

A Discriminative Model-Constrained Graph Cuts Approach to Fully Automated Pediatric Brain Tumor Segmentation in 3-D MRI

Michael Wels¹, Gustavo Carneiro², Alexander Ablas³, Martin Huber⁴, Dorin Comaniciu², and Joachim Hornegger¹

¹ Chair of Pattern Recognition, University Erlangen-Nuremberg, Germany
michael.wels@informatik.uni-erlangen.de

² Siemens Corporate Research, IDS, Princeton, NJ, USA

³ Institute of Radiology, University Medical Center, Erlangen, Germany

⁴ Siemens, CT SE SCR 2, Erlangen, Germany

Abstract. In this paper we present a fully automated approach to the segmentation of pediatric brain tumors in multi-spectral 3-D magnetic resonance images. It is a top-down segmentation approach based on a Markov random field (MRF) model that combines probabilistic boosting trees (PBT) and lower-level segmentation via graph cuts. The PBT algorithm provides a strong discriminative observation model that classifies tumor appearance while a spatial prior takes into account the pair-wise homogeneity in terms of classification labels and multi-spectral voxel intensities. The discriminative model relies not only on observed local intensities but also on surrounding context for detecting candidate regions for pathology. A mathematically sound formulation for integrating the two approaches into a unified statistical framework is given. The proposed method is applied to the challenging task of detection and delineation of pediatric brain tumors. This segmentation task is characterized by a high non-uniformity of both the pathology and the surrounding non-pathologic brain tissue. A quantitative evaluation illustrates the robustness of the proposed method. Despite dealing with more complicated cases of *pediatric brain tumors* the results obtained are similar to those reported for current state-of-the-art approaches to 3-D MR brain tumor segmentation in adult patients. Processing one multi-spectral data set takes 5 minutes on average including pre-processing and does not require any user interaction.

1 Introduction

Detection and delineation of pathology, such as cancerous tissue, within multi-spectral brain magnetic resonance (MR) volume sequences is an important problem in medical image analysis. For example, a precise and reliable segmentation of brain tumors present in the childlike brain is regarded critical when aiming for the automatic extraction of diagnostically relevant quantitative or more abstract findings. This may include the volume of the tumor or its relative location. Once these findings are obtained they can be used both for guiding computer-aided diagnosis and therapy planning as well as for traditional decision making. However, the manual labeling of volumetric data is usually time consuming, which has the

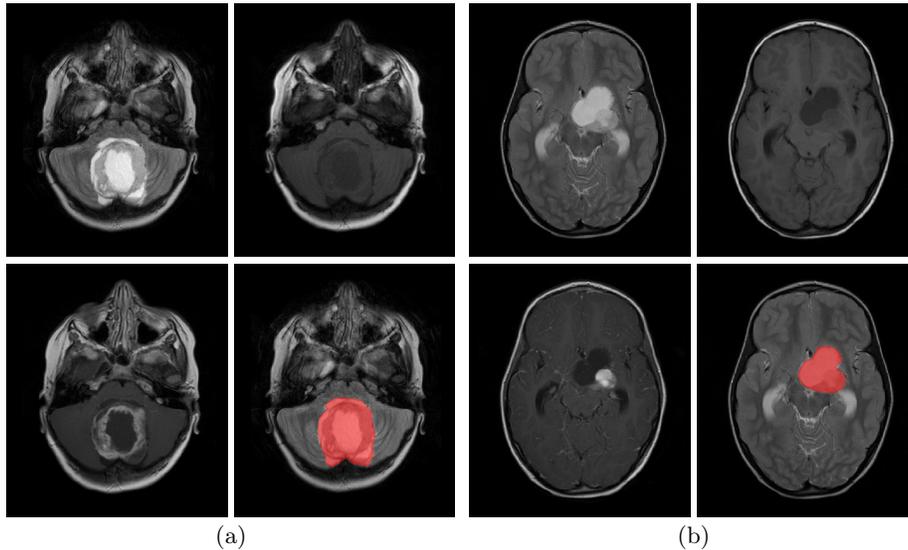


Fig. 1. Two different cases of pediatric brain tumors exhibiting heterogeneous shape and appearance. Columns (a) and (b) show axial slices of the typically acquired pulse sequences (row-wise from left to right: T2-weighted, T1-weighted, and T1-weighted after contrast enhancement) and the expert annotated ground-truth overlaid to the T2-weighted pulse sequence.

potential to delay clinical workflow, such that there is a need for fully automatic segmentation tools in this particular context. Furthermore, manual annotations may vary significantly among experts as a result of individual experience and interpretation.

As multi-spectral 3-D magnetic resonance imaging (MRI) is the method of choice for the examination of neurological pathology such as brain cancer in pediatric patients, automatic approaches first have to be capable of dealing with the characteristic artifacts of this imaging modality: Rician noise, partial volume effects, and intra-/inter-scan intensity inhomogeneities. Second and more importantly, they have to be robust enough to handle the heterogeneous shape and appearance of pediatric brain tumors in different patients (see Fig. 1).

In this paper, we propose a fully automatic solution based on a novel top-down segmentation scheme that uses a statistical model of the pathology appearance as a constraint for a sub-sequent optimization problem. The statistical model is provided by a machine learning technique that is able to work with high-dimensional feature vectors allowing to encode characteristic voxel contexts. The optimization problem itself is stated as a search for an MAP estimate of the most-likely binary image segmentation, which permits efficient computation of a solution by means of a max-flow/min-cut optimization procedure.

Approaches in the field of MR brain tumor segmentation rarely rely on pure data-driven approaches due to the complexity in terms of tumor shape and appearance of the segmentation task. The vast majority of methods make use

of domain knowledge in varying representations in combination with low-level imaging techniques. Fletcher-Heath et al. [1] use unsupervised fuzzy clustering followed by 3-D connected components with an intermediate step incorporating knowledge about the usual distribution of cerebral spinal fluid and location of the ventricular system. Gering et al. [2] use trained models for intensity distributions of non-pathologic brain tissue to detect model outliers on the voxel level that are considered tumor voxels in a multi-layer Markov random field framework. In a similar manner Prastawa et al. [3] detect outliers based on refined intensity distributions for healthy brain tissue initially derived from a registered probabilistic atlas, which introduces structural domain knowledge. Registration is also used in combination with voxel intensities in the adaptive template-moderated classification algorithm by [4]. More recent approaches try to enrich low level segmentation techniques, like level set evolution [5] or hierarchical clustering [6], by using supervised machine learning on higher dimensional feature sets associated with each image voxel. These feature sets are capable of representing a more general variety of domain knowledge on different levels of abstraction. In a similar manner we make use of the recently proposed technique of probabilistic boosting trees (PBT) [7] for supervised learning, which has proven its robustness and its capability for efficient training and classification in numerous applications [8, 9]. The probability estimates provided by PBT are then used to constrain the highly efficient computation of minimum cuts [10] for image segmentation based on a Markov random field (MRF) prior model. It takes into account both coherence of classification labels as well as intensity similarities within voxel neighborhoods. To the best of our knowledge, this is the first paper giving an integrated formulation for combining PBT classification and computation of minimum cuts. Opposed to [5, 6] there is no involvement of a time consuming bias field correction step in data pre-processing. In the case of [6] this seems to be done by FAST [11], which relies on an HMRF-EM segmentation approach. In the presence of abnormal tissue types this requires the determination of the number of different intensity regions expected within each scan. Furthermore, the inherent low level segmentation might bias the final segmentation result. In contrast we build discriminative models, i.e. PBTs, whose generalization capabilities are strong enough to implicitly handle those intra-patient intensity non-uniformities. Moreover, we apply our method to the more complicated task of segmenting pediatric brain tumors where not only pathology underlies significant variation in shape and appearance but also the non-pathological “background”, which is caused by ongoing myelination of white matter during maturation [12].

2 Discriminative Model-Constrained Graph Cuts Segmentation

Our segmentation method relies on the integrated formulation of an objective function that is subject to optimization via the efficient graph cuts algorithm [10]. In the following we derive this objective function from the general MAP framework for image segmentation.

In general, the problem of segmenting an image can be stated as the search for an MAP estimate of the most likely class labels given appropriate prior and observation models in terms of probability density functions. Let $\mathcal{S} = \{1, 2, \dots, N\}$ be a set of indices to image voxels. At each index $s \in \mathcal{S}$ there are two random

variables: $y_s \in \mathcal{Y} = \{-1, 1\}$ and $\mathbf{x}_s \in \mathcal{X}$. The former denotes the unobservable binary segmentation of voxel s into fore- and background, whereas the latter states the observable vector of associated features, that are assumed to be causally linked to the underlying class labels y by a unified observation model defined by a probability density function $P(\mathbf{x}|y)$. The emergence of the class labels themselves is described by a prior model $P(y)$. The segmentation task at hand can now be stated as the search for an MAP estimate

$$\mathbf{Y}^* = \underset{\mathbf{Y}}{\operatorname{argmax}} p(\mathbf{Y}|\mathbf{X}) \quad (1)$$

where $p(\mathbf{Y}|\mathbf{X})$ is the joint posterior probability over the image domain. Using the Bayes rule, and assuming a uniform distribution $P(\mathbf{X})$, we have:

$$\mathbf{Y}^* = \underset{\mathbf{Y}}{\operatorname{argmax}} \ln p(\mathbf{X}|\mathbf{Y}) + \ln p(\mathbf{Y}). \quad (2)$$

To concretize this optimization problem a region-specific probability term and an appropriate prior need to be identified. In our method $P(\mathbf{X}|\mathbf{Y})$ is provided by a PBT classifier. The machine learning technique of PBT recursively groups boosted ensembles of weak classifiers to a tree structure during learning from expert annotated data. Training a PBT resembles inducing a multivariate binary regression tree as the final strong classifier

$$H(\mathbf{x}) = \sum_{t=1}^T h_t(\mathbf{x}) \quad (3)$$

generated within each inner node for a feature vector \mathbf{x} through a combination of real-valued contributions $h_t(\mathbf{x})$ of $T \in \mathbb{N}$ weak classifiers asymptotically approaches the additive logistic regression model [13]

$$H(\mathbf{x}) \approx \frac{1}{2} \ln \frac{p(y = +1|\mathbf{x})}{p(y = -1|\mathbf{x})}. \quad (4)$$

Accordingly, at each inner node v of the resulting PBT there are current approximations of the posterior probability distribution $\tilde{P}_v(y|\mathbf{x})$. During classification those values are used to guide tree traversing and combined propagation of posteriors in order to get a final approximation $\tilde{P}(y|\mathbf{x})$ of the true posterior probabilities $P(y|\mathbf{x})$ at the tree's root node.

As mentioned above, assuming \mathbf{X} to be distributed uniformly, and also to be independently and identically distributed, we have $P(\mathbf{x}|y) \propto P(y|\mathbf{x})$ and therefore $p(\mathbf{X}|\mathbf{Y}) \approx \prod_{s \in \mathcal{S}} p(y_s|\mathbf{x}_s)$. We furthermore assume conditional independence of the observed features \mathbf{x} for PBT classification and of the multi-spectral intensities \mathbf{i} such that we can say $p(\mathbf{X}|\mathbf{Y}) \approx \prod_{s \in \mathcal{S}} \tilde{p}(y_s|\mathbf{x}_s)p(\mathbf{i}_s|y_s)$ in (2) with approximations $\tilde{p}(y_s|\mathbf{x}_s)$ from the classifier and $P(\mathbf{i}|y)$ being the foreground/background intensity distribution. The latter can be estimated non-parametrically via histograms from the initial hard classification for $\tilde{p}(y_s|\mathbf{x}_s) \geq 0.5$ obtained by the PBT.

The feature vectors \mathbf{x} used for PBT classification consist of individual multi-spectral intensities, inter-spectra intensity gradients, and 2-D Haar-like features

[14, 15] computed on an intra-axial 2-D context surrounding the voxel of interest. The Haar-like features are derived from a subset of the extended set of Haar-like feature prototypes by [16] and are represented only implicitly in memory by so-called (Rotated) Integral Images. This allows for fast re-computation of the features with respect to a given voxel when they are actually assessed. We decided on 2-D Haar-like features in contrast to the full set of 3-D Haars because of their lower computational cost and memory requirements. Also, as we intend to capture a discriminative representation of the full context, and not only of local edge characteristics at the central voxel, Haar-like feature values are computed according to the given prototypes on every valid origin and scale within the chosen voxel context.

For the prior distribution we assume a Markov random field prior model

$$P(\mathbf{Y}) \propto \exp(-U(\mathbf{Y}|\frac{1}{\lambda})) \quad (5)$$

formulated, according to the Hammersly-Clifford Theorem, as a Gibbs distribution with energy function

$$U(\mathbf{Y}|\frac{1}{\lambda}) = \frac{1}{\lambda} \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{N}_s} V_{st}(y_s, y_t) \quad (6)$$

where $\frac{1}{\lambda}$ with $\lambda \in]0.0; +\infty[$ controls the relative influence of the spatial prior over the observation model and \mathcal{N}_s states the neighborhood of voxel s . Inspired by [17] the interaction potentials are

$$V_{st}(y_s, y_t) = \exp\left(-\frac{1}{2} \sum_{l=1}^L \frac{(i_{sl} - i_{tl})^2}{\sigma_l^2}\right) \cdot \frac{\delta(y_s, y_t)}{dist(s, t)} \quad (7)$$

where vectors \mathbf{i} denote the observed intensities taken from L aligned input pulse sequences and

$$\delta(y_s, y_t) = \begin{cases} 1 & \text{if } y_s \neq y_t \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

The function $dist(s, t)$ denotes the physical distance between voxels s and t in patient space, which varies when working on image volumes with anisotropic voxel spacing. The model emphasizes homogeneous segmentations among neighboring voxels but weights penalties for heterogeneity according to intensity similarities of the voxels involved.

With the equality

$$\begin{aligned} \mathbf{Y}^* &= \operatorname{argmax}_{\mathbf{Y}} \sum_{s \in \mathcal{S}} \ln \tilde{p}(y_s | \mathbf{x}_s) p(\mathbf{i}_s | y_s) - \frac{1}{\lambda} \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{N}_s} V_{st}(y_s, y_t) = \\ & \operatorname{argmin}_{\mathbf{Y}} \lambda \cdot \left(\sum_{s \in \mathcal{S}} -\ln \tilde{p}(y_s | \mathbf{x}_s) p(\mathbf{i}_s | y_s) \right) + \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{N}_s} V_{st}(y_s, y_t) \end{aligned} \quad (9)$$

the initial maximization problem can be transformed into a minimization problem that is in a suitable form for optimization by the graph cuts algorithm [10].

Note that the reciprocal of the regularization parameter in (6) can equivalently be used to weight the influence of the observation model over the prior model. Details on how to set up the graph and how to define the edge capacities for the associated max-flow/min-cut problem instance based on (9) can be found in [17].

3 Experimental Setting and Results

For quantitative evaluation of the proposed methods there were six multi-spectral expert annotated data sets available. Each scan consists of three 3-D images acquired at different pulse sequences (T2-weighted, T1-weighted, and T1-weighted after contrast enhancement). The resolution is $512 \times 512 \times 20$ with a voxel spacing of $0.449219 \times 0.449219 \times 6.0$. Where necessary due to patient movement during image acquisition the pulse sequences were co-aligned by means of the MedINRIA affine registration tool (www-sop.inria.fr/asclepios/software/MedINRIA). All the sequences were further pre-processed by the following pipeline: skull stripping by the Brain Extraction Tool (BET) [18], gradient anisotropic diffusion filtering (www.itk.org), and inter-scan intensity standardization by Dynamic Histogram Warping (DHW) [19]. Note that all of the pre-processing steps involved, including co-alignment, can be performed fully automatically without any user interaction.

The PBT voxel classifiers built were restricted to a maximum depth of 10 with 10 weak classifiers per tree node. The 2-D voxel context considered was of size 11×11 on volumes down-sampled to a voxel spacing of $2.0\text{mm} \times 2.0\text{mm} \times 6.0\text{mm}$. The graph cuts optimization is carried out on the original image resolution with \mathcal{N}_s defined to be a standard 6-neighborhood on the 3-D image lattice. The standard deviation σ for the interaction potentials in (7) was estimated offline as “camera noise” within manually delineated homogeneous regions throughout the patient volumes. A leave-one-out cross validation on the patient data sets and their accompanying PBT models yielded best average segmentation scores in terms of the Jaccard coefficient ($TP/(TP + FP + FN)$) for $\lambda \in]0.1, 0.5[$ such that finally $\lambda = 0.2$ was chosen for computing the results depicted in Fig. 2. In order to remove small regions of false positive voxels only the largest connected component of the graph cuts result is considered to be the final segmentation. With Jaccard coefficients of 0.78 ± 0.17 they are better than those published by [5] and, except for one case, in a similar range as those of [6] who all work with adult patient data sets and partly on four pulse sequences. However, comparability of results is limited because of different characteristics between the data sets used by the mentioned scientists, e.g., pediatric patients versus adult patients, presence of necrotic tissue within the tumors, restriction to a certain histological type of tumor, etc.

It takes about 1–2 minutes to process one of the MRI volumes in a non-optimized C++ implementation of our segmentation method on a Fujitsu Siemens Computers notebook equipped with an Intel Pentium M 2.0 GHz processor and 2 GB of memory. With the same hardware as above training one classifier takes about 4 hours. Preprocessing the images takes about 3 minutes so a total amount of 5 minutes is needed for the processing of one patient data set.

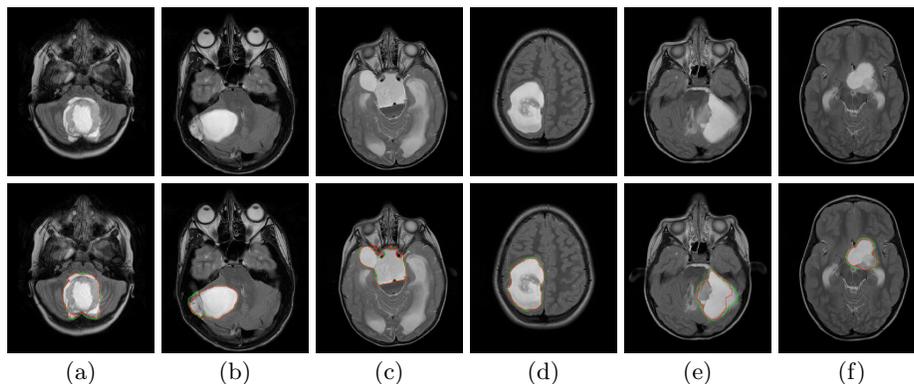


Fig. 2. Segmentation results obtained by leave-one-patient-out cross validation. The first row shows selected slices of the T2-weighted pulse sequences of the six available patient data sets. The second row shows the associated segmentation results (red) and the ground-truth segmentation (green) overlaid on the T2-weighted pulse sequence. Please view in color.

4 Conclusions

The contribution of this paper is two-fold: starting from the well-known MAP framework for image segmentation we have derived a constrained minimization problem suitable for max-flow/min-cut optimization that incorporates an observation model provided by a discriminative PBT classifier into the process of segmentation. Secondly, we successfully applied the method to the difficult problem of fully automatic pediatric brain tumor segmentation in multi-spectral 3-D MRI. The experimental results obtained are comparable to those recently published for fully automatic segmentation in adult patients.

Since tumor is a rather general concept in medicine, limitations of the proposed approach might become apparent as soon as unforeseen pathologic tissue types that could not adequately be captured by the discriminative model appear in previously unseen patients. Especially secondary tumors might be composed of an enormous variety of tissue types depending on the primary tumor site. Future work will therefore investigate possibilities to handle this issue. The goal is to detect, to segment, and to identify most types of pathological tissue that occur within pediatric brain tumors.

References

1. Fletcher-Heath, L., Hall, L.O., Goldgof, D.B., Murtagh, F.R.: Automatic segmentation of non-enhancing brain tumors in magnetic resonance images. *Artif. Intell. Med.* **21**(1-3) (2001) 43-63
2. Gering, D.T., Grimson, W.E.L., Kikinis, R.: Recognizing deviations from normalcy for brain tumor segmentation. In: *Int. Conf. Med. Image Comput. and Comp.-Assist. Interv.*, Tokyo, Japan. (Sept. 2002) 388-395
3. Prastawa, M., Bullitt, E., Ho, S., Gerig, G.: A brain tumor segmentation framework based on outlier detection. *Med. Image Anal.* **8**(3) (Sept. 2004) 275-283

4. Kaus, M.R., Warfield, S.K., Nabavi, A., Black, P.M., Jolesz, F.A., Kikinis, R.: Automated segmentation of MR images of brain tumors. *Radiology* **218** (2001) 586–591
5. Cobzas, D., Birkbeck, N., Schmidt, M., Jagersand, M., Murtha, A.: 3D variational brain tumor segmentation using a high dimensional feature set. In: Proceedings of the Mathematical Methods in Biomedical Image Analysis (MMBIA) Workshop (11th IEEE Int. Conf. Comp. Vis., Rio de Janeiro, Brazil). (Oct. 2007) 1–8
6. Corso, J.J., Yuille, A., Sicotte, N.L., Toga, A.: Detection and segmentation of pathological structures by the extended graph-shifts algorithm. In: Int. Conf. Med. Image Comput. and Comp.-Assist. Interv., Brisbane, Australia. Volume 4791. (Oct. 2007) 985–993
7. Tu, Z.: Probabilistic boosting-tree: Learning discriminative models for classification, recognition, and clustering. In: IEEE Int. Conf. Comp. Vis., Beijing, China. (Oct. 2005) 1589–1596
8. Tu, Z., Zhou, X., Comaniciu, D., Bogoni, L.: A learning based approach for 3D segmentation and colon detagging. In: Europ. Conf. Comp. Vis., Graz, Austria. (May 2006) 436–448
9. Carneiro, G., Georgescu, B., Good, S., Comaniciu, D.: Automatic fetal measurements in ultrasound using constrained probabilistic boosting tree. In: Int. Conf. Med. Image Comput. and Comp.-Assist. Interv., Brisbane, Australia. (Oct. 2007)
10. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE T. Pattern Anal.* **26**(9) (Sept. 2004) 1124–1137
11. Zhang, Y., Brady, M., Smith, S.: Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE T. Med. Imag.* **20**(1) (Jan. 2001) 45–57
12. Murgasova, M., Dyet, L., Edwards, A.D., Rutherford, M.A., Hajnal, J.V., Rueckert, D.: Segmentation of brain MRI in young children. In Larsen, R., Nielsen, M., Sporring, J., eds.: Int. Conf. Med. Image Comput. and Comp.-Assist. Interv., Copenhagen, Denmark. Volume 4190 of Lect. Notes in Comp. Sci., Springer (Oct. 2006) 687–694
13. Friedman, J., Hastie, T., Tibshirani, R.: Additive logistic regression: a statistical view of boosting. Stanford University, Department of Statistics, Technical Report (Aug. 1998)
14. Oren, M., Papageorgiou, C., Sinha, P., Osuna, E., Poggio, T.: Pedestrian detection using wavelet templates. In: IEEE Comp. Soc. Conf. Comp. Vis. and Pat. Recog., San Juan, Puerto Rico. (June 1997) 193–199
15. Viola, P., Jones, M.: Robust real-time object detection. In: 2nd International Workshop on Statistical Learning and Computational Theories of Vision, Vancouver, Canada. (July 2001) 1–25
16. Lienhart, R., Kuranov, A., Pisarevsky, V.: Empirical analysis of detection cascades of boosted classifiers for rapid object detection. In Michaelis, B., Krell, G., eds.: DAGM-Symposium. Volume 2781 of Lect. Notes in Comp. Sci., Springer (Sept. 2003) 297–304
17. Boykov, Y., Funka-Lea, G.: Graph cuts and efficient n-d image segmentation. *Int. J. Comput. Vision* **70**(2) (2006) 109–131
18. Smith, S.M.: Fast robust automated brain extraction. *Hum. Brain Mapp.* **17**(3) (2002) 143–155
19. Cox, I.J., Hingorani, S.L.: Dynamic histogram warping of image pairs for constant image brightness. In: IEEE Int. Conf. on Image Proc., Vol. II, Washington, D.C., USA. (Oct. 1995) 366–369