

Recognizing Deviations from Normalcy for Brain Tumor Segmentation

David T. Gering¹, W. Eric L. Grimson¹, R. Kikinis²

¹ Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, 02139 USA

{gering, welg}@ai.mit.edu

<http://www.ai.mit.edu/people/gering>

² Brigham and Women's Hospital, Harvard Medical School, Boston MA, USA
kikinis@bwh.harvard.edu

Abstract. A framework is proposed for the segmentation of brain tumors from MRI. Instead of training on pathology, the proposed method trains exclusively on healthy tissue. The algorithm attempts to recognize deviations from normalcy in order to compute a fitness map over the image associated with the presence of pathology. The resulting fitness map may then be used by conventional image segmentation techniques for honing in on boundary delineation. Such an approach is applicable to structures that are too irregular, in both shape and texture, to permit construction of comprehensive training sets. The technique is an extension of EM segmentation that considers information on five layers: voxel intensities, neighborhood coherence, intra-structure properties, inter-structure relationships, and user input. Information flows between the layers via multi-level Markov random fields and Bayesian classification. A simple instantiation of the framework has been implemented to perform preliminary experiments on synthetic and MRI data.

1 Introduction

The literature is rich with techniques for segmenting healthy brains – a task simplified by the predictable appearance, size, and shape of healthy structures. Many of these methods fail in the presence of pathology – the very focus of segmentation for image-guided surgery [1]. Furthermore, the techniques that are intended for tumors leave significant room for increased automation and applicability. Specifically, we consider the task of segmenting large brain tumors such as gliomas, meningiomas, astrocytomas, glioblastoma multiforme, cavernomas, and Arteriovenous Malformations (AVM). In practice, segmentation of this class of tumors continues to rely on manual tracing and low-level computer vision tools such as thresholds, morphological operations, and connective component analysis. Automatic techniques tend to be either region- or contour-based.

Region-based methods usually reduce operator interaction by automating some aspects of applying the low-level operations [2]. Threshold selection can be assisted through histogram analysis, and logic can be applied to the application of low-level

vision techniques through a set of rules to form a knowledge-based system [3]. Since statistical classification alone may not allow differentiation between non-enhancing tumor and normal tissue, anatomic information derived from a digital atlas has been used to identify normal anatomic structures. Of these approaches, the most successful has been the iteration of statistical classification and template matching as developed in [4,5,6]. However, the use of morphological operations has the drawback of making assumptions about the radius parameter that are both application-dependent (anatomy) and scan-dependent (voxel size). Such operations destroy fine details and commit to irreversible decisions at too low of a level to benefit from all available information – thus violating Marr’s principle of least commitment.

Contour-based methods evolve a curve based on internal forces (e.g.: curvature) and external forces (e.g.: image gradients) to delineate the boundary of a tumor. Since they experience similar drawbacks as the region-based approaches, they tend to apply only to tumors that are easily separable from their surroundings [7,8]. Level-set based curve evolution [9, 10] has the advantage over region-based approaches in that the connectivity constraint is imposed implicitly rather than through morphological operations. However, 3D level-sets find limited use in medical practice due to their reliance on the operator to set the sensitive parameters that govern the evolution’s stopping criteria. The more heterogeneous the tumor, the more user interaction is required.

Both region- and contour-based segmentation methods have largely ignored the bias field, or patient-specific, signal inhomogeneities present in MRI. The bias field is slowly varying, and therefore its computation from the regions of healthy tissue could be extrapolated over tumor tissue to provide some degree of benefit. Methods for segmenting healthy brains have incorporated the EM algorithm [11] to simultaneously arrive at both a bias field and segmentation into healthy tissue classes [12]. There have been several extensions, such as collecting all non-brain tissue into a single class [13], handling thermal noise with a Markov random field [14], using a mean-field solution to the Markov random field [15], incorporating geometric constraints [15], using a digital brain atlas as a spatially-varying prior [16], automating the determination of the tissue class parameters [17], and identifying MS lesions as hyper-intense outliers from white matter [18].

In contrast to existing methods for tumor segmentation, the hypothesis underlying our work is that we can segment brain tumors by focusing not on what typically represents tumor, but on what typically represents healthy tissue. Our method extends EM-based segmentation to compute a fitness map over the image to be associated with the probability of pathology. In this paper, we present our proposed framework and the preliminary results that we achieved with a simple implementation. We hope that more detailed, future implementations of this framework will be able to address many of the drawbacks to the existing region- and contour-based methods.

2 Method

Inherent ambiguity necessitates the incorporation of contextual information into the brain segmentation process. Consider the example of non-enhancing tumor tissue that

garners the same intensity classification of healthy gray matter, but is too thick to be gray matter. An algorithm’s low-level computer vision techniques could first classify the tissue as gray matter, and a higher-level stage – through its broader understanding of context -- could correct the classifications of the first-pass. This example motivates the introduction of hierarchical context into the segmentation process. A voxel’s classification could be considered on several levels: the voxel itself, the voxel’s immediate (Markov) neighborhood, the voxel’s region (entire connected structure), the global setting (position of the voxel’s structure relative to other structures), and user guidance. Just as a voxel-wise classification must be computed prior to a neighborhood-wise refinement, a voxel’s region must be classified before features regarding the size and shape (or other intrinsic properties) of that region can be computed.

One way to approach the above example would be to employ a combination of morphological operations and a high-level expert system (such as [3]) to simultaneously switch the classification of every voxel in the mistaken connected mass from gray matter to tumor. However, in keeping with our goal of adhering to the principle of least commitment, we propose an alternative approach where voxels toward the center of the mass could be first classified as tumor based on their unusually high distance from their structure’s boundary. This tumor classification would subsequently flow outward throughout the mass over several iterations in a *probabilistic flow*. The flow is driven by our introduction of novel multi-layer Markov random fields. A given voxel would change its high-level classification in the evolving presence of tumor if the attributes of lower-level layers shared strong similarities.

Table 1. The Layered Vision framework features no decisions made by certain layers that permanently (and perhaps adversely) affect other layers. Information flows between the layers (bidirectionally depending on implementation details) while converging toward a final solution

#	Layer	Definition	Computation
5	User (<i>oracle</i>)	Spatially specific points clicked on by the user either at initialization, or on the fly between iterations as a corrective action.	Quickly and crudely drawn line during initialization.
4	Inter-structure (<i>global</i>)	Relative position of a voxel’s structure to other structures.	Distance from scalp
3	Intra-structure (<i>region</i>)	Relative position of a voxel within its own structure.	Distance from boundary. Multi-layer MRF.
2	Neighborhood (<i>local</i>)	Classification of a voxel’s immediate neighbors.	Mean Field MRF
1	Voxel (<i>point</i>)	Classification based on voxel’s intensity.	EM, ML or MAP

Table 1 combines these observations into our proposed hierarchy of five processing layers. The rightmost column describes each layer’s computation for initial experiments, allowing for more elaborate implementations in the near future. The key to the algorithm’s novelty is the third layer. Observe that some of the existing segmentation algorithms employ the following combinations of layers:

Wells96:	1
Held97:	1, 2
Kapur99:	1, 2, 4
Kaus01:	1, 4, 5 (and 2 by pre-processing)

2.2 Layer 1: EM Segmentation

EM segmentation models the image intensities as visible variables, Y , tissue classifications as hidden variables, Γ , and the bias field as governed by model parameters, β . We would like to choose the parameters that maximize the log likelihood of the data, $\log p(Y, \Gamma|\beta)$, but we do not know this likelihood because Γ 's invisibility renders $p(Y, \Gamma|\beta)$ to be a random variable. Thus, although we cannot maximize it, we can *maximize its expectation*. This results in the following two iterative steps until convergence to a local minimum. In our implementation, we followed [12].

E-Step: Compute the expectation $\sum_{\Gamma} p(\Gamma|Y, \beta) \log p(\Gamma, Y|\beta)$ using the current β .

M-Step: Find new $\beta^{(t+1)}$ to maximize the expectation, assuming $p(\Gamma|Y, \beta^t)$ is correct.

2.2 Layer 2: Spatial Coherence with Markov Random Field

Following [15], the prior knowledge of spatial coherence over a configuration, W , of segmented voxels is modeled with a Gibbs distribution, $P(W)$. The energy function, $U(W)$ is an Ising model generalized to the case of discrete, multi-valued labels [19]:

$$P(W) = \frac{\exp(-U(W))}{\sum_{W'} \exp(-U(W'))} \quad (1)$$

Given M possible label values, let W_s be an M -length binary vector of classification at the voxel indexed by s . Let N_s be the set of voxels in the neighborhood of s (which, in our case, are the 6 closest voxels in 3D). Let the superscripts refer to the layer of processing from Table 1.

$$U^2(W) = V_1(W_s^2) + \sum_{n \in N_s} V_2(W_s^2, W_n^2) \quad (3)$$

Then V_1 is the clique potential of all cliques of size 1. In other words, V_1 encodes our prior knowledge about an isolated voxel prior to viewing the image data (the tissue class prior probability). V_2 is the potential over all cliques of size 2, and represents the tendency of two classified voxels to be neighbors. That tendency is encoded in the $M \times M$ Class Interaction Matrix, J^2 , and is computed from a segmented scan offered as training data.

$$V_1(W_s^2) = - \sum_{s \in S} p(\Gamma) W_s^2 \quad (3)$$

$$V_2(W_s^2) = -\sum_{s \in S} \sum_{n \in N_s} (W_s^2)^T J^2 W_n^2 \quad (4)$$

For tractable computation, our implementation follows the Mean-Field approximation as derived in [15]. Thus, Layer 2 effectively relaxes Layer 1's E-Step weights.

2.3 Layer 3: Region Properties Propagated with a Multi-Level MRF

We now derive our simple implementation of Layer 3 where the feature computed over the output of Layer 2 is the radius to each structure's own boundary. The per-class probability distributions, $p(r|\Gamma)$ are readily computed from a sample segmented scan presented as training data. We perform a Maximum A Posteriori (MAP) classification of the features (just radius at present) computed over Layer 2's output. Recall that the EM algorithm of Layer 1 must compute $p(\Gamma|Y, \beta)$ at each E-Step. It can be shown that $p(\Gamma|Y, \beta, r)$ can be computed with the same update equation except for an extra multiplicative term, $p(r|\Gamma)$. Therefore, the posterior probabilities for the Layer 3 MAP classification are equal to the relaxed weights of Layer 2 multiplied by this new likelihood. That is, the Layer 2 weights provide the spatially varying prior for the Layer 3 MAP classification.

Next, we desire the MAP result (corrections to Layer 2's classifications) to propagate over regions that are homogenous at Layer 2, as demonstrated in Figure 2. We introduce a multi-level Markov random field, and define the Gibb's energy function to encode our prior knowledge of its behavior. Compare equations 5-7 with their Layer 2 counterparts:

$$U^3(W) = V_1(W_s^2, W_s^3) + \sum_{n \in N_s} V_2(W_s^2, W_s^3, W_n^2, W_n^3) \quad (5)$$

$$V_1(W_s^2, W_s^3) = -\sum_{s \in S} \bar{W}_s^2 W_s^3 \quad (6)$$

$$V_2(W_s^2, W_s^3, W_n^2, W_n^3) = -\sum_{s \in S} \sum_{n \in N_s} (W_s^3)^T J^3 \bar{W}_n^2 \quad (7)$$

The MxM square similarity matrix, J^3 is chosen to drive voxels classified to structures with large radii to propagate over voxels associated with structures with small radii. The bar over W denotes the Mean Field approximation, thus a vector of probabilities.

2.4 Layer 4: Global, Inter-structure Relationships

Broader context regarding relationships between structures can be incorporated in several ways. The stationary prior used in the calculation of the posterior probabilities by Layer 1 can be replaced by a spatially varying prior. [15] uses joint distances from ventricles and skin, while [16-18] use a rigidly registered digital atlas. Another approach, experimented with in Figure 1, is to associate a probability distribution in a similar manner as the distance to structure boundaries in Layer 3.

2.5 Layer 5: User Interaction

User time (such as in [6]) can be spent either initializing the segmentation, or correcting it to resolve ambiguities reached by the algorithm. Moreover, very small tumors in early stages of growth may require input from the 5th Layer (oracle) to disambiguate them from their surroundings. For our initial experiments, we have used Layer 5 for quick initialization of the Gaussian models for healthy tissue classes.

2.6 Expressing Tumors as Deviations from Normalcy

While the algorithm employs models of healthy tissue classes, pathology is only modeled by its deviation from normalcy. The degree to which each voxel belongs is its minimum class distance expressed in standard deviations, which is commonly known as the Mahalanobis-distance. Besides providing the basis for the fitness map, this computation also weights the evaluation of the bias field in Layer 1.

$$\min_{\Gamma} \sqrt{\frac{(I - \mu_{\Gamma})^2}{\sigma_{\Gamma}^2}} \quad (8)$$

3 Preliminary Experiments

Figure 1 depicts results of determining outliers with respect to each layer separately, and in total combination. Observe that a single layer is insufficient for recognition. Figure 2 explores a “toy” example to clearly illustrate the effects of layered vision.

Fig. 1. From left to right: Input MRI, Map of outliers based on intensity only, Map of outliers based on structure size only, Map of outliers based on relative position only, Combined map of outliers based on voxel intensity, structural size, and relative position to other structures. (Note: neck structures become removed upon rigid registration to a digital atlas.)





Fig. 2. The “toy” volume consists of 2 small, dark spheres and 2 large bright ones corrupted with Gaussian noise. The top, somewhat dark, and large sphere is ambiguous, and it is classified incorrectly by the lower-level layers of MAP and MRF. The 3rd layer then identifies that the center voxels are too distant from the boundary, and corrects their classification. The multi-level MRF propagates this information across the structure because its lower-level segmentation is mostly homogenous. From left to right: Original, Result after Layer 2, Result after 15 iterations of Layer 3’s multi-level MRF, Result after 50 iterations of Layer 3’s multi-level MRF.



4 Discussion

The contributions of this paper are two-fold. First, we proposed segmenting large brain tumors by training exclusively on healthy brains to recognize deviations from normalcy. Second, we designed a framework for layered vision that incorporates context at various levels. We extended EM-based segmentation with region-level properties, and we derived the novel multi-level MRF. While preliminary experiments are encouraging, there is future development to be performed before a clinical validation can be run. We are now investigating how to handle image noise and partial volume artifacts. Layer 3 is sensitive to noise in the lower layers, which is why the algorithm depends on the MRF relaxation of Layer 2. Additionally, voxels that contain tissue belonging to more than one tissue class display an intensity value along the linear combination of the classes’ distributions. While partial volume artifacts always present somewhat of an obstacle to segmentation, their effect becomes much more pronounced in our algorithm because the entire interface between structures incorrectly appears abnormal. We believe these issues can be solved, and that this approach, in general, represents a step in a new direction towards automated tumor segmentation.

References

1. D. Gering, A. Nabavi, R. Kikinis, N. Hata, L. Odonnell, W. Eric L. Grimson, F. Jolesz, P. Black, W. Wells III. "An Integrated Visualization System for Surgical Planning and Guidance Using Image Fusion and an Open MR". *Journal of Magnetic Resonance Imaging* June 2001; 13:967-975.

2. B.N. Joe, M.B. Fukui, C.C. Meltzer, Q. Huang, R.S. Day, P.J. Greer, M.E. Bozik. "Brain Tumor Volume Measurement: Comparison of Manual and Semiautomated Methods". *Radiology* 1999; 212:811-816.
3. M.C. Clark, L.O. Hall, D.B. Goldgof, R. Velthuizen, F.R. Murtagh, M.S. Silbiger. "Automatic Tumor Segmentation Using Knowledge-Based Techniques". *IEEE Transactions on Medical Imaging* April 1998; 17:238-251.
4. S. Warfield, J. Dengler, J. Zaers, C.R.G. Guttmann, W.M. Wells III, G.J. Ettinger, J. Hiller, R. Kikinis. "Automatic Identification of Grey Matter Structures from MRI to Improve the Segmentation of White Matter Lesions". *Journal of Image Guided Surgery* 1995; 6:326-338.
5. S.K. Warfield, M. Kaus, F.A. Jolesz, R. Kikinis. "Adaptive, Template Moderated, Spatially Varying Statistical Classification". *Medical Image Analysis* October 2000; 4:43-55.
6. M.R. Kaus, S.K. Warfield, A. Nabavi, P.M. Black, F.A. Jolesz, R. Kikinis. "Automated Segmentation of MR Images of Brain Tumors". *Radiology* 2001; 218:586-591.
7. Y. Zhu, H. Yan. "Computerized Tumor Boundary Detection Using a Hopfield Neural Network". *IEEE Transactions on Medical Imaging* February 1997; 16:55-67.
8. A.X. Falcao, J.K. Udupa, S. Samarasekera, S. Sharma. "User-Steered Image Segmentation Paradigms: Live Wire and Live Lane". *Graphical Models and Image Processing* 1998; 60:233-260.
9. A. Yezzi, S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum. "A Geometric Snake Model for Segmentation of Medical Imagery". *IEEE Transactions on Medical Imaging* April 1997; 16:199-209.
10. S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, A. Yezzi. "Gradient Flows and Geometric Active Contour Models". In: *International Conference on Computer Vision*. Boston: ; 810-815.
11. A.P. Dempster, N.M. Laird, D.B. Rubin. "Maximum Likelihood from Incomplete Data via the EM Algorithm". *Journal Royal Statistical Society* 1977; 39:1-38.
12. W.M. Wells III, W.E.L. Grimson, R. Kikinis, F.A. Jolesz. "Adaptive Segmentation of MRI Data". *IEEE Transactions on Medical Imaging* 1996; 15(4):429-443.
13. R. Guillemaud, M. Brady. "Estimating the Bias Field of MR Images". *IEEE Transactions on Medical Imaging* June 1997; 16:238-251.
14. K. Held, E.R. Kops, B.J. Krause, W.M. Wells III, R. Kikinis, H.-W. Muller-Gartner. "Markov Random Field Segmentation of Brain MR Images". *IEEE Transactions on Medical Imaging* December 1997; 16:878-886.
15. T. Kapur. "Model Based Three Dimensional Medical Image Segmentation". Ph.D. Thesis, Massachusetts Institute of Technology, 1999.
16. K.V. Leemput, F. Maes, D. Vandermeulen, P. Suetens. "Automated Model-Based Bias Field Correction of MR Images of the Brain". *IEEE Transactions on Medical Imaging* October 1999; 18:885-896.
17. K.V. Leemput, F. Maes, D. Vandermeulen, P. Suetens. "Automated Model-Based Tissue Classification of MR Images of the Brain". *IEEE Transactions on Medical Imaging* October 1999; 18:897-908.
18. K.V. Leemput, F. Maes, D. Vandermeulen, P. Suetens. "Automated Segmentation of Multiple Sclerosis Lesions by Model Outlier Detection". *IEEE Transactions on Medical Imaging* August 2001; 20:677-688.
19. S.Z.Li. *Markov Random Field Modeling in Image Analysis*. Springer-Verlag, 2001.