

In Martin A. Styner, Elsa D. Angelini (Eds.):
Medical Imaging 2016: Image Processing (San Diego, CA, February 2016)
SPIE Vol. 9784, 97842S, 2016.

The final publication is available in the SPIE Digital Library

Copyright 2016 Society of Photo Optical Instrumentation Engineers. One print or electronic copy may be made for personal use only. Systematic reproduction and distribution, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

Automatic brain tumor segmentation with a fast Mumford-Shah algorithm

Sabine Müller^{a,b}, Joachim Weickert^b, and Norbert Graf^a

^aDepartment of Pediatric Oncology and Hematology, Saarland University Hospital,
Homburg, Germany

^bMathematical Image Analysis Group, Saarland University, Campus E1.7,
Saarbrücken, Germany

ABSTRACT

We propose a fully-automatic method for brain tumor segmentation that does not require any training phase. Our approach is based on a sequence of segmentations using the Mumford-Shah cartoon model with varying parameters. In order to come up with a very fast implementation, we extend the recent primal-dual algorithm of Strelakovsky et al. (2014) from the 2D to the medically relevant 3D setting. Moreover, we suggest a new confidence refinement and show that it can increase the precision of our segmentations substantially. Our method is evaluated on 188 data sets with high-grade gliomas and 25 with low-grade gliomas from the BraTS14 database. Within a computation time of only three minutes, we achieve Dice scores that are comparable to state-of-the-art methods.

Keywords: segmentation, Mumford-Shah, brain tumor, MRI

1. INTRODUCTION

Brain tumors account only for a very small fraction of all types of cancer, but are also among the most fatal forms of cancer. Gliomas, developing from the glial cells, are the most frequent primary brain tumors. The fast growing and more aggressive types of gliomas called high-grade gliomas, come with an median overall survival rate up to 15 months.¹ The standard diagnosis technique for brain tumor is magnetic resonance imaging (MRI)² providing detailed information about the tumor and the surrounding brain. Tumor segmentation is a crucial task in surgical and treatment planning. The clinicians' standard technique is still manual tumor segmentation, which tends to inter- and intra-rater variability.³ Moreover, the time required to manually annotate and segment the data is high.

Therefore, much research is performed to develop methods for automatic brain tumor segmentation; see⁴ and the references therein. Fully automated segmentation is a challenging task, especially for high-grade gliomas as they usually show diffuse and irregular boundaries and have intensities overlapping with normal brain tissue. Moreover, acquisition parameters are not standardized, and different parameter settings can have a substantial impact on the visual appearance of the tumor. This makes it difficult to compare the quality of different methods

Further author information: (Send correspondence to S. Müller)

S. Müller: E-mail: smueller@mia.uni-saarland.de, Telephone: +49 (0)681 302 57353

J. Weickert: E-mail: weickert@mia.uni-saarland.de, Telephone: +49 (0)681 302 57341

N. Graf: E-mail: norbert.graf@uniklinikum-saarland.de, Telephone: +49 (0)6841 16 28411

for brain tumor segmentation. As a step towards an unbiased performance evaluation, the Multimodal Brain Tumor Image Segmentation (BraTS) database has been created,⁵ and many recent approaches report benchmark results on either the full data set or parts of it; see e.g.^{6,7} Most of the better methods, however, are relatively complex or time-consuming or require a training phase.

Brain tumor segmentation methods can be roughly categorized into two groups: semi-automatic and fully-automatic segmentation. Most of the fully-automatic segmentation approaches include classification⁶ or clustering algorithms or rely on atlas-based methods. The majority of semi-automatic segmentation methods are based on active contour models. Active contours are often modeled implicitly as level-sets and can be further divided in edge- and region-based methods. Edge-based active contours heavily depend on the image gradient. Hence, they are less suited for segmenting edematous regions of high-grade gliomas. However, region-based active contours are more robust when parts of the object boundaries are diffuse, since they reward homogeneity within the segment.

The most popular approach for region-based active contours is the Chan-Vese model.⁸ In its basic formulation, it segments an image into fore- and background. The granularity of this segmentation is steered by a single parameter that weights the length of the segment boundaries. The global behavior of this segmentation model is both a blessing and a curse - a blessing because it is very robust to noise and initialization; a curse because it is prone to false segments and intensity inhomogeneities. The Chan-Vese model can be seen as a simplification of the cartoon limit of the Mumford-Shah functional.^{9,10} So far, most algorithms for the Mumford-Shah cartoon model were rather time consuming. This situation has been improved substantially by recent results of Strekalovskiy et al.,¹¹ who proposed an efficient primal-dual algorithm to minimize this functional.

However, the crux in tumor segmentation approaches with this functional is to find an appropriate parameter setting and to identify segments with tumor tissue. If the steering parameter is too small, the result contains many small segments such that it is hard to determine the appropriate ones. In contrast, if the parameter is too large, the segmentation is too coarse, and the segmentation does not approximate the tumor boundaries very well.

We address the before mentioned problems by proposing a fast and fully-automatic method for brain tumor segmentation. It is simple, transparent, and makes consequent use of a single ingredient: the Mumford-Shah cartoon model. In a first step, we segment the complete tumor. We tackle the difficulty of finding an appropriate parameter setting by solving the Mumford-Shah cartoon model iteratively with varying boundary weights until we reach a segmentation with only few segments. We refine these tumor boundaries afterwards by introducing a new confidence measure that increases the precision significantly. In a second step, we use this segmentation to determine the boundaries of the tumor subcomponents (edema, necrosis, enhancing-, and non-enhancing parts). We evaluate our approach with the training data of the MICCAI BraTS14 challenge.⁵ We compare our method to other state-of-the-art algorithms by means of the test data of the MICCAI BraTS13 contrast-enhanced T_{1c} , T_2 , and T_2 -Flair scans.

2. METHOD

To describe our algorithm we first discuss the preprocessing steps. Afterwards we explain the sequence of Mumford-Shah segmentations and their refinement with a suitable confidence measure. Finally we address the problem of identifying the fine-structure of the tumor.

2.1 Preprocessing

Since MRI scans may suffer from non-uniformities within each data set, we first apply the N4ITK filter¹² to all scans to correct for these artifacts. Then we reduce the influence of the background on the segmentation outcome by replacing all gray values smaller than the average gray value μ with μ . Finally, we compute and equalize the histograms of each 3D scan to exploit the contrast in the data in the best possible way. All data channels are rescaled to intensity values in the interval $[0, 255]$.

2.2 The Mumford-Shah cartoon model

Let us consider a cubic data domain $\Omega \subset \mathbb{R}^3$ and some volumetric data set $\mathbf{f} : \Omega \rightarrow \mathbb{R}^m$. For our application, its m channels describe different MRI modalities such as T_1 , T_{1c} , T_2 and T_2 -Flair. Then a segmentation of \mathbf{f} by

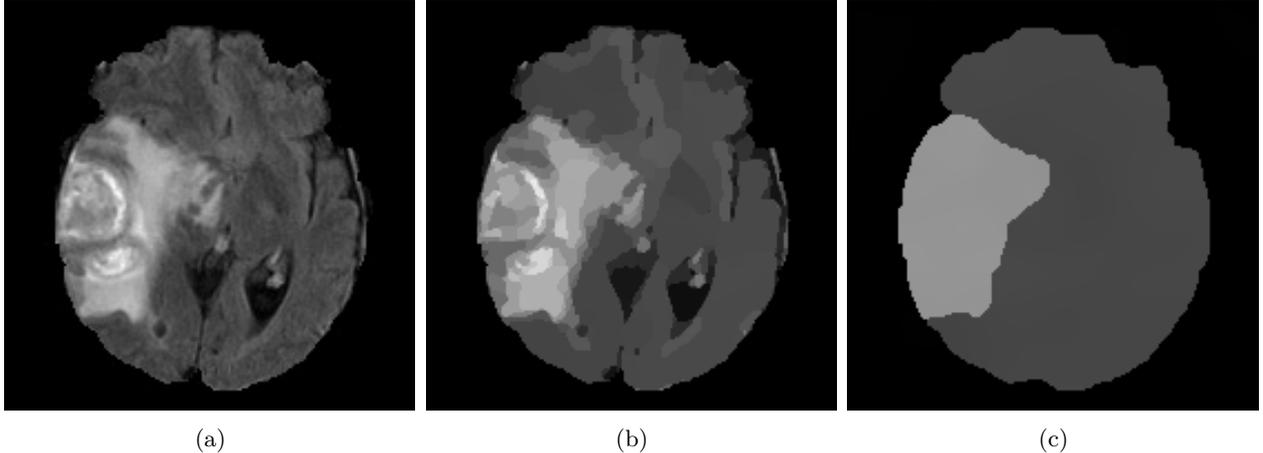


Figure 1: Exemplary results for different penalizations of the boundary length. From left to right: (a) T_2 -Flair input image. (b) Result for $\nu = 1000$. (c) Result for $\nu = 340000$.

means of the Mumford–Shah cartoon model^{9,10} minimizes the energy functional

$$E(\mathbf{u}, C) = \sum_i \int_{\Omega_i} \|\mathbf{u} - \mathbf{f}\|^2 d\mathbf{x} + \nu \ell(C). \quad (1)$$

Here the a priori unknown number of segments Ω_i partition the data domain Ω , the function \mathbf{u} denotes a piecewise constant approximation of \mathbf{f} , $\|\cdot\|$ is the Euclidean norm in \mathbb{R}^m , and the segment boundaries C have a (Hausdorff) length of $\ell(C)$. The first term of the energy is a data term that penalizes fluctuations within each segment, while the second term favors short segment boundaries. The parameter $\nu > 0$ allows to weight the boundary length in relation to the inhomogeneities within each segment. Obviously the choice of ν is of crucial importance: The higher the value of this parameter, the less segments are contained in the final result. In Fig. 1, the number of segments decreases with increasing penalization of the boundary length. At the same time, the inhomogeneities within individual segments increases.

While early algorithms for the Mumford–Shah cartoon model are based on region merging concepts,¹³ a recent approach by Strelakovsky et al.¹¹ describes a very fast approximation of this model by means of primal–dual optimization ideas. Due to its intrinsic parallelism, it is well-suited for parallel processing hardware such as GPUs. In its original formulation, this algorithm has been specified for 2D data sets. For our framework, we have extended this approach in a straightforward way to the segmentation of volumetric data.

2.3 Segmentation of the complete tumor

On MRI T_2 -Flair scans, high-grade gliomas contain areas that are brighter than the brain tissue. To detect them, we segment for a bright outlier in intensity in the following way: We also include T_1 data in our segmentation, since our experiments show that this makes the segmentation process more robust against small distortions. We start with the parameter $\nu = 400,000$ and check if this gives a segmentation into two areas: the tumor and the background. Since the algorithm of Strelakovsky et al. might give more than one segment, we postprocess the result with an Otsu thresholding.¹⁴ If the area of the thresholded tumor is larger than 50% of the brain volume, this is an indication that ν was too large such that the tumor has been merged with its background. In this case, we reduce ν by 15% and start the procedure again. This approach is repeated recursively until we have a segmentation where the tumor volume is below 50% of the brain volume.

2.4 Confidence refinement

The total tumor segmentation from the previous subsection tends to give an area that is somewhat larger than the real tumor area: It favors sensitivity over precision. To refine this segmentation with a confidence refinement

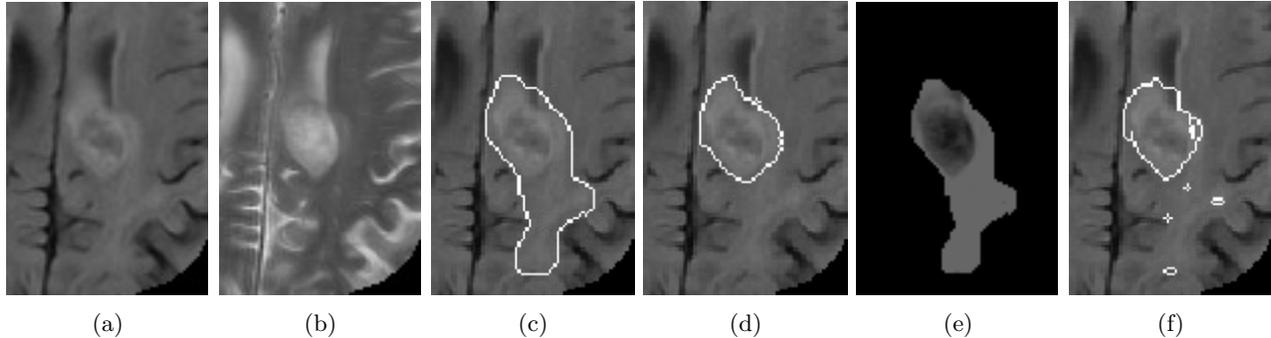


Figure 2: Illustration of the confidence refinement procedure. From left to right: (a) T_2 -Flair input image for (1). (b) T_2 input image for consistency refinement. (c) Segment boundaries before refinement. Dice score: 0.51, sensitivity: 0.82, precision: 0.37 (3D tumor volume). (d) Boundaries of the ground truth segmentation. (e) Difference map $d(\mathbf{x})$. (f) Segment boundaries after refinement. Dice score: 0.77, sensitivity: 0.7, precision: 0.86 (3D tumor volume).

postprocessing, we investigate how well the data term of the Mumford-Shah cartoon model is fulfilled locally. It is sufficient to do this only in the T_2 channel, since this channel always contains the tumor and it has not been used before. Thus, in any location \mathbf{x} of the tumor area, we measure the difference $d(\mathbf{x}) := |u(\mathbf{x}) - f(\mathbf{x})|$, where f and u denote the original resp. segmented intensity values of the T_2 channel. Since u and f have values in $[0, 255]$, the difference d lies in the same interval. Large values of $d(\mathbf{x})$ indicate that the local confidence in our segmentation result should be low. We quantize the maximal d -range $[0, 255]$ into 256 bins, and we discard those voxels from the segment where the distance lies in the highest bin with nonvanishing contributions.

This confidence refinement is a trade-off between gaining precision and losing sensitivity: In general, we lose a small amount of sensitivity but gain remarkably in precision: In Fig. 2, the Dice score¹⁵ of the segmentation improved from 0.51 to 0.77, and the precision increased even from 0.37 to 0.86, while the sensitivity deteriorated only mildly from 0.82 to 0.7. This illustrates the usefulness of our confidence refinement.

2.5 Segmentation of the tumor fine-structure

We use this first segmentation to determine further tumor subcomponents: We minimize the Mumford-Shah cartoon model again with a very small boundary penalization ($\nu = 1$), but this time exclusively on the T_{1c} scans and in the previously defined segment. Afterwards we use Otsu’s thresholding to identify the active tumor, i.e. enhancing- and non-enhancing tumor core. In this way, we get a splitting of the complete tumor region into active tumor and necrosis/edema. To get the final subcomponents, we apply Otsu’s method on both subcomponents and split the first component, i.e. active tumor, into its enhancing and non-enhancing part and the second subcomponent into necrosis and edema.

3. EXPERIMENTS

We evaluate the performance of the proposed method on image data of the high- and low-grade gliomas of the MICCAI BraTS14 training data set.⁵ Unfortunately, BraTS14 test data changed over time which makes it impossible to compare our method with other approaches on this data sets. Therefore, we decided to benchmark our approach additionally on BraTS13 test data.

Both challenges contain skull-stripped and spatially registered multimodal MR images (T_1 , T_{1c} , T_2 , and T_2 -Flair) with a voxel size of $1mm$ in every direction. Tumors are of different shape, size and location in each data set. The MRI scans of the benchmark data differ significantly in quality parameters such as contrast, noise, and blur.

For an evaluation of our segmentation results, we measure Dice score, sensitivity, and precision in terms of positive predictive value.¹⁹ The results are displayed in Tab. 1. The first label *Complete* includes edema, necrosis and both enhancing and non-enhancing tumor. The second label *Tumor Core* equals the complete tumor but

Table 1: BraTS14 evaluation for high- (HGG) and low-grade glioma (LGG).

Region	Dice Score	Sensitivity	Precision
Complete (HGG)	0.83 ± 0.15	0.87 ± 0.12	0.81 ± 0.14
Tumor Core (HGG)	0.75 ± 0.10	0.81 ± 0.08	0.69 ± 0.15
Active (HGG)	0.69 ± 0.13	0.76 ± 0.11	0.63 ± 0.21
Complete (LGG)	0.78 ± 0.17	0.75 ± 0.18	0.86 ± 0.17
Tumor Core (LGG)	0.54 ± 0.14	0.81 ± 0.15	0.40 ± 0.19
Active (LGG)	0.45 ± 0.23	0.77 ± 0.22	0.34 ± 0.21

Table 2: BraTS13 evaluation for high-grade glioma (Dice scores for complete tumor).

Method	Dice Score	Comp. Time	Architecture
Tustinson ⁶	0.87	100 min	Cluster
Our approach	0.85	0.3 min	CPU/GPU
Zhao ¹⁶	0.84	20 min	CPU
Cordier ¹⁷	0.84	20 min	Cluster
Reza ¹⁸	0.83	90 min	CPU

without the edema. The third label *Active* contains enhancing and non-enhancing tumor. The results show a good segmentation of the complete tumor, both for high-grade and low-grade tumors. The tumor fine-structure is also segmented well for high-grade tumors and suffers from precision for low-grade tumors. All these observations are comparable to other state-of-the-art methods.

Our method, however, differs from many other approaches by its simplicity and its computational efficiency: Regarding simplicity, note that it is basically not more than a consequent application of the Mumford-Shah cartoon model with automatically varying parameters. No training phase and no user interaction are required. With respect to computational efficiency, we strongly benefit from the primal-dual algorithm of Strelakowski et al. To exploit its inherent parallelism, we have implemented the 3D and multimodal minimization of the Mumford-Shah cartoon model in CUDA. The average computation time of our approach per patient is 130.9 s for low-grade gliomas and 137.03 s for high-grade gliomas. We conducted our experiments on a desktop PC with 3.4 GHz Intel Core i7 CPU, 16GB RAM, and a NVidia GeForce GTX 970 graphics card with 4GB RAM and 1664 CUDA cores. As can be seen in Tab. 2, our computation times belong to the fastest approaches in the literature that use comparable hardware and achieve results of similar quality.

We show some exemplary results of our method in Fig. 3. Intuitively one would expect that the segmentation performance would deteriorate when the tumor is very similar to the background. Remarkably, this is not the case: We can determine exact tumor boundaries for low-contrast regions, even when the tumor is very small. Apart from that, our method is also able to identify subcomponents properly.

4. SUMMARY AND CONCLUSIONS

With our paper we have enriched the tool box of algorithms for segmenting brain tumors in MRI data by an approach that differs from existing methods in a number of aspects: First of all, it is fully automatic, but does not require any training phase. Secondly, it is simple due to its consequent use of the Mumford-Shah cartoon model in combination with an intuitive confidence refinement. Thirdly, its computational efficiency benefits from very recent algorithmic advances by means of primal-dual methods and their efficient parallel implementation

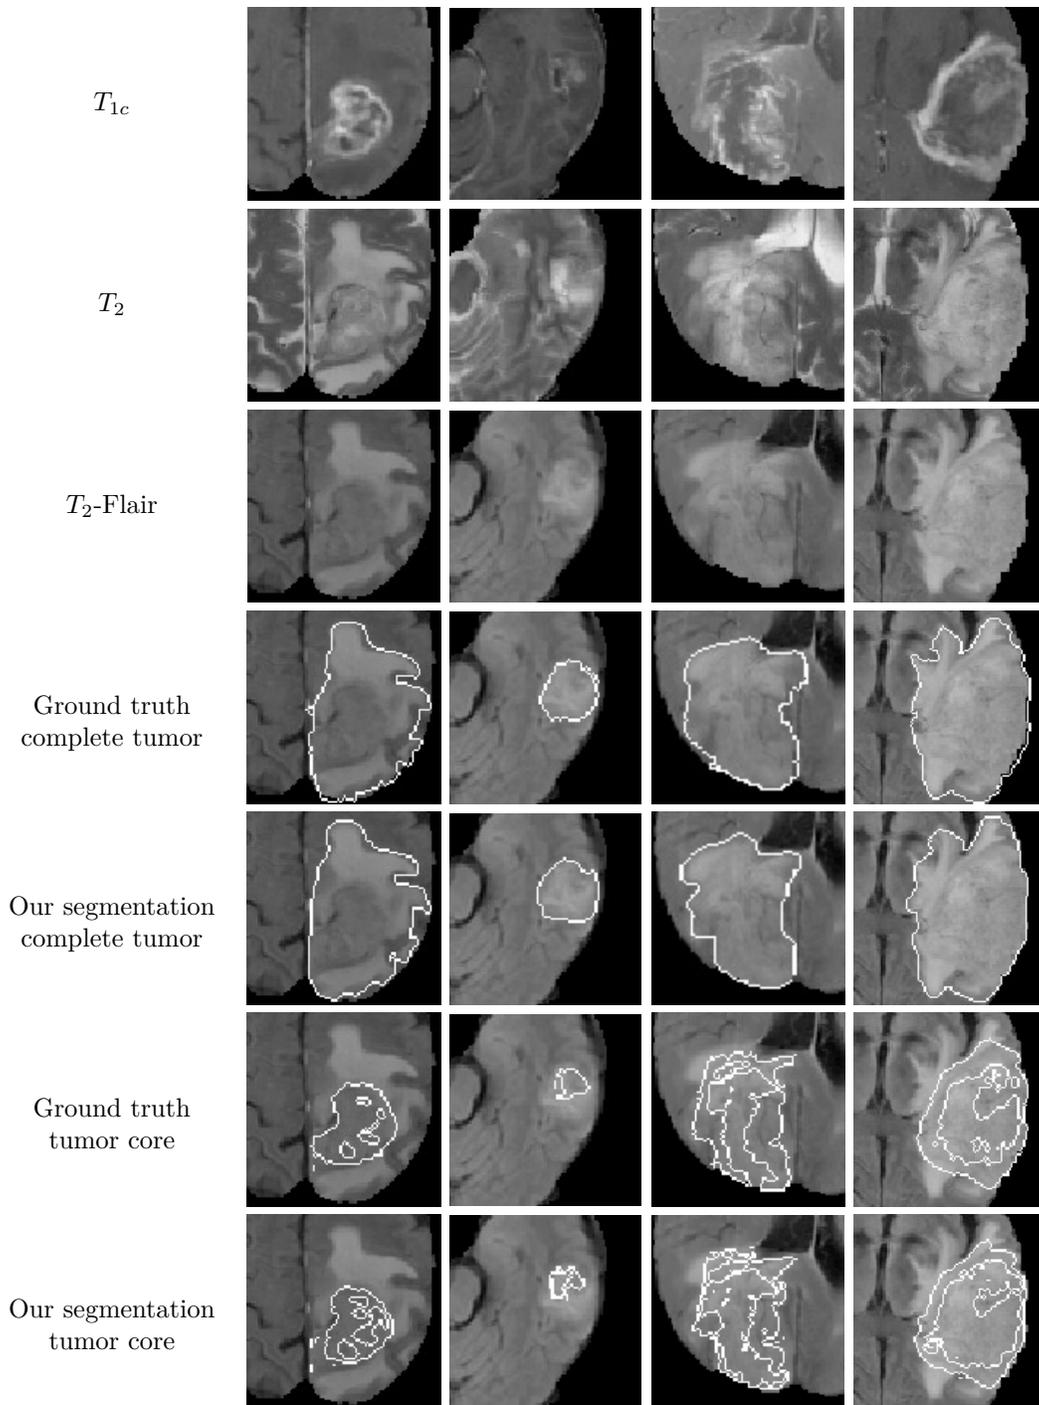


Figure 3: Exemplary results for high grade brain tumors.

on low-cost GPUs. Last but not least, we have shown that these advantages are possible without compromising segmentation quality.

So far we have focused on the important but specific problem of segmenting brain tumors. However, it should be noted that our approach is fairly generic and does not rely very much on particular properties of glioma. Thus,

it is likely that it can also be adapted without too many modifications to other tumor segmentation problems, e.g. kidney tumors. This is a topic of our ongoing research.

ACKNOWLEDGMENTS

This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement No 600841.

REFERENCES

- [1] Marsh, J. C., Goldfarb, J., Shafman, T. D., and Diaz, A. Z., "Current status of immunotherapy and gene therapy for high-grade gliomas," *Cancer Control* **20**, 43–48 (Jan. 2013).
- [2] Wen, P. Y., Macdonald, D. R., Reardon, D. A., Cloughesy, T. F., Sorensen, A. G., Galanis, E., DeGroot, J., Wick, W., Gilbert, M. R., Lassman, A. B., et al., "Updated response assessment criteria for high-grade gliomas: response assessment in neuro-oncology working group," *Journal of Clinical Oncology* **28**, 1963–1972 (Mar. 2010).
- [3] Mazzara, G. P., Velthuizen, R. P., Pearlman, J. L., Greenberg, H. M., and Wagner, H., "Brain tumor target volume determination for radiation treatment planning through automated MRI segmentation," *International Journal of Radiation Oncology - Biology - Physics* **59**, 300–312 (May 2004).
- [4] Bauer, S., Wiest, R., Nolte, L.-P., and Reyes, M., "A survey of MRI-based medical image analysis for brain tumor studies," *Physics in Medicine and Biology* **58**(13), R97–R129 (2013).
- [5] Menze, B., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., et al., "The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS)," *IEEE Transactions on Medical Imaging* (2014).
- [6] Tustison, N., Wintermark, M., Durst, C., and Avants, B., "Ants and Arboles," in [*Proc. MICCAI BraTS Workshop*], MICCAI Society, Nagoya, Japan (Sept. 2013).
- [7] Urban, G., Bendszus, M., Hamprecht, F. A., and Kleesiek, J., "Multi-modal brain tumor segmentation using deep convolutional neural networks," in [*MICCAI BraTS Challenge. Proceedings*], 31–35 (2014).
- [8] Chan, T. F. and Vese, L. A., "Active contours without edges," *IEEE Transactions on Image Processing* **10**, 266–277 (Feb. 2001).
- [9] Mumford, D. and Shah, J., "Boundary detection by minimizing functionals, I," in [*Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*], 22–26, IEEE Computer Society Press, San Francisco, CA (June 1985).
- [10] Mumford, D. and Shah, J., "Optimal approximations by piecewise smooth functions and associated variational problems," *Communications on Pure and Applied Mathematics* **42**(5), 577–685 (1989).
- [11] Strelakovski, E. and Cremers, D., "Real-time minimization of the piecewise smooth mumford-shah functional," in [*Computer Vision – ECCV 2014*], Fleet, D., Pajdla, T., Schiele, B., and Tuytelaars, T., eds., *Lecture Notes in Computer Science* **8692**, 127–141, Springer (2014).
- [12] Tustison, N. J., Avants, B. B., Cook, P. A., Zheng, Y., Egan, A., Yushkevich, P. A., and Gee, J. C., "N4ITK: improved N3 bias correction," *IEEE Transactions on Medical Imaging* **29**, 1310–1320 (June 2010).
- [13] Morel, J.-M. and Solimini, S., [*Variational Methods in Image Segmentation*], vol. 12 of *Progress in Nonlinear Differential Equations and Their Applications*, Birkhäuser, Basel (1994).
- [14] Otsu, N., "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man and Cybernetics* **9**, 62–66 (Jan. 1979).
- [15] Dice, L. R., "Measures of the amount of ecologic association between species," *Ecology* **26**, 297–302 (July 1945).
- [16] Zhao, L., Sarikaya, D., and Corso, J. J., "Automatic brain tumor segmentation with mrf on supervoxels," in [*Multimodal Brain Tumor Segmentation*], 51 (2013).
- [17] Cordier, N., Menze, B., Delingette, H., and Ayache, N., "Patch-based segmentation of brain tissues," in [*MICCAI Challenge on Multimodal Brain Tumor Segmentation*], 6–17, IEEE (2013).
- [18] Reza, S. and Iftekharuddin, K., "Multi-class abnormal brain tissue segmentation using texture features," in [*MICCAI Challenge on Multimodal Brain Tumor Segmentation*], 38–42, IEEE (2013).
- [19] Fawcett, T., "An introduction to ROC analysis," *Pattern Recognition Letters* **27**, 861–874 (June 2006).