Segmenting Glioma in Multi-Modal Images using a Generative Model for Brain Lesion Segmentation

Bjoern H. Menze[1,2], Koen Van Leemput[3], Danial Lashkari[4]
Marc-André Weber[5], Nicholas Ayache[2], and Polina Golland[4]

[1] Computer Vision Laboratory, ETH Zurich, Switzerland
[2] Asclepios Research Project, INRIA Sophia-Antipolis, France
[3] Radiology, Massachusetts General Hospital, Harvard Medical School, USA
[4] Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, USA
[5] Diagnostic Radiology, Heidelberg University Hospital, Germany

1 Introduction

In this paper, we evaluate a fully automated method for channel-specific tumor segmentation in multi-dimensional images proposed by us in [1]. The method represents a tumor appearance model for multi-dimensional sequences that provides channel-specific segmentation of the tumor. Its generative model shares information about the spatial location of the lesion among channels while making full use of the highly specific multi-modal signal of the healthy tissue classes for segmenting normal tissues in the brain. In addition to tissue types, the model includes a latent variable for each voxel encoding the probability of observing tumor at that voxel, based on the ideas from [2, 3]. This extends the general "EM segmention" algorithm for situations when specific spatial structures cannot be described sufficiently through population priors. Different from [1], we now use a simplified EM algorithm for estimating the tissue state that also allows us to enforce additional constraints for segmenting lesions that are either hyper- or hypo-intense with respect to other tissues visible in the same image.

2 Generative Tumor Model

We use a generative modeling approach, in which we first build an explicit statistical model of image formation and subsequently use this model to derive a fully automatic segmentation algorithm. We follow the description of the model from [1].

Normal state We model the *normal state* of the healthy brain using a spatially varying probabilistic prior $\boldsymbol{\pi}_k$ for each of the K tissue classes. This population prior (atlas) is estimated from prior examples and is assumed to be known. At each voxel i, the atlas defines a multinomial distribution for the tissue label k_i:

$$p(k_i = k) = \pi_{ki}. \tag{1}$$

The normal state k_i is shared among all C channels at voxel i. In our experiments we assume $K = 3$, representing gray matter (GM), white matter (WM) and cerebrospinal fluid (CSF).

Tumor state We model the *tumor state* using a spatially varying "latent" probabilistic atlas $\boldsymbol{\alpha}$, similar to [2], that is specific to the given image data set or patient. At each voxel i, this atlas provides a scalar parameter α_i that defines the probability of observing tumor at that voxel. Parameter α_i is unknown and is estimated as part of the segmentation process. We define a latent tumor state $t_i^c \in \{0, 1\}$ that indicates the presence of tumor in channel c at voxel i and model it as a Bernoulli random variable with parameter α_i. We form a binary tumor state vector $\boldsymbol{t}_i = [t_i^1, \ldots, t_i^C]^T$ indicating the tumor presence for all c observations at voxel i, with probability

$$p(\boldsymbol{t}_i; \alpha_i) = \prod_c p(t_i^c; \alpha_i) = \prod_c \alpha_i^{t_i^c} \cdot (1 - \alpha_i)^{1 - t_i^c}. \tag{2}$$

Observation model The *image observations* y_i^c are generated by Gaussian intensity distributions for each of the K tissue classes and the C channels, with mean μ_k^c and variance v_k^c, respectively. In tumor tissue (i.e., if $t_i^c = 1$) the normal observations are replaced by intensities from another set of channel-specific Gaussian distributions with mean μ_T^c and variance v_T^c, representing the tumor class. Letting $\boldsymbol{\theta}$ denote the set of all mean and variance parameters, and $\boldsymbol{y}_i = [y_i^1, \ldots, y_i^C]^T$ denote the vector of the intensity observations at voxel i, we define the data likelihood:

$$\begin{aligned} p(\boldsymbol{y}_i | \boldsymbol{t}_i, k_i; \boldsymbol{\theta}) &= \prod_c p(y_i^c | t_i^c, k_i; \boldsymbol{\theta}) \\ &= \prod_c \left[\mathcal{N}(y_i^c \ ; \ \mu_{k_i}^c, v_{k_i}^c)^{1 - t_i^c} \cdot \mathcal{N}(y_i^c \ ; \ \mu_T^c, v_T^c)^{t_i^c} \right], \end{aligned} \tag{3}$$

where $\mathcal{N}(\cdot \ ; \ \mu, v)$ is the Gaussian distribution with mean μ and variance v.

Joint model Finally, the *joint probability* of the atlas, the latent tissue class and the observed variables

$$p(\boldsymbol{y}_i, \boldsymbol{t}_i, k_i; \boldsymbol{\theta}, \alpha_i) = p(\boldsymbol{y}_i | \boldsymbol{t}_i, k_i; \boldsymbol{\theta}) \cdot p(\boldsymbol{t}_i; \alpha_i) \cdot p(k_i) \tag{4}$$

is the product of the components defined in Eqs. (1-3).

3 Inference

Maximum Likelihood Parameter Estimation We seek Maximum Likelihood estimates of the model parameters $\{\boldsymbol{\theta}, \boldsymbol{\alpha}\}$:

$$\langle \widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\alpha}} \rangle = \arg \max_{\langle \boldsymbol{\theta}, \boldsymbol{\alpha} \rangle} \ p(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_N; \boldsymbol{\theta}, \boldsymbol{\alpha}) = \arg \max_{\langle \boldsymbol{\theta}, \boldsymbol{\alpha} \rangle} \ \prod_{i=1}^{N} p(\boldsymbol{y}_i; \boldsymbol{\theta}, \boldsymbol{\alpha}),$$

where N is the number of voxels in the volume and

$$p(\boldsymbol{y}_i;\boldsymbol{\theta},\boldsymbol{\alpha}) = \sum_{\boldsymbol{t}_i}\sum_{k_i} p(\boldsymbol{y}_i,\boldsymbol{t}_i,k_i;\boldsymbol{\theta},\boldsymbol{\alpha}) = \sum_{\boldsymbol{s}_i} p(\boldsymbol{y}_i,\boldsymbol{s}_i;\boldsymbol{\theta},\boldsymbol{\alpha}).$$

For summing over values of $\boldsymbol{t}_i$ and k_i in Eq. (4), we follow the same approach as in [1], but – rather than summing over the two parameters individually – we now sum over tissue state vector $\boldsymbol{s}_i$ that is obtained by expanding $\boldsymbol{t}_i$ and k_i into one state vector. This state vector $\boldsymbol{s}$ indicates tumor $s_i^c = T$ in all channels with $t_i^c = 1$, and normal tissue $s_i^c = k_i$ for all other channels. As an example, with $\boldsymbol{t}_i = [0,0,1,1]$ and $k_i = WM$ indicating tumor in channels 3 and 4 and a white matter image intensity in all healthy channels, we obtain a tissue state vector $\boldsymbol{s}_i = [WM, WM, T, T]$. Letting $\{\widetilde{\boldsymbol{\theta}}, \widetilde{\boldsymbol{\alpha}}\}$ denote the current parameter estimates, we can compute the posterior probability of any of the resulting $K * 2^C$ tissue state vectors $\boldsymbol{s}_i$ that may characterize the multimodal image intensity pattern observed in voxel i. Writing out the components of Eq. (4) we obtain for $\boldsymbol{s}_i$ (using the corresponding $\boldsymbol{t}_i(\boldsymbol{s}_i)$ and $k_i(\boldsymbol{s}_i)$ for simplicity of the notation):

$$p(\boldsymbol{s}_i|\boldsymbol{y}_i;\widetilde{\boldsymbol{\theta}},\widetilde{\boldsymbol{\alpha}}) \;\propto\; \pi_{ki}\prod_c \left[\mathcal{N}(y_i^c \;;\; \tilde{\mu}_T^c, \tilde{v}_T^c)^{t_i^c}\alpha_i^{t_i^c}\cdot\mathcal{N}(y_i^c \;;\; \tilde{\mu}_k^c, \tilde{v}_k^c)^{1-t_i^c}(1-\alpha_i)^{1-t_i^c}\right] \tag{5}$$

As an additional constraint we only consider state vectors $\boldsymbol{s}_i$ that are biologically reasonable. We rule out, for example, state vectors that indicate at the same time CSF and tumor, or that correspond to observing tumor-specific changes in the T1gad channel (that is characteristic for the tumor core), while T2 and FLAIR do not show tumor specific changes in the same location. Choosing appropriate constraints reduces the total number of states $|\boldsymbol{S}|$ to be summed over in Eq. 5 significantly. Similar to the double EM-type minorization from [1] – that updated $\boldsymbol{t}_i$ and k_i iteratively – we arrive at closed-form update expressions that guarantee increasingly better estimates of the model parameters. The updates are intuitive: the latent tumor prior $\tilde{\alpha}_i$ is an average of the corresponding posterior estimated and the intensity parameters $\tilde{\mu}_k^c$ and $\tilde{v}_k^c$ are updated with the weighted statistics of the data for the healthy tissues and for the tumor class. We iterate the estimation of the parameters $\{\widetilde{\boldsymbol{\theta}}, \widetilde{\boldsymbol{\alpha}}\}$ and the computation of the posterior probabilities $p(\boldsymbol{s}_i|k_i,\boldsymbol{y}_i;\widetilde{\boldsymbol{\theta}},\widetilde{\boldsymbol{\alpha}})$ until convergence that is typically reached after 10-15 updates. During the iterations we enforced that tumor voxels are hyper- or hypo-intense with respect to the current average μ_k^c of the white matter tissue (hypo-intense for T1, hyper-intense for T1gad, T2, FLAIR) by reducing the class weight for observations that do not comply with this constraint, similar to [4].

Spatial regularization Little spatial context is used in the basic model, as we assume the tumor state $\boldsymbol{t}_i$ in each voxel to be independent from the state of other voxels Eq. 3). It is only the atlas $\boldsymbol{\pi_k}$ that encourages smooth classification for the healthy tissue classes by imposing similar priors in neighboring voxels. To encourage a similar smoothness of the tumor labels, we extend the latent atlas α to include a Markov Random Field (MRF) prior, relaxing the MRF to

a mean-field approximation with an efficient approximate algorithm. Different from [1], we now use channel-specific regularization parameters β^c that are all in the range of .3 to 1.

Channel-specific tumor probabilities, and semantic interpretation Once we have an estimate of the model parameters $\{\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\alpha}}\}$, we can evaluate the probability that tumor is visible in channel c of voxel i by summing over all the configurations $\boldsymbol{t}_i$ for which $s_i^c = T$ or $t_i^c = 1$, respectively:

$$p(t_i^c = 1|\boldsymbol{y}_i; \widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\alpha}}) = p(s_i^c = T|\boldsymbol{y}_i; \widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\alpha}}) = \sum_{\boldsymbol{s}_i} \delta(s_i^c, T)\ p(\boldsymbol{s}_i|\boldsymbol{y}_i; \widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\alpha}}), \qquad (6)$$

where δ is the Kroneker delta that is 1 for $s_i^c = T$ and 0 otherwise.

We then assign channel c of voxel i to tumor if $p(t_i^c = 1|\boldsymbol{y}_i; \widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\alpha}}) > 0.5$. For a semantic interpretation that is in line with the class definitions of the segmentation challenge, we label voxels that show tumor specific changes in the T2 channel as *edema*, and voxels that show hyper-intense tumor specific changes as *tumor core*. All other image voxels are considered to be normal. Moreover, we remove any isolated regions that is smaller than .5 cm^3 in size.

4 Experiments

We evaluate our model on a 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 using an affine registration. We segment the volume into the three healthy and an outlier class using a freely available implementation of the EM segmentation with bias correction [5, 4]. Outliers are defined as being more than three standard deviations away from the centroid of any of the three normal tissue classes. We apply our algorithm to the bias field corrected volumes and initialize intensity parameters with values estimated in the initial segmentation. We initialize the latent atlas $\boldsymbol{\alpha}$ to 0.7 time the local prior for the presence of gray or white matter.

Channels-specific segmentations returned by our algorithm are transformed to *Edema* and *Core* classes as detailed above. Exemplary segmentations are shown in Figure 1 and quantitative results from a leave-one-out cross-validation are shown in Table 1. Note that the definition of "core" labels differs between ground truth (where it also includes the T1 hypo-intense center of the tumor) and the algorithm tested (where it is only the T1gad hyper-intense area of the tumor) leading to misleading evaluation scores for low-grade cases and, to some degree, for high-grade core labels. Please note that another submission to the BRATS challenge [6] deals with further processing the probability maps presented here.

Acknowledgements. This work was supported by the German Academy of Sciences Leopoldina (Fellowship Programme LPDS 2009-10), the Academy of Finland (133611), INRIA CompuTumor, NIH NIBIB NAMIC U54-EB005149, NIH NCRR NAC P41-RR13218, NIH NINDS R01-NS051826, NIH R01-NS052585, NIH R01-EB006758, NIH R01-EB009051, NIH P41-RR014075 and the NSF CAREER Award 0642971.

ID	Dice1	Sens1	Spec1	Dice2	Sens2	Spec2
BRATS_HG0027	0.633	0.649	0.995	0.728	0.619	0.999
BRATS_HG0026	0.681	0.616	0.998	0.443	0.369	0.999
BRATS_HG0025	0.643	0.704	0.996	0.154	0.087	1
BRATS_HG0024	0.652	0.685	0.998	0.71	0.639	0.999
BRATS_HG0022	0.689	0.683	0.998	0.463	0.311	1
BRATS_HG0015	0.762	0.699	0.998	0.691	0.534	1
BRATS_HG0014	0.217	0.453	0.989	0.457	0.303	1
BRATS_HG0013	0.429	0.647	0.999	0.74	0.983	1
BRATS_HG0012	0.373	0.58	0.997	0.068	0.043	1
BRATS_HG0011	0.606	0.464	0.999	0.57	0.54	0.998
BRATS_HG0010	0.381	0.792	0.996	0.724	0.77	1
BRATS_HG0009	0.697	0.594	0.997	0.486	0.38	0.997
BRATS_HG0008	0.652	0.56	0.996	0.697	0.556	1
BRATS_HG0007	0.542	0.492	0.997	0.775	0.727	0.999
BRATS_HG0006	0.649	0.621	0.997	0.65	0.505	1
mean	0.574	0.616	0.997	0.557	0.491	0.999
median	0.643	0.621	0.997	0.65	0.534	1

ID	Dice1	Sens1	Spec1	Dice2	Sens2	Spec2
BRATS_LG0015	0.373	0.523	0.998	0	0	1
BRATS_LG0014	0.182	0.335	0.998	0	0	1
BRATS_LG0013	0.185	0.324	0.995	0.17	0.099	1
BRATS_LG0012	0.42	0.79	0.997	0.005	0.002	1
BRATS_LG0011	0.344	0.777	0.996	0.001	0	1
BRATS_LG0008	0.471	0.386	1	0.675	0.547	1
BRATS_LG0006	0.625	0.809	0.998	0.591	0.507	1
BRATS_LG0004	0.75	0.764	0.998	0.011	0.006	1
BRATS_LG0002	0.584	0.622	0.991	0.109	0.059	1
BRATS_LG0001	0.3	0.495	0.997	0.838	0.777	1
mean	0.423	0.582	0.997	0.24	0.2	1
median	0.396	0.572	0.998	0.06	0.0325	1

Table 1. 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.

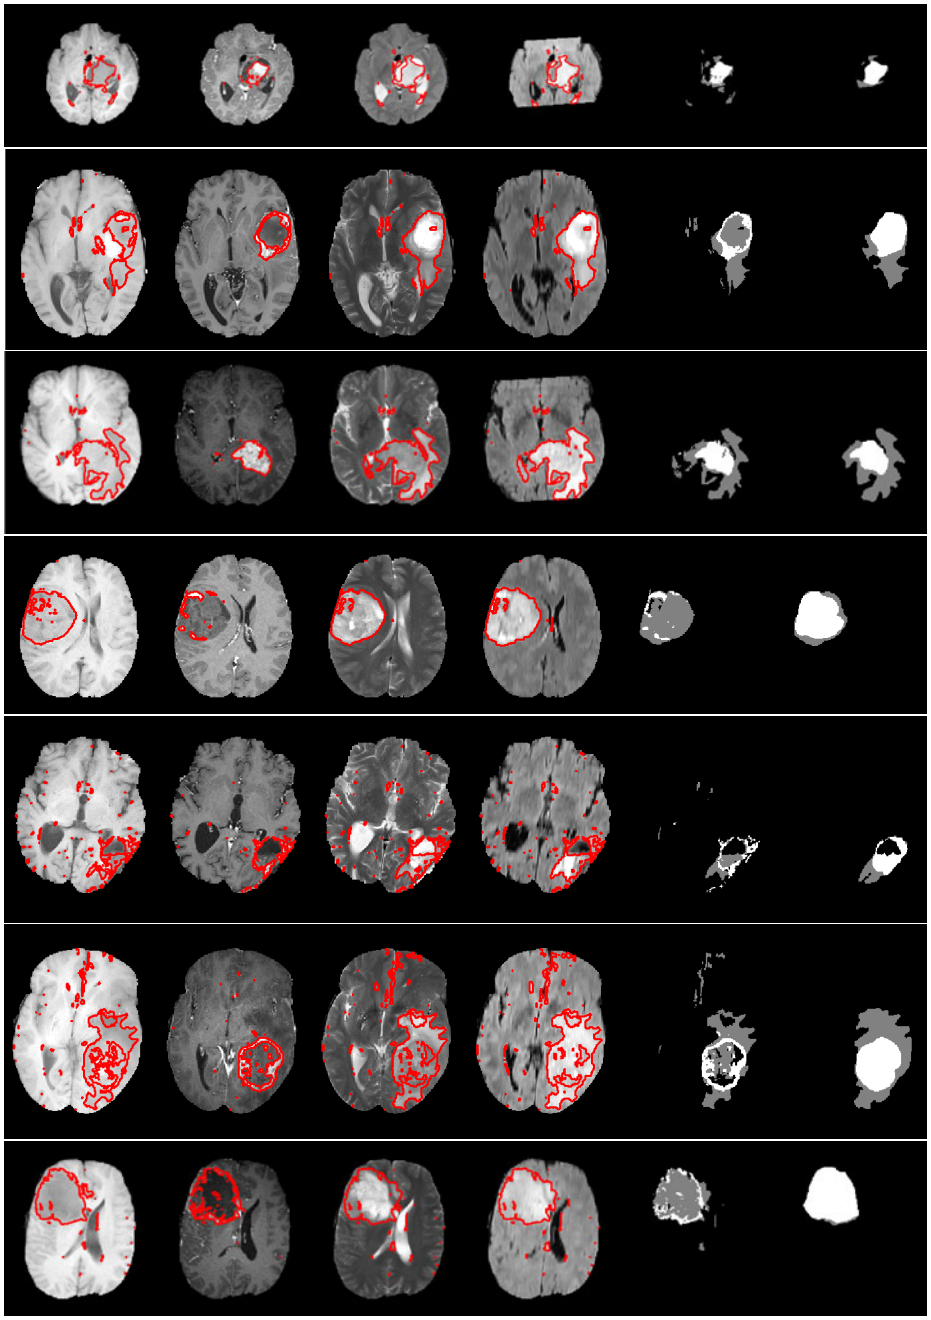

Fig. 1. Representative results of the channel-wise tumor segmentation. Shown are the MR images together with the most likely tumor areas (outlined red). The first four columns show T1, T1gad, T2 and FLAIR MRI, lesions are hyper-intense with respect to the gray matter for T1gad, T2 and FLAIR, they are hypo-intense in T1. The last two columns show the labels inferred from the channel-specific tumor segmentation (column 5), and the ground truth (column 6). The examples show that expert annotation may be disputable in some cases (e.g., rows 4, 5, 6).

References

1. Menze, B.H., Van Leemput, K., Lashkari, D., Weber, M.A., Ayache, N., Golland, P.: A generative model for brain tumor segmentation in multi-modal images. In: Proc MICCAI. (2010) 151–159
2. Riklin-Raviv, T., Menze, B.H., Van Leemput, K., Stieltjes, B., Weber, M.A., Ayache, N., Wells, W.M., Golland, P.: Joint segmentation via patient-specific latent anatomy model. In: Proc MICCAI-PMMIA (Workshop on Probabilistic Models for Medical Image Analysis). (2009) 244–255
3. Riklin-Raviv, T., Van Leemput, K., Menze, B.H., Wells, 3rd, W.M., Golland, P.: Segmentation of image ensembles via latent atlases. Med Image Anal **14** (2010) 654–665
4. Van Leemput, K., Maes, F., Vandermeulen, D., Colchester, A., Suetens, P.: Automated segmentation of multiple sclerosis lesions by model outlier detection. IEEE T Med Imaging **20** (2001) 677–688
5. Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P.: Automated model-based bias field correction of MR images of the brain. IEEE T Med Imaging **18** (1999) 885–896
6. Menze, B.H., Geremia, E., Ayache, N., Szekely, G.: Segmenting glioma in multi-modal images using a generative-discriminative model for brain lesion segmentation. In: Proc MICCAI-BRATS (Multimodal Brain Tumor Segmentation Challenge). (2012) 8p