Fast and Robust Brain-Tissue Segmentation

S. P. Awate[1], and J. C. Gee[1]

[1]Radiology, University of Pennsylvania, Philadelphia, Pennsylvania, United States

Purpose: We present a fast and robust method for brain-tissue (white matter, gray matter, cerebrospinal fluid) segmentation in magnetic resonance (MR) images [1]. The method relies on a *spatially-adaptive nonparametric Markov random field* (MRF) image model. This general formulation enables the method to easily adapt to various kinds of MR images and the associated acquisition artifacts. It implicitly accounts for the intensity nonuniformity and *performs reasonably well without any inhomogeneity correction*. The method minimizes an information-theoretic metric on the probability density functions associated with image neighborhoods to produce an optimal segmentation. It automatically tunes its internal parameters based on the information content of the data. Experiments indicate that the *performance of the method is very robust to small changes in the values of the internal parameters*. Thus, the proposed method is easy to apply to a variety of clinical data and applications. The method incorporates a probabilistic atlas for initialization and as a prior during *Bayesian segmentation;* other priors are easily applicable. Our current implementation *segments a 1x1x1 mm^3 adult brain volume (about 2 million voxels inside the brain) in about 4 minutes on a standard 2.66 GHz Pentium-IV processor*. This is an order-of-magnitude improvement over the implementation in [1]. Experiments on simulated (BrainWeb) and real (IBSR) data demonstrate the advantages of the method over the current state-of-the-art.

Methods: The proposed method constructs a segmentation strategy based on a Markov statistical image model that it learns automatically from the input data. It formulates the segmentation problem as an optimization (maximization) of the mutual information between the segmentation labels and the Markov image statistics. Loosely speaking, the mutual information between two random variables quantifies the degree of functional dependence between them. For functionally dependent random variables, each variable uniquely determines the other, and the mutual information is maximized. On the other hand, independent random variables convey no information about each other, and their mutual information is zero (minimal). A "good" segmentation is one in which the voxel-neighborhood intensity values provide the most information about the class labels. Likewise, knowing the voxel class should provide a reliable estimate of the intensities in the voxel neighborhood.

Formulation: Consider a discrete random variable L that maps each voxel v to the class to which it belongs, i.e. $L(v)=k$ if voxel v is in class k. Let us define a mutually-exclusive and collectively-exhaustive decomposition of the image domain V into K segments, $\{v:L(v)=k\}$, with the Markov probability density functions (PDFs) $P(Z|L(v)=k)$ modeling the intensities in neighborhoods. Assume that at each voxel t, an instance $(l(v),z(v))$ is drawn from the joint PDF $P(Z,L)$. What we observe are, however, only the intensity vectors $z(v)$. We define the optimal segmentation as the one that maximizes the mutual information between L and Z, i.e. $I(L,Z)$.

Algorithm Overview: (Step 1) Given a segmentation, we estimate the Markov probabilities associated with image neighborhoods using nonparametric Parzen-window density estimation. (Step 2) Using these Markov probabilities, we reassign voxel labels to increase the mutual information $I(L,Z)$; each voxel is assigned the class k that maximizes the Markov probability. For most clinical studies, *convergence occurs in one iteration*.

To compute the Markov probabilities, nonparametrically, we draw a random sample of image neighborhoods from the image. The random selection results in a stochastic approximation for the Markov PDFs that alleviates the effects of spurious local maxima introduced in the finite-sample-size Parzen-window density estimate. To account for spatially-varying brain structure and intensity inhomogeneity, we use a local sampling strategy where, for each voxel v, we draw a unique random sample from an isotropic 3D Gaussian PDF on the image-coordinate space, with mean at the voxel v. Experiments demonstrate that the *method performs very robustly for a variety of choices of the sample size and Gaussian variance that produce more than several hundred voxels*. We use a neighborhood comprising 7 voxels corresponding to two voxel neighbors in each of the three cardinal directions.

Results and Validation: Our current implementation employs an *intelligent approximation of the algorithm* described above that provides an *order-of-magnitude speedup over a naïve implementation, while retaining the same level of performance*. The *algorithm scales linearly with the number of processors/cores* on a shared-memory machine. Our implementation relies on the Insight Toolkit (www.itk.org). We show results on the IBSR data (www.cma.mgh.harvard.edu/ibsr) comprising clinical MR images (including pathological cases and severe imaging artifacts) with manual segmentations. In the images on the right, we see an axial slice of the data followed by the automatic and manual segmentations. The graphs show the Dice overlaps for the segmentations with different strengths of the atlas prior (average Dice overlap: white matter 0.88; gray matter is 0.80).

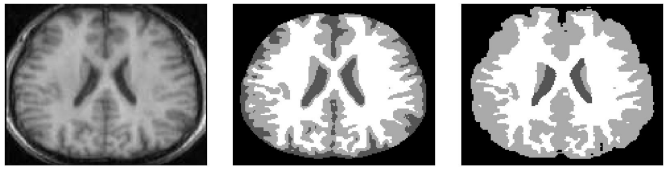

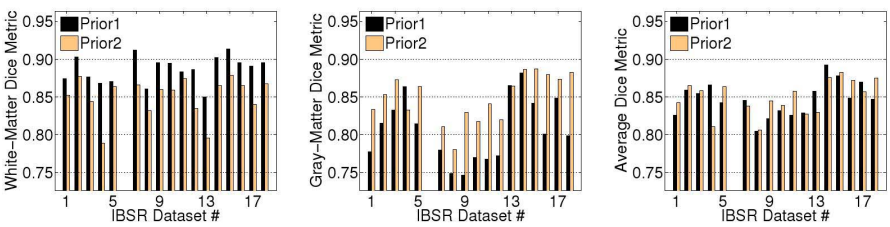

References: (1) Suyash P. Awate, Tolga Tasdizen, Norman Foster, Ross T. Whitaker. *Medical Image Analysis* 2006, 10(5):726-739

Acknowledgement: We are grateful to the NIH for supporting this work via grants NS045839, HD046159, HD042974, MH068066.