

Multi-modal Brain Tumor Segmentation via Latent Atlases

Tammy Riklin Raviv^{1,2,3}, Koen Van Leemput^{4,5}, and Bjoern H. Menze^{6,7}

¹ Psychiatry and NeuroImaging Laboratory, Harvard Medical School, Brigham and Women’s Hospital, Boston, USA

² Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, USA

³ The Broad Institute of MIT and Harvard, Cambridge, USA

⁴ Harvard Medical School, Mass General Hospital, Boston, USA

⁵ Technical University of Denmark, Lyngby, Denmark

⁶ Computer Vision Laboratory, ETH Zurich, Switzerland

⁷ Asclepios Research Project, INRIA Sophia-Antipolis, France

Abstract. In this work, a generative approach for patient-specific segmentation of brain tumors across different MR modalities is presented. It is based on the latent atlas approach presented in [7, 8]. The individual segmentation of each scan supports the segmentation of the ensemble by sharing common information. This common information, in the form of a spatial probability map of the tumor location is inferred concurrently with the evolution of the segmentations. 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.

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. The work is based on the latent atlas approach presented in [7, 8]. While the methodology can be applied to a variety of applications, here we demonstrate our algorithm on a problem of multimodal segmentation of brain tumors. 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.

The method we propose is almost fully automatic. 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 used to generate a sphere that initializes the segmentations. All model parameters, spatial and intensity, are inferred from the patient scans alone. The output of the algorithm consist of

individual segmentations for each modality. This is in contrast to many discriminative methods, e.g., [9], that use multimodal 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.

2 Problem definition and probabilistic model

This section summarizes the formulation of [6–8] for the joint segmentation of N aligned MR images. The input consists of N scans of a specific patient acquired via different imaging protocols. Our objective is to extract a brain tumor 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. 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:

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

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

3 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 $C_n = \{\mathbf{x} \in \Omega \mid \phi_n(\mathbf{x}) = 0\}$ defines the interface that partitions the image space of I_n into two disjoint regions ω and $\Omega \setminus \omega$. Similar to [4, 5] 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 [1]. 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.

Cost functional for segmentation The joint estimation problem of the hidden variables Γ and the unknown model parameters $\{\theta_\Gamma, \theta_\Gamma^n\}$ can be solved as an energy minimization problem. As in [6–8], we establish the correspondence between the log probability and the level-set energy terms. We also look for the fuzzy labeling functions $H(\phi_n)$ rather than the hard segmentations Γ_n .

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

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

where $Z(\theta_\Gamma)$ is the partition function and $\mathcal{N}(v)$ is the set of the closest neighbors of voxel v . We define the spatial energy term E_S based on the singleton term in Eq. (8). Using the level-set formulation we obtain:

$$E_S(\phi_n, \Theta) = - \int_{\Omega} [\log \theta_\Gamma(\mathbf{x}) H(\phi_n(\mathbf{x})) + \log(1 - \theta_\Gamma(\mathbf{x})) (1 - H(\phi_n(\mathbf{x})))] d\mathbf{x}. \quad (9)$$

The dynamically evolving latent atlas θ_Γ is obtained by optimizing the sum of the energy terms that depend on θ_Γ :

$$\hat{\theta}_\Gamma(\mathbf{x}) = \frac{1}{N} \sum_{n=1}^N \tilde{H}(\phi_n(\mathbf{x})). \quad (10)$$

The standard smoothness term used in level-set framework:

$$E_{\text{LEN}}(\phi_n) = \int_{\Omega} |\nabla H(\phi_n(\mathbf{x}))| d\mathbf{x}, \quad (11)$$

can be obtained as an approximation of the pairwise term in Eq. (8).

The energy term $E_I(\phi_n, \theta_I^n)$ corresponds to the image likelihood term in Eq. (3):

$$E_I(\phi_n, \Theta) = - \int_{\Omega} [\log p_{\text{in}}(I_n; \theta_{I,n}^{\text{in}}) H(\phi_n(\mathbf{x})) + \log p_{\text{out}}(I_n; \theta_{I,n}^{\text{out}}) (1 - H(\phi_n(\mathbf{x})))] d\mathbf{x}. \quad (12)$$

We assume that the intensities of the structure of interest are drawn from a normal distribution such that the pair of scalars $\theta_{I,n}^{\text{in}} = \{\mu_n^{\text{in}}, \sigma_n^{\text{in}}\}$ are the mean and standard deviation of the foreground intensities. We use a local-intensity model for the background intensity distributions in the spirit of [3], where $\theta_{I,n}^{\text{out}}(\mathbf{x}) = \{\mu_n^{\text{out}}(\mathbf{x}), \sigma_n^{\text{out}}(\mathbf{x})\}$ are the local mean and standard deviation of a small neighbourhood of \mathbf{x} that exclusively belongs to the background.

We construct the cost functional for $\phi_1 \dots \phi_N$ and the mode parameters by combing Eq. (12), (11) and (9):

$$E(\phi_1 \dots \phi_N, \Theta) = \gamma E_{\text{LEN}} + \beta E_I + \alpha E_S \quad (13)$$

where α , β and γ are positive scalars.

Gradient descent and parameter estimation. We optimize Eq. (13) by a set of alternating steps. For fixed model parameters the update of each level-set function ϕ_n in each iteration is determined by the following gradient descent equation:

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

where $\delta(\phi_n)$ is the derivative of $H(\phi_n)$ with respect to ϕ_n . For fixed segmentations ϕ_n , the model parameters are recovered by differentiating the cost functional in Eq. (13) with respect to each parameter.

4 Experiments

We evaluate our model on the BRATS challenge data set of 25 patients with glioma. The data set comprises T_1 , T_2 , FLAIR-, and post-Gadolinium T_1 MR images, all images are skull stripped and co-registered. The tumor is initialized through a sphere of 1-3 cm diameter, that is placed in the center of the tumor.

Exemplary segmentations that are returned from the present approach are shown in Figure 1 and quantitative results from a leave-one-out cross-validation are shown in Table 3. Note that the definition of “core” labels differs between ground truth (where it also includes the T_1 hypo-intense center of the tumor) and the algorithm tested (where it is only the T_1 gad hyper-intense area of the tumor) which results in misleading evaluation scores for the “core” class in low-grade cases.

5 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. Promising segmentation results on scans of 25 patients with Glioma were demonstrated.

Acknowledgements. This work was supported by NCCR Co-Me of the Swiss National Science Foundation, and INRIA CompuTumor.

References

1. T.F. Chan and L.A. Vese. Active contours without edges. *IEEE TIP*, 10(2):266–277, 2001.
2. L. Dice. Measure of the amount of ecological association between species. *Ecology*, 26(3):297–302, 1945.
3. S. Lankton and A. Tannenbaum. Localizing Region-Based Active Contours. *IEEE Transactions on Image Processing*, 17(11):2029–2039, 2008.
4. 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.
5. Pohl *et al.* Using the logarithm of odds to define a vector space on probabilistic atlases. *Medical Image Analysis*, 11(6):465–477, 2007.
6. T. Riklin Raviv, K. Van Leemput, W.M. Wells, and P. Golland. Joint segmentation of image ensembles via latent atlases. In *Medical Image Computing and Computer-Assisted Intervention*, pages 272–280, 2009.
7. T. Riklin Raviv, B.H. Menze, K. Van Leemput, B. Stieltjes, N. Weber, M.A. Ayache, W.M. Wells, and P. Golland. Joint segmentation via patient-specific latent anatomy model. In *Proc. of MICCAI Workshop on Probabilistic Models for Medical Image Analysis*, 2009.




Fig. 1. Representative results of the tumor segmentation. Shown are the segmentations for the four different modalities (columns 1-4), the labels inferred from the channel-wise segmentation (column 5), and the ground truth (column 6). The examples show that expert annotation may be disputable in some cases.

ID	Dice1	Sens1	Spec1	Dice2	Sens2	Spec2
BRATS_HG0027	0.65	0.752	0.993	0.813	0.812	0.998
BRATS_HG0026	0.678	0.602	0.998	0.413	0.293	0.999
BRATS_HG0025	0.59	0.933	0.991	0.083	0.043	1
BRATS_HG0024	0.659	0.873	0.996	0.825	0.779	1
BRATS_HG0022	0.699	0.82	0.997	0.608	0.498	0.999
BRATS_HG0015	0.756	0.908	0.991	0.831	0.731	1
BRATS_HG0014	0.27	0.665	0.987	0.59	0.45	0.999
BRATS_HG0013	0.684	0.713	1	0.894	0.996	1
BRATS_HG0012	0.637	0.709	0.999	0.098	0.077	1
BRATS_HG0011	0.798	0.742	0.998	0.882	0.971	0.998
BRATS_HG0010	0.097	0.145	0.997	0.276	0.945	0.996
BRATS_HG0009	0.795	0.8	0.995	0.548	0.377	1
BRATS_HG0008	0.734	0.771	0.992	0.841	0.885	0.998
BRATS_HG0007	0.407	0.361	0.996	0.278	0.298	0.996
BRATS_HG0006	0.648	0.843	0.991	0.817	0.716	1
mean	0.607	0.709	0.995	0.586	0.591	0.999
median	0.659	0.752	0.996	0.608	0.716	0.999

ID	Dice1	Sens1	Spec1	Dice2	Sens2	Spec2
BRATS_LG0015	0.37	0.712	0.997	0.116	0.066	1
BRATS_LG0014	0	0	1	0	0	1
BRATS_LG0013	0.326	0.631	0.995	0.452	0.293	1
BRATS_LG0012	0.563	0.721	0.999	0.822	0.762	1
BRATS_LG0011	0.262	0.958	0.993	0	0	1
BRATS_LG0008	0	0	1	0	0	1
BRATS_LG0006	0.556	0.985	0.996	0.73	0.832	1
BRATS_LG0004	0.513	0.492	0.997	0.022	0.072	0.997
BRATS_LG0002	0.636	0.734	0.989	0.242	0.178	0.997
BRATS_LG0001	0.345	0.648	0.997	0.843	0.774	1
mean	0.357	0.588	0.996	0.323	0.298	0.999
median	0.358	0.68	0.997	0.179	0.125	1

Table 1. Real data. Performance measures as returned by the online challenge tool (challenge.kitware.com/midas/) indicating Dice score, sensitivity and specificity (top: high-grade cases; bottom: low-grade cases). Class “1”, with results shown in the left column, refers to the “edema” labels. Class “2”, with results shown in the right column, refers to the “tumor core” labels (for both low and high grade cases). Note that this definition differs somewhat from the labels returned by the algorithm that only indicates T1gad hyper-intense regions as class 2, irrespectively of the grading (low/high) of the disease.

ID	Dice1	Sens1	Spec1	Dice2	Sens2	Spec2
SimBRATS_HG0025	0.023	0.105	0.996	0.631	0.593	0.998
SimBRATS_HG0024	0.217	0.993	0.999	0.968	0.948	1
SimBRATS_HG0023	0.007	0.078	0.997	0.454	0.419	0.997
SimBRATS_HG0022	0.689	0.682	0.996	0.002	0.001	0.999
SimBRATS_HG0021	0.312	0.214	0.998	0.02	0.011	1
SimBRATS_HG0020	0.138	0.127	0.996	0.315	0.249	0.998
SimBRATS_HG0019	0.45	0.349	0.998	0	0	1
SimBRATS_HG0018	0.01	0.047	0.996	0.579	0.552	0.997
SimBRATS_HG0017	0.147	0.179	0.998	0.499	0.348	1
SimBRATS_HG0016	0.033	0.091	0.995	0.681	0.667	0.997
SimBRATS_HG0015	0.36	0.289	0.998	0.234	0.186	0.998
SimBRATS_HG0014	0.362	0.3	0.998	0.451	0.406	0.998
SimBRATS_HG0013	0.623	0.564	0.996	0.004	0.002	0.999
SimBRATS_HG0012	0.44	0.36	0.999	0.035	0.022	1
SimBRATS_HG0011	0.453	0.518	0.997	0.366	0.235	1
SimBRATS_HG0010	0.528	0.867	0.999	0.974	0.978	1
SimBRATS_HG0009	0.762	0.788	1	0.958	0.977	1
SimBRATS_HG0008	0.381	0.352	0.996	0.454	0.386	0.999
SimBRATS_HG0007	0.635	0.689	0.995	0.559	0.75	0.997
SimBRATS_HG0006	0.011	0.037	0.998	0.373	0.274	0.999
SimBRATS_HG0005	0.63	0.615	0.996	0.019	0.015	0.999
SimBRATS_HG0004	0.33	0.311	0.996	0.485	0.475	0.998
SimBRATS_HG0003	0.63	0.593	0.998	0.317	0.314	0.999
SimBRATS_HG0002	0.405	0.819	0.999	0.924	0.875	1
SimBRATS_HG0001	0.592	0.856	0.999	0.971	0.982	1
mean	0.367	0.433	0.997	0.451	0.427	0.999
median	0.381	0.352	0.998	0.454	0.386	0.999

Table 2. Simulated data (high grade). Performance measures as returned by the online challenge tool (challenge.kitware.com/midas/) indicating Dice score, sensitivity and specificity (top: high-grade cases; bottom: low-grade cases). Class “1”, with results shown in the left column, refers to the “edema” labels. Class “2”, with results shown in the right column, refers to the “tumor core” labels (for both low and high grade cases). Note that this definition differs somewhat from the labels returned by the algorithm that only indicates T1gad hyperintense regions as class 2, irrespectively of the grading (low/high) of the disease.

ID	Dice1	Sens1	Spec1	Dice2	Sens2	Spec2
SimBRATS_LG0025	0.042	0.528	0.993	0.05	0.026	1
SimBRATS_LG0024	0.404	0.997	0.993	0.137	0.074	1
SimBRATS_LG0023	0.662	0.74	0.997	0.008	0.004	1
SimBRATS_LG0022	0.404	0.551	0.997	0.007	0.004	1
SimBRATS_LG0021	0	0	1	0	0	1
SimBRATS_LG0020	0.367	0.702	0.997	0.023	0.012	1
SimBRATS_LG0019	0.378	0.367	0.998	0.015	0.008	1
SimBRATS_LG0018	0	0	1	0	0	1
SimBRATS_LG0017	0.632	0.678	0.998	0	0	1
SimBRATS_LG0016	0.699	0.858	0.993	0.08	0.045	1
SimBRATS_LG0015	0.46	0.578	0.997	0.01	0.005	1
SimBRATS_LG0014	0.02	0.416	0.996	0.025	0.013	1
SimBRATS_LG0013	0.402	0.447	0.997	0.014	0.007	1
SimBRATS_LG0012	0	0	1	0	0	1
SimBRATS_LG0011	0	0	1	0	0	1
SimBRATS_LG0010	0.078	0.694	0.991	0.21	0.117	1
SimBRATS_LG0009	0.394	0.507	0.997	0.035	0.018	1
SimBRATS_LG0008	0.051	0.998	0.994	0.13	0.07	1
SimBRATS_LG0007	0	0	1	0	0	1
SimBRATS_LG0006	0.395	0.857	0.995	0.054	0.028	1
SimBRATS_LG0005	0.483	0.994	0.993	0.089	0.047	1
SimBRATS_LG0004	0.317	0.316	0.998	0.002	0.001	1
SimBRATS_LG0003	0.359	0.546	0.997	0.007	0.004	1
SimBRATS_LG0002	0	0	1	0	0	1
SimBRATS_LG0001	0.489	0.39	0.998	0.001	0.001	1
mean	0.281	0.487	0.997	0.0359	0.0194	1
median	0.367	0.528	0.997	0.01	0.005	1

Table 3. Simulated data (low grade). Performance measures as returned by the online challenge tool (challenge.kitware.com/midas/) indicating Dice score, sensitivity and specificity (top: high-grade cases; bottom: low-grade cases). Class “1”, with results shown in the left column, refers to the “edema” labels. Class “2”, with results shown in the right column, refers to the “tumor core” labels (for both low and high grade cases). Note that this definition differs somewhat from the labels returned by the algorithm that only indicates T1gad hyperintense regions as class 2, irrespectively of the grading (low/high) of the disease.

8. T. Riklin Raviv, K. Van Leemput, B.H. Menze, W.M. Wells, and P. Golland. Segmentation of image ensembles via latent atlases. *Medical Image Analysis*, 14(5):654–665, 2010.
9. Van Leemput *et al.* Automated segmentation of multiple sclerosis lesions by model outlier detection. *IEEE TMI*, 20:677–688, 2001.