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33(3):341–355, 2002.
9. D. Gering, W.E.L. Grimson, and R. Kikinis. Recognizing deviations from normalcy for brain tumor segmentation. In *MICCAI*, vol. 2488, 388 – 395, 2002.
10. Grlitz *et al.* Semi-supervised tumor detection in magnetic resonance spectroscopic images using discriminative random fields, 2007.
11. S. Ho, E. Bullitt, and G. Gerig. Level-set evolution with region competition: automatic 3-D segmentation of brain tumors. In *ICPR*, vol. 1, p532–535, 2002.
12. M. Kass, A.P. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, Jan. 1988.
13. Kaus *et al.* Automated segmentation of MR images of brain tumors. *Radiology*, 218:586 – 591, 2001.
14. Konukoglu *et al.* A recursive anisotropic fast marching approach to reaction diffusion equation: Application to tumor growth modeling. In *IPMI*, p687–699, 2007.
15. S.Z. Li. *Markov Random Field Modeling in Computer Vision*. Springer-Verlag, 1995.
16. Mohamed *et al.* Deformable registration of brain tumor images via a statistical model of tumor-induced deformation. *Medical Image Analysis*, 10:752–763, 2006.

17. Moon *et al.* Model-based brain and tumor segmentation. In *ICPR*, vol. 1, 528–531, 2002.
18. D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, 42:577–684, 1989.
19. N. Paragios and R. Deriche. Geodesic active regions: A new paradigm to deal with frame partition problems in computer vision. *JVCIR*, 13:249–268, 2002.
20. K.M. Pohl and R. Kikinis and W.M. Wells. Active Mean Fields: Solving the Mean Field Approximation in the Level Set Framework. *IPMI*, 4584:26–37, 2007.
21. Pohl *et al.* A bayesian model for joint segmentation and registration. *NeuroImage*, 31(1):228 – 239, 2006.
22. Pohl *et al.* A hierarchical algorithm for MR brain image parcellation. *TMI*, 26(9):1201–1212, 2007.
23. Pohl *et al.* Using the logarithm of odds to define a vector space on probabilistic atlases. *Medical Image Analysis*, 11(6):465–477, 2007.
24. Prastawa *et al.* Automatic brain tumor segmentation by subject specific modification of atlas priors. *Academic Radiology*, 10:1341–1348, 2003.
25. Prastawa *et al.* A brain tumor segmentation framework based on outlier detection. *Medical Image Analysis*, 8:275–283, 2004.
26. Rey *et al.* Automatic detection and segmentation of evolving processes in 3D medical images: Application to multiple sclerosis. *Medical Image Analysis*, 6(2):163–179, 2002.
27. Riklin Raviv *et al.* Joint segmentation of image ensembles via latent atlases. In *MICCAI*, 2009. accepted.
28. T. Riklin-Raviv, N. Sochen, and N. Kiryati. Shape-based mutual segmentation. *International Journal of Computer Vision*, 79:231–245, 2008.
29. J.P. Thirion and G. Calmon. Deformation analysis to detect and quantify active lesions in three-dimensional medical image sequences. *IEEE TMI*, 18(5):429–441, 1999.
30. N. Toussaint, J.C. Souplet, and P. Fillard. Medinria: Medical image navigation and research tool by inria. In *Proc. of MICCAI Workshop on Interaction in medical image analysis and visualization*, 2007.
31. Van Leemput *et al.* Automated model-based tissue classification of MR images of the brain. *IEEE TMI*, 18(10):897–908, 1999.
32. Van Leemput *et al.* Automated segmentation of multiple sclerosis lesions by model outlier detection. *IEEE TMI*, 20:677–688, 2001.
33. Wels *et al.* A discriminative model-constrained graph cuts approach to fully automated pediatric brain tumor segmentation in 3-D MRI. In *MICCAI*, vol. 5241, 67 – 75, 2008.
34. Zacharaki *et al.* Orbit: A multiresolution framework for deformable registration of brain tumor images. *IEEE TMI*, 27(8):1003–1017, 2008.
35. S.C. Zhu and A.L. Yuille, Region Competition: Unifying Snakes, Region Growing, and Bayes/MDL for Multiband Image Segmentation. *PAMI*, 18(9):884–900, 1996.