

Joint Segmentation via Patient-Specific Latent Anatomy Model

T. Riklin Raviv¹, B. 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 Helsinki University of Technology, Finland

³ Department of Neurology, 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 or for pathology detection. We present a generative approach to 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. The latent anatomy determines a Markov Random Field (MRF) prior on the segmentations. 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, 17, 18, 27] 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 [12, 15, 20]. Other methods detect tumors

from differences in images acquired at different time points [22, 25]. Both approaches heavily rely on priors such as tumor shape, intensities, growth and expected evolution [5, 12, 14, 15, 20, 22, 25, 30].

Discriminative approaches [4, 9–11, 21, 28, 29] segment lesions by constructing feature distributions that characterize healthy subjects, so that the pathologies 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. While many discriminative methods use multimodal datasets for multivariate feature extraction to achieve a better discrimination between tissue types, e.g., [28], they may not be applicable in a clinical setting. For example, longitudinal studies that are conducted to explore anatomical deformation due to aging or to model the growth and appearance of a pathology frequently focus on specific modalities.

Here we propose and demonstrate a method that is almost ‘free’ of prior information. All model parameters, spatial and intensity related, are inferred from the individual patient scans alone. The segmentations for the first time point are initialize as a sphere centered approximately at the tumor center. The sphere’s location and diameter can be provided by the user with a couple of mouse clicks. Tumor segmentations at a given time point are used to initialize the segmentation 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 clinically valueable since each imaging protocol exposes different characteristics of the tumor, which allows finer analysis of growth and shape deformation. Our latent-anatomy segmentation model is based on probabilistic principles but is solved using partial differential equations (PDEs) and energy minimization criteria. We derive a statistically-driven level set algorithm that expresses segmentation uncertainty via the logistic function of the associated level set values, similar to [19]. We developed this approach in [23] 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 group-wise 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 [23] 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 $\Gamma_n: \Omega \rightarrow \{0,1\}$ be the unknown segmentation of the image I_n . We assume that each segmentation Γ_n is generated iid from a probability distribution $p(\Gamma | \theta_\Gamma)$ where θ_Γ is a set of unknown 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 . We assign a specific set of intensity parameters to each image since it could be acquired under different conditions.

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 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 value of the model parameters $\hat{\Theta}$, Equation (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:

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

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

2.1 Image likelihood term

We use by p_{in} and p_{out} to denote the probability distributions of the foreground and background intensities of a particular image I_n , respectively. The image likelihood term in Equation (2) then becomes:

$$\log p(I_n | \Gamma_n, \hat{\theta}_{I,n}) = \sum_{\{v | \Gamma_n^v = 1\}} \log p_{\text{in}}(I_n^v; \theta_{I,n}) + \sum_{\{v | \Gamma_n^v = 0\}} \log p_{\text{out}}(I_n^v; \theta_{I,n}). \quad (6)$$

2.2 Spatial prior term

We now consider the prior probability $p(\Gamma_n | \theta_\Gamma)$ in Equation (2). Specifically, we construct an MRF prior for segmentations:

$$p(\Gamma_n | \theta_\Gamma) = \frac{1}{Z(\theta_\Gamma)} \prod_{v=1}^V (\theta_\Gamma^v)^{\Gamma_n^v} (1 - \theta_\Gamma^v)^{(1 - \Gamma_n^v)} e^{-f(\Gamma_n^v, \Gamma_n^{\mathcal{N}(v)})}, \quad (7)$$

where $Z(\theta_\Gamma)$ is the partition function and $\mathcal{N}(v)$ are the closest neighbors of voxel v . The function $f(\cdot)$ accounts for the interactions between neighboring voxels. If we omit the pairwise term in Equation (7), 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 will hereafter omit the partition function, using only an approximation to the distributions of the random variables Γ_n [2]. The logarithm of the expression above takes the following form:

$$\begin{aligned} \log p(\Gamma_n | \theta_\Gamma) &= \sum_{v=1}^V [\Gamma_n^v \log(\theta_\Gamma^v) + (1 - \Gamma_n^v) \log(1 - \theta_\Gamma^v)] \\ &\quad - \sum_{v=1}^V f(\Gamma_n^v, \Gamma_n^{\mathcal{N}(v)}). \end{aligned} \quad (8)$$

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 [19] we define the level set function ϕ_n using the log-odds formulation instead of the conventional signed distance function:

$$\phi_n(\mathbf{x}) \triangleq \epsilon \logit(p) = \epsilon \log \frac{p(\mathbf{x} \in \omega)}{1 - p(\mathbf{x} \in \omega)} = \epsilon \log \frac{p(\mathbf{x} \in \omega)}{p(\mathbf{x} \in \Omega \setminus \omega)}, \quad (9)$$

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:

$$H(\phi_n) = \frac{1}{2} \left(1 + \tanh \left(\frac{\phi_n}{2\epsilon} \right) \right) = \frac{1}{1 + e^{-\phi_n/\epsilon}}. \quad (10)$$

Note, that $H(\phi_n)$ is similar, though not identical, to the regularized Heaviside function introduced by Chan and Vese [3]. The logistic function $H(\phi_n)$ is used in the proposed level-set formulation as the Heaviside function in [3]. 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

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

for some positive scalars α and β .

As in [23] we show the correspondence between the log of the probability distributions and the level-set energy terms. Let $E_I(\phi_n, \Theta)$ denote the term corresponding to the image likelihood in Eq. (3). Then,

$$E_I(\phi_n, \Theta) = -\beta \int_{\Omega} [\log p_{\text{in}}(I_n; \theta_{I,n}) H(\phi_n(\mathbf{x})) + \log p_{\text{out}}(I_n; \theta_{I,n}) (1 - H(\phi_n(\mathbf{x})))] d\mathbf{x}. \quad (11)$$

If we use, for example, Gaussian densities for p_{in} and p_{out} we get the familiar minimal variance term [3, 16]. 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 Equation (2) and its corresponding energy terms. Specifically, we show in detail, extending the derivation in [23], the equivalence of the length term E_{LEN} and the pairwise interaction term in Equation (8). We define $f(\Gamma_n^v, \Gamma_n^{\mathcal{N}(v)}) = c\sqrt{k(\Gamma_n^v, \Gamma_n^{\mathcal{N}(v)})}$ where c is some positive scalar and

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

The function $f(\cdot)$ 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 [13]. In 1D, for example, denoting v^l, v^r as the left and the right neighbors of v respectively, f is a forward difference if $w_{(v,v^l)} = 0$ and $w_{(v,v^r)} = -1$. The extension to higher dimensions is straightforward, although a slight modification of the function f is needed to account for derivatives based on central differences. The finite difference approximation together with the approximations above allow 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:

$$E_{\text{LEN}}(\phi_n) = \gamma \int_{\Omega} |\nabla H(\phi_n(\mathbf{x}))| d\mathbf{x}, \quad (12)$$

where $\gamma = c\beta$.

We define the spatial energy term E_S based on the singleton term of the MRF model presented in Equation (8). Using the level-set formulation we get:

$$E_S(\phi_n, \Theta) = - \int_{\Omega} [\log \theta_I(\mathbf{x}) H(\phi_n(\mathbf{x})) + \log(1 - \theta_I(\mathbf{x})) (1 - H(\phi_n(\mathbf{x})))] d\mathbf{x}. \quad (13)$$

The cost functional for $\phi_1 \dots \phi_N$ and the parameters Θ is constructed by combining Equations (11), (13) and (12):

$$E(\phi_1 \dots \phi_N, \Theta) = \gamma E_{\text{LEN}} + \beta E_I + \alpha E_S \quad (14)$$

where $\alpha = 1 - \beta - \gamma$. As in [24] we tune the weights such that the contributions of the energy terms E_{LEN} , E_I and E_S to the overall cost is balanced.

4 Gradient descent and parameter estimation

We now describe an algorithm for optimizing Equation (14) in 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:

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

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:

$$\begin{aligned} \frac{\partial \phi_n}{\partial t} = \delta(\phi_n) \left\{ \gamma \operatorname{div} \left(\frac{\nabla \phi_n}{|\nabla \phi_n|} \right) + \beta [\log p_{\text{in}}(I_n(\mathbf{x}); \theta_{I,n}) - \log p_{\text{out}}(I_n(\mathbf{x}); \theta_{I,n})] \right. \\ \left. + \alpha [\log \theta_I - \log(1 - \theta_I)] \right\}, \end{aligned} \quad (15)$$

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 (14) 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:

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

where w_n^k is the mixing proportion of component k in the mixture. The Gaussian mixture model parameters are estimated using expectation maximization (EM) method [6].

4.2 Spatial parameters

We estimate the spatial function $\theta_\Gamma(\mathbf{x})$, constructing 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

$$\hat{\theta}_\Gamma(\mathbf{x}) = \frac{1}{N} \sum_{n=1}^N \tilde{H}(\phi_n(\mathbf{x})). \quad (16)$$

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 affine registration using the MedINRIA software [26]. We calculated fractional anisotropy (FA) and apparent diffusion coefficient (ADC) maps from the diffusion tensor images (DTI) using the same software.

Fig. 1 presents axial slices of the available image volumes together with the boundaries of the automatic 3D segmentation. To evaluate the accuracy of the method we compared the automatic segmentations with the manual segmentations for three slices using the Dice scores [7]. The rater created three separate manual segmentations based on three visualizations with differently adjusted image contrast. 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 similarity between the automatic and the manual segmentations as measured by the Dice coefficients. We compared the automatic segmentation results to the three manual segmentations in each of the three slices. The first number in each cell in Table 1 reports the mean and the standard deviation of these 9 Dice scores. We also compared one of the manual segmentations with the other two in each of the three slices. The second number in each cell in Table 1 reports the mean and the standard deviation of these six measurements. We marked with asterisks cells that shows similar Dice scores, as measured by a t-test at 0.05 significance level. The overall average Dice scores for the automatic segmentation is above 0.85 while the average Dice scores for one of the manual segmentations with respect to the other two is 0.91.




Fig. 1. Axial slice of the tumor volumes and of the automatic 3D segmentation across 6 modalities and 10 time points. Automatic segmentations of the tumor are shown in red. The corresponding Dice scores are shown in Table 1. Note that not all the modalities were acquired at each time point.

We compared the segmentation results obtained by the proposed method with segmentation results of the multivariate EM algorithm for lesion segmentation [28]. The multivariate EM method detects lesions as outliers of healthy brain tissues. The top plot in Fig. 3 presents the average Dice scores over all the modalities at a given time point obtained by our method (red) and the Dice scores of the multivariate tissue classification of [28] (green). The comparison shows that the segmentations obtained via the latent anatomy method are consistently better. The bottom plot in Fig. 3 compares the correspondence between the manual segmentations for each individual modality with the correspondence between the manual segmentations for all the modalities together. Correspondence, here, is defined as the mean Dice score between the three manual segmentations, as described above. This plot suggests that even the manual segmentations vary significantly across modalities for the same time point, providing justification for our approach of generating separate segmentations for each image.

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 corresponding to the images shown in Fig. 1. 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 different 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 which did not differ significantly from the manual ones are indicated by the asterisk.




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 forth image show the corresponding section of a 3D level-set function related to one of the scans. The black line indicates the zero level. Gray dashed lines indicate the tumor boundaries of the other 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 [28] (green). Note that the segmentation results obtained by the proposed latent anatomy method are consistently better. Bottom: comparison of the correspondence among the manual segmentations for each individual modality (red) with the correspondence among the manual segmentations for all the modalities considered 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 extend the applicability of the proposed methodology for tracking mild deformation of cortical and subcortical structures due to illness or aging.

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. Grilz *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. Kaus *et al.* Automated segmentation of MR images of brain tumors. *Radiology*, 218:586 – 591, 2001.
13. S.Z. Li. *Markov Random Field Modeling in Computer Vision*. Springer-Verlag, 1995.
14. Mohamed *et al.* Deformable registration of brain tumor images via a statistical model of tumor-induced deformation. *Medical Image Analysis*, 10:752–763, 2006.
15. Moon *et al.* Model-based brain and tumor segmentation. In *ICPR*, vol. 1, 528–531, 2002.
16. 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.
17. Pohl *et al.* A bayesian model for joint segmentation and registration. *NeuroImage*, 31(1):228 – 239, 2006.

18. Pohl *et al.* A hierarchical algorithm for MR brain image parcellation. *TMI*, 26(9):1201–1212, 2007.
19. Pohl *et al.* Using the logarithm of odds to define a vector space on probabilistic atlases. *Medical Image Analysis*, 11(6):465–477, 2007.
20. Prastawa *et al.* Automatic brain tumor segmentation by subject specific modification of atlas priors. *Academic Radiology*, 10:1341–1348, 2003.
21. Prastawa *et al.* A brain tumor segmentation framework based on outlier detection. *Medical Image Analysis*, 8:275–283, 2004.
22. 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.
23. Riklin Raviv *et al.* Joint segmentation of image ensembles via latent atlases. In *MICCAI*, 2009. accepted.
24. T. Riklin-Raviv, N. Sochen, and N. Kiryati. Shape-based mutual segmentation. *International Journal of Computer Vision*, 79:231–245, 2008.
25. 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.
26. 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.
27. Van Leemput *et al.* Automated model-based tissue classification of MR images of the brain. *IEEE TMI*, 18(10):897–908, 1999.
28. Van Leemput *et al.* Automated segmentation of multiple sclerosis lesions by model outlier detection. *IEEE TMI*, 20:677–688, 2001.
29. 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.
30. E. I. Zacharaki, D. Shen, S. K. Lee, and C. Davatzikos. Orbit: A multiresolution framework for deformable registration of brain tumor images. *IEEE TMI*, 27(8):1003–1017, 2008.