

Rapid Mode Estimation for 3D Brain MRI Tumor Segmentation

Haithem Boussaid, Iasonas Kokkinos, Nikos Paragios

Center for Visual Computing, Ecole Centrale de Paris, France
Galen, INRIA Saclay, France

{haithem.boussaid,iasonas.kokkinos,nikos.paragios}@ecp.fr

Abstract. In this work we develop a method for the efficient automated segmentation of brain tumors by developing a rapid initialization method. Brain tumor segmentation is crucial for brain tumor resection planning, and a high-quality initialization may have a significant impact on segmentation quality.

The main contribution of our work is an efficient method to initialize the segmentation by casting it as nonparametric density mode estimation, and developing a Branch and Bound-based method to efficiently find the mode (maximum) of the density function. Our technique is exact, has guaranteed convergence to the global optimum, and scales logarithmically in the volume dimensions by virtue of recursively subdividing the search space through Branch-and-Bound. Our method employs the Dual Tree data structure originally developed for nonparametric density estimation, and recently used for object detection with branch-and-bound. In this work we ‘close the loop’, and use the Dual Tree data structure for finding the mode of a density.

This estimated mode provides our system with an initial tumor hypothesis which is then refined by graph-cuts to provide a sharper outline of the tumor area. We demonstrate a 12-fold acceleration with respect to a standard mean-shift implementation, allowing us to accelerate tumor detection to a level that would facilitate high-quality brain tumor resection planning.

1 Introduction

The most common treatment of brain tumors is surgical resection, where the quality of the intervention can be largely affected by the efficient identification of the surgical margins during planning. Conventional segmentation techniques rely on prior knowledge and smoothness constraint to overcome the enormous variability of tumors both in terms of location as well as in terms of geometric characteristics. Even though recent studies [1] indicate statistically preferable locations for tumors in the brain and [2] proved that using this information improves substantially the results, in our work we take a more agnostic approach, using a clustering-based method for tumor detection.

Clustering, segmentation and nonparametric density mode estimation are related problems whose combination has been particularly studied in 2D images




Fig. 1. First row: patient MRI annotated with ground truth segmentations, Second row: Adaboost scores, Third row: detection bounding boxes are shown in yellow, our segmentations in red and ground truth segmentations in green. From left to right: horizontal plane, sagittal plane and coronal plane.

in the thread of works developed around the Mean-Shift method [3]. This method is also used as a component in a number of diagnosis tools such as vessel tracking [4], Multiple Sclerosis brain segmentation [5] and MRI brain clustering [6], and is a general tool that applies transversally to a host of problems in medical imaging.

In our work we develop a method to rapidly initialize a segmentation by relying on nonparametric density mode estimation. The mode of a nonparametric density estimate is used to pinpoint the center of the tumor, and thereby initialize a 3D segmentation implemented using graph-cuts. Our main contribution lies in addressing the computational complexity of the mode estimation problem.

The original Mean Shift method [3] is iterative, scales linearly in the number of points used in the Kernel Density Estimation (KDE) (as it follows the trajectory of every one of them) and can require careful checking of convergence. Faster variants of Mean Shift exist including Medoid Shift [7], Quick Shift [8], Fast Gauss transforms [9] as well as the Dual Tree variant of Mean Shift [10]. However, those of them that are exact [9, 10] are ‘dense’ i.e. evaluate the KDE over a dense set of locations; as such they may be inappropriate for application to 3D medical image volumes. Alternatively, those that focus on the modes [7, 8] are only approximate and have complexity proportional to $O(KN)$ where N is the number of pixels and K is the typical neighborhood size.

The main contribution of our work is a rapid mode estimation technique for 3D MRI image segmentation. Dealing with three dimensional data challenges

several algorithms which are reasonably efficient for 2D medical image analysis. In this paper, we leverage upon recent developments using Branch-and-Bound (BB) in object detection [11], which demonstrated that detection is possible in time sub-linear in the image size.

The main thrust of our work is the adaptation of this idea to the mode finding problem in KDE, typically addressed through Mean Shift. We propose an algorithm that can find the mode of the density with best-case complexity being logarithmic in the size of the search space.

We apply our algorithm to the setting of 3D brain tumor segmentation. Our algorithm takes the scores of a discriminatively trained classifier for tumor voxels and uses them to construct weights for a KDE-based estimate of the tumor location. Using standard mean shift would require tracking the trajectory of each voxel, and identifying the largest basin of attraction. Instead our algorithm narrows down the location of the maximum through an iterative branch-and-bound algorithm. In specific, we construct bounds on the value of the KDE over intervals of the search space, and use these bounds to devise a prioritized search strategy for the density’s mode. We demonstrate substantial speedups when comparing to the standard mean-shift algorithm.

Furthermore, we couple the mode estimation results with a post-processing step using graph-cuts, which allows us to boost the segmentation performance of our algorithm. Systematic evaluations are performed on clinical datasets demonstrating a 12-fold acceleration in speed over classical Mean-Shift while at the same time achieving robust tumor detection and segmentation.

2 3D Structure Detection

Our goal is to devise an algorithm that can quickly detect the largest region corresponding to a class (tumor in our case) within a 3D volume. This problem is particularly challenging for standard segmentation algorithms as it is hard to define exact boundaries between tumor and normal tissue [12]. Moreover, relying only on a classifier to separate the tumor class from the remaining structures in the MRI volume is challenging, due to the similarity between tumor and normal tissue and the high diversity in appearance of tumor tissue among different patients [12].

We start by phrasing our problem as mode seeking for a Kernel Density Estimate and then proceed to describing our Branch-and-Bound based optimization algorithm. We note that even though we focus on tumor segmentation, the same approach could potentially be useful to other maximization problems in 3D space.

2.1 Problem Formulation

We consider that we are provided with a scoring function that estimates the probability w_i with which a voxel x_i in \mathbb{R}^3 can belong to the considered class (i.e. a tumor vs non-tumor classifier). Namely, we have a mapping:

$$f : \mathbb{R}^3 \rightarrow [0, 1], x_i \mapsto w_i, \tag{1}$$

where f encapsulates the feature extraction around x_i and the subsequent formation of the class posterior. In specific, this score can be the output of a soft classifier or a likelihood function on the density distribution of the tumor class.

In order to pool information from multiple voxels and obtain a regularized labeling of the 3D volume, we phrase our problem in terms of a Kernel-based Density Estimator of the form:

$$KDE(x) = \sum_{i=1}^n w_i K_h(x - x_i) \quad (2)$$

We consider that K_h is a separable decreasing kernel, with the parameter h determining the used amount of smoothing. In the context of our application, we work with the finite-support Epanechnikov kernel [13]:

$$K_h(x - x_i) = \begin{cases} 0 & \text{if } \|x - x_i\| > h \\ \frac{3}{4} \left(1 - \left(\frac{\|x - x_i\|}{h}\right)^2\right) & \text{else,} \end{cases} \quad (3)$$

even though any other separable decreasing kernel could be used, e.g. an uniform or a Gaussian kernel [2]. We also note that in principle we should normalize the w_i elements to have unit sum, but the subsequent tasks are unaffected by this normalization. We address the problem of region detection in terms of mode estimation for the KDE above, namely we set out to find:

$$x^* = \operatorname{argmax}_{\mathbb{R}^3} S(x) = \operatorname{argmax}_{\mathbb{R}^3} \sum_{i=1}^n w_i K_h(x - x_i) \quad (4)$$

Instead of the iterative procedure employed by Mean-Shift, we now describe how Branch-and-Bound can be used to directly recover the solution of Eq.4.

2.2 Branch and Bound Algorithm

Branch-and-Bound is an optimization method that searches for the global maximum of a function $S(x)$. To this end, the algorithm employs a recursive subdivision of an interval of solutions X in its prioritized search for the maximum. The priority of an interval is determined by the function's upper bound \bar{S} within it. So, if we consider the maximum value of function S within the interval X , say $S(X) = \max_{x \in X} S(x)$, we bound it with $\bar{S}(X) \geq S(X)$. Moreover, we require $\bar{S}(\{x\}) = S(x)$

At each iteration a candidate domain X is popped from the priority queue, and split into subintervals. A new upper bound for each subinterval is computed and inserted in the priority queue. The bounding function drives BB to the most promising intervals until the first singleton interval, say x , is reached. Since the bound is tight for singletons, we know that the solutions contained in the remaining intervals of the priority queue will score below x , since the upper bound of their scores is below the returned singleton's score. This guarantees that once a singleton is popped from the priority queue, it will be the global maximum of S .

2.3 Bounding the KDE score

Having phrased the general setting of Branch-and-Bound, we now turn to how we can apply it to mode finding for Kernel Density Estimation; the main mathematical construct that we need is a bound on the score of a KDE within an interval of solutions. Namely, we need a function $\bar{S}(X)$ which gives us for an interval X an upper bound to the score of the KDE score within X :

$$\bar{S}(X) \geq \max_{x_j \in X} \sum_{\substack{n \\ x_i \in X'}} w_i K_h(x_j - x_i) = S(X). \quad (5)$$

We call points contained in X' the source locations and points in X the domain locations, with the intuition that the points in X' contribute to a score in X [14].

We now decompose the computation of the upper bound in Eq. 5 into smaller parts by using the partitions $X = \cup_{d \in D} X_d$ and $X' = \cup_{s \in S} X_s$. Our decomposition is based on the fact that $\max_x f(x) + g(x) \leq \max_x f(x) + \max_x g(x)$. For Eq. 5 this means that if we separately maximize some of the summands and add them back, this gives us something that will be larger than $S(X)$ (and as such, a valid candidate for $\bar{S}(X)$).

Based on this observation we introduce the quantity:

$$\mu_d^s = \max_{x_j \in X_d} \sum_{x_i \in X_s} w_i K_h(x_j - x_i) \quad (6)$$

and upper bound the right-hand side of Eq. 5 as:

$$S(X) \leq \max_d \sum_s \mu_d^s \quad (7)$$

where we have brought the summation over s outside the maximization over x_j . This means that if we can construct separately upper bounds to μ_d^s , we can add them up and obtain a valid upper bound for $S(X)$. This will then be used by Branch-and-Bound to prioritize the search over intervals that contain the density's mode.

In particular, we can upper bound $\mu_X^{X'}$ with $\bar{\mu}_X^{X'}$ as follows:

$$\bar{\mu}_X^{X'} \doteq \left(\sum_{i \in X'_s} w_i \right) \max_{i \in X} \max_{j \in X'} K(x_i, x_j) \quad (8)$$

Thus, the upper bound $\bar{S}(X)$ for $S(X)$ can be written as:

$$\bar{S}(X) \doteq \max_X \sum_{X'} \bar{\mu}_X^{X'} \geq S(X) \quad (9)$$

The first term in Eq. 8 can be computed rapidly over large intervals using fine-to-coarse summation. The second term can also be efficiently computed by exploiting the fact that X and X' are cubes, as also illustrated in the right of Fig. 2, and detailed in [14, 11].




Fig. 2. Left: an example of Dual Trees interaction: points belonging to Source node 6 have insignificant contribution to the objective sum computed in Domain points in node A. Right: Distance bounds between nodes in dual trees.

Our algorithm largely follows the one in [11]; it consists in performing a dual recursion, where the domain intervals X and source intervals X' are simultaneously refined in a coarse-to-fine manner. Namely starting with two intervals X and X' that correspond to the whole signal domain, at each iteration of the dual recursion these are split, bounded, and pruned if possible. A conservative pruning criterion can guarantee that there is no (or controllable loss in accuracy), while at the same time ensuring the number of bounds computed at every level of the refinement. The data structure used to implement this is called a dual-tree, as it involves two kd-trees corresponding to the two domain and source intervals respectively.

3 3D Brain Tumor Segmentation

Once a region of interest is efficiently selected, we proceed to segmentation in order to delineate the tumor from the normal tissue. Many segmentation methods have been proposed in the literature for tumor segmentation. MRF based-segmentation [15] has proved its performance and robustness in many applications. Therefore, we formulate the task of tumor segmentation from the 3D volume of interest as a discrete energy minimization problem. The 3D volume V is viewed as a lattice $\{\vartheta, \varepsilon\}$ where each voxel v_p forms a node in the graph. The MRF energy is written as:

$$E(V) = \sum_{p \in \vartheta} \Theta_p(v_p) + \sum_{(p,q) \in \varepsilon} \Theta_{pq}(v_p, v_q) \quad (10)$$

The function f serves as the unary potential energy. In order to improve the classification results, we use a regularization expressed by the binary potential energy. The conventional 4-neighborhood system is extended in 3D so that each voxel has 8 neighbors. We consider, in this work, the Potts model modulated by the contrast of normalized intensities as our regularization function.




Fig. 3. First row: a) patient MRI, b) adaboost segmentation result, c) graph cut result considering the whole volume. Second row: a) our method: bounding box b) our method segmentation results, c) ground truth segmentation.

This global criterion measures both the total dissimilarity among the two groups and the total similarity inside them. Global minimum of the considered energy is efficiently computed with the graph cut/max flow minimization algorithm[15, 16].

4 Experimental Evaluation

image size	profile	our method	Mean-Shift	exhaustive search
256x256x24	detection	2.5s	31s	60s
	overall time	17s	46.5s	75.5s
512x512x33	detection	8s	223s	319s
	overall time	93s	293.5s	389s

Table 1. Average computational time comparison

To evaluate our method on a real dataset, we use adaboost to provide us with the scores of individual voxels. It is based on the idea that a combination of weak classifiers such as decision trees can create a strong classifier. We use 40 randomly selected images to train the classifier with the following features: normalized intensities, locations (x,y,z), intensities of smoothed image at 3 half octave scales, gradient magnitude, Laplacian of Gaussian features at 3 half octave scales, absolute of Laplacian of Gaussian features at 3 different scales. Our classifier was trained with 50 rounds of boosting and we employed Decision Tress of Depth 3.




Fig. 4. Boxplots of the Dice values. From left to right: segmentation results with boosting only, boosting and pairwise regularization, boosting, rapid mode estimation and pairwise regularization

4.1 The Dataset

We did our experiments on a dataset composed of 113 patients with low grade gliomas. The patient age ranged from 21 to 65 years, and tumor size between 3.5 and 250 cm^3 . The MRI volume size varied from 256x256x24 to 512x512x33. The voxel resolution ranged from 0.4x0.4 to 0.9x0.9 mm^2 in the (x,y) plane and 5.3 to 6.4 mm in the z plane. The 3D images were rigidly aligned using *medInria* [17]. The dataset comes with a manual segmentation of the gliomas tumor done by experts, which is considered as our ground truth data.

4.2 Validation Methodology

In order to assess the quality of the segmentation results, we compute the Dice similarity coefficient, which reflects the overlapping rate between the segmented volume and the volume defined by experts. We evaluate the efficiency of our algorithm by comparing the computational time of the detection part with the Mean-Shift procedure and convolving the 3D volume with the kernel. Since the tumor size can achieve 250 cm^3 we convolve with an Epanechnikov kernel whose scale equals nearly the quarter of the volume size. This value matches the maximal size of the ground truth segmented gliomas. We use the most efficient available CPU version of convolution. The used package detects automatically if the kernel is separable and optimizes the convolution computation.

4.3 Results

The average Dice computed on our database is 0.73 (cf. Fig. 4) which is comparable to the results produced by the state of the art methods [18, 2]. we report from [18] that the computational cost is between 20 and 120 seconds and the

average DICE coefficient is 0.77. Our average computational time is 19 seconds. Mode estimation is a principal ingredient of the proposed method, as the results become less accurate if we only use either adaboost classifier or graph-cuts (cf. Fig. 4, Fig. 4). We compare the computational time between our work, a standard implementation of Mean Shift and an exhaustive search over volume locations after evaluating KDE in all locations cf. Table 1. We run the algorithms on a 4-core Intel Xeon computer which frequency is 2.67GHz. We use, though, a single core in the computation.

5 Discussion

While our method was inspired from previous work [14, 11], we recall that [14] does not use Branch-and-Bound and provides a technique for the efficient computation of a KDE score everywhere, on all ‘domain points’. Similarly, the multipole method [19] evaluates a KDE on all candidate locations, and is thus linear in the number of points. The aforementioned methods are excellent for KDE evaluation- but for mode estimation they perform an ‘overkill’, since they exactly evaluate the score everywhere, while we only want the location of the maximum. Instead our technique searches directly for the maximum location, and effectively ‘ignores’ less promising locations. In particular we discard chunks of points quickly by using cheaply computable upper bounds to their score. This allows us to work in time sublinear (practically logarithmic) in the number of possible locations. This is crucial for 3D medical data, where increasing the resolution by a factor of 2 will result in an 8-fold slowdown for Multipole/Dual Trees, but will require only 3 more iterations for our method (cf. Table. 1).

To the best of our knowledge, branch and bound has not been used before for mode estimation of KDEs. It was used in [11] for Object Detection, but with a different score function. We thus expect that our work will be of interest to other researchers working on mode estimation.

6 Conclusion

In this paper, we have presented a Branch-and-Bound based method for efficient mode estimation in KDE. We used our method for brain tumor detection and segmentation of 3D MR images. We demonstrate that our method results in a 12-fold speedup over standard Mean-Shift. Our approach is more robust than applying graph cut on the whole volume. The largest part of the computational time is taken by feature computation which can easily be accelerated with graphic processing units. Future directions include evaluating and adapting the proposed approach to the 3D liver tumor tracking in radiation therapy where real time is crucial.

7 Acknowledgements

This work has been funded by grant ANR-10-JCJC-0205; and by the EU Project MOBOT FP7-ICT-2011-600796. We would like to thank Hugues Duffau for providing us with the 3D MRI brain tumor dataset.

References

1. Duffau, H., C.L.: Preferential brain locations of low-grade gliomas. *Cancer* **100** (2004) 2622–2626.
2. Parisot, S., Duffau, H., Chemouny, S., Paragios, N.: Graph-based detection, segmentation & characterization of brain tumors. In: *CVPR*. (2012) 988–995.
3. Comaniciu, D., Meer, P.: Mean shift: a robust approach toward feature space analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**(5) (2002) 603–619.
4. van Walsum, T., Schaap, M., Metz, C.T., van der Giessen, A.G., Niessen, W.J.: Averaging centerlines mean shift on paths. (MICCAI 2008.)
5. Prima, D.G.L.S., Collins, D.L., Arnold, D.L., Morrissey, S.P., Barillot, C.: Combining robust expectation maximization and mean shift algorithms for multiple sclerosis brain segmentation. (MICCAI workshop on Medical Image Analysis on Multiple Sclerosis 2008.)
6. Mayer, A., Greenspan, H.: An adaptive mean-shift framework for mri brain segmentation. *IEEE TRANSACTIONS ON MEDICAL IMAGING* (2009.)
7. Sheikh, Y.A., E.Khan, Kanade, T.: Mode-seeking by medoidshifts. In: *ICCV 2007*. (Number 141.)
8. Vedaldi, A., Soatto, S.: Quick shift and kernel methods for mode seeking. (In: *ECCV 2008*.)
9. Yang, C., Duraiswami, R., Gumerov, N.A., Davis, L.S.: Improved fast gauss transform and efficient kernel density estimation. In: *ICCV*. (2003) 464–471.
10. Wang, P., Lee, D., Gray, A.G., Rehg, J.M.: Fast mean shift with accurate and stable convergence. *Journal of Machine Learning Research - Proceedings Track* **2** (2007) 604–611.
11. Kokkinos, I.: Rapid deformable object detection using dual-tree branch-and-bound. In Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., Weinberger, K., eds.: *NIPS*. (2011) 2681–2689.
12. Birkbeck, N., Cobzas, D., Jägersand, M., Murtha, A., Kesztyues, T.: An interactive graph cut method for brain tumor segmentation. In: *WACV*. (2009) 1–7.
13. Scott, D.: *Multivariate density estimation: theory, practice, and visualization*. Wiley series in probability and mathematical statistics: Applied probability and statistics. (Wiley. 1992.)
14. Gray, A.G.: Nonparametric density estimation: toward computational tractability. (In: *In SIAM International Conference on Data Mining 2003*.)
15. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. Pattern Anal. Mach. Intell.* **26** (2004) 1124–1137
16. Kolmogorov, V., Zabih, R.: What energy functions can be minimized via graph cuts. *PAMI* **26** (2004) 65–81.
17. Toussaint, N., Souplet, J.C., Fillard, P.: *Medinria: Medical image navigation and research tool by inria*. (In *Proc. of MICCAI'07 Workshop on Interaction in medical image analysis and visualization* (2007).)
18. Bauer, S., Nolte, L.P., Reyes, M.: Fully automatic segmentation of brain tumor images using support vector machine classification in combination with hierarchical conditional random field regularization. In: *MICCAI. MICCAI'11, Berlin, Heidelberg, Springer-Verlag* (2011) 354–361.
19. Engheta, N., Murphy, W.D., Rokhlin, V., Vassiliou, M.: The fast multipole method for electromagnetic scattering computation. *IEEE Transactions on Antennas and Propagation* **40** (1985) 634–641.




Fig. 5. Segmentation results: each 3 rows represent a patient case. First row: patient MRI with ground truth segmentations, Second row: Adaboost scores, Third row: detection bounding boxes are shown in yellow, our segmentations in red and ground truth segmentations in green. From left to right: horizontal plane, sagittal plane and coronal plane.