

A Statistical Framework for Partial Volume Segmentation

Koen Van Leemput^{1,2}, Frederik Maes¹, Dirk Vandermeulen¹, and Paul Suetens¹

¹ Katholieke Universiteit Leuven, Medical Image Computing, Radiology-ESAT, UZ Gasthuisberg, Herestraat 49, B-3000 Leuven, Belgium

² Helsinki University of Technology, Laboratory of Biomedical Engineering, P.O. Box 2200, FIN-02015 HUT, Finland
`koen.vanleemput@hut.fi`

Abstract. Accurate brain tissue segmentation by intensity-based voxel classification of MR images is complicated by partial volume (PV) voxels that contain a mixture of two or more tissue types. In this paper ¹, we present a statistical framework for PV segmentation that combines and extends existing techniques. We think of a partial volumed image as a downsampled version of a fictive higher-resolution image that does not contain partial voluming, and we estimate the model parameters of this underlying image using an Expectation-Maximization algorithm. This leads to an iterative approach that interleaves a statistical classification of the image voxels using spatial information and an according update of the model parameters. We illustrate the performance of the method on simulated data and on 2-D slices of real MR images. We demonstrate that the use of appropriate spatial models not only improves the classification, but is often indispensable for robust parameter estimation as well.

1 Introduction

Previously we presented a fully automated model-based approach for intensity-based tissue classification of MR images of the brain [1]. The method uses an Expectation-Maximization (EM) algorithm [2] to iteratively estimate during classification the parameters of tissue-specific Gaussian intensity models, of a polynomial model for MR bias field correction, and of a Markov random field (MRF) prior that models spatial interactions between neighboring voxels. Prior knowledge about the expected distribution of the various tissue classes in the image is derived from a digital brain atlas that is co-registered with the image under study, which allows full automation of the method.

Whereas this technique assigns each voxel to a single tissue type, the limited spatial resolution of MR imaging and the complex shape of the tissue interfaces in the brain imply that a large part of the voxels in MR brain images are so-called partial volume (PV) voxels, i.e. voxels that contain not a single tissue, but

¹ This paper is a short version of a technical report KUL/ESAT/PSI/0102. Available: <http://bilbo.esat.kuleuven.ac.be>

rather a mixture of two or more tissue types. In this paper, we therefore extend our approach by explicitly including a model for such PV voxels. We derive an EM algorithm for assessing the Maximum Likelihood (ML) parameters of the resulting image model, and demonstrate that this leads to a general PV segmentation framework that enables estimation of tissue-specific means and covariance matrices guided by spatial information, while classifying the image voxels at the same time.

2 Image Model

Let $\mathbf{L} = \{l_j, j = 1, 2, \dots, J\}$ be a label image with a total of J voxels, where $l_j \in \{1, 2, \dots, K\}$, denotes the one of K non-mixed tissue types to which each voxel j belongs. These labels are assumed to be drawn according to some probability distribution $f(\mathbf{L} | \Phi_L)$ with parameters Φ_L to be specified further that imposes certain spatial constraints. Suppose that a non-mixed intensity image $\mathbf{Y} = \{\mathbf{y}_j, j = 1, 2, \dots, J\}$ is generated from \mathbf{L} by drawing a sample from a probability distribution $f(\mathbf{Y} | \mathbf{L}, \Phi_Y)$ parameterized by Φ_Y . We assume that the intensity of each class k is normally distributed with mean $\boldsymbol{\mu}_k$ and covariance $\boldsymbol{\Sigma}_k$, such that $\Phi_Y = \{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k, k = 1, 2, \dots, K\}$ and $f(\mathbf{Y} | \mathbf{L}, \Phi_Y) = \prod_j f(\mathbf{y}_j | l_j, \Phi_Y) = \prod_j G_{\boldsymbol{\Sigma}_{l_j}}(\mathbf{y}_j - \boldsymbol{\mu}_{l_j})$ with $G_{\boldsymbol{\Sigma}}(\cdot)$ the zero-mean normal distribution with covariance matrix $\boldsymbol{\Sigma}$. With the exception of an explicit model for the MR bias field, which we assume here not to be present for the sake of simplicity, this was the image model used in our previous work [1].

Now an extra step is added, where \mathbf{Y} is not directly observed, but downsampled by a factor M to yield a partial volumed MR image $\tilde{\mathbf{Y}} = \{\tilde{\mathbf{y}}_i, i=1, 2, \dots, I\}$ with only $I = J/M$ voxels. The observed intensity $\tilde{\mathbf{y}}_i$ in voxel i of the downsampled image $\tilde{\mathbf{Y}}$ is modeled as the sum of the intensities \mathbf{y}_j of all the subvoxels in the original image \mathbf{Y} that underlie i . In voxels where not all subvoxels belong to the same tissue type, this causes partial voluming. Let \mathbf{t}_i be a vector that contains the relative amount of each class $k, k=1, 2, \dots, K$ in voxel i . A value of $t_{ik} = 1$ for some class k means that all the subvoxels underlying voxel i belong to class k , whereas a value of $t_{ik} = 0$ indicates that i does not contain class k at all. It can be shown that the observed intensity $\tilde{\mathbf{y}}_i$ in a voxel is governed by a normal distribution that only depends on its mixing proportions \mathbf{t}_i :

$$f(\tilde{\mathbf{y}}_i | \mathbf{t}_i, \Phi_Y) = G_{\tilde{\boldsymbol{\Sigma}}(\mathbf{t}_i)}(\tilde{\mathbf{y}}_i - \tilde{\boldsymbol{\mu}}(\mathbf{t}_i)) \tag{1}$$

with $\tilde{\boldsymbol{\mu}}(\mathbf{t}_i) = M \cdot \sum_k t_{ik} \boldsymbol{\mu}_k$ and $\tilde{\boldsymbol{\Sigma}}(\mathbf{t}_i) = M \cdot \sum_k t_{ik} \boldsymbol{\Sigma}_k$.

This model is illustrated in figure 1 on a 2-D example for two classes and $M = 3^2$ subvoxels per voxel. Figure 1f also shows the normal distributions $f(\tilde{\mathbf{y}}_i | \mathbf{t}_i, \Phi_Y)$ for each mixture $t = (\alpha, 1-\alpha)$ with $\alpha \in \{0, 1/M, 2/M, \dots, (M-1)/M, 1\}$, weighted by the number of times each mixture occurs in the image. Note that these mixture distributions do not all have the same weight, indicating that not all mixtures are equally likely. Summing the distributions of all non-pure mixtures yields a model for the intensity distribution of PV voxels.

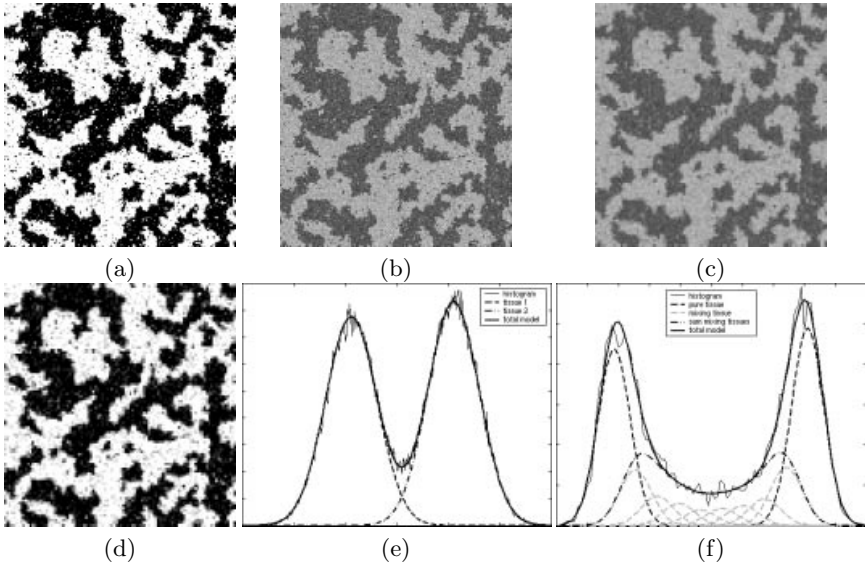


Fig. 1. PV image model: (a) label image \mathbf{L} with 2 classes drawn according to some statistical model; (b) intensity image \mathbf{Y} obtained from \mathbf{L} assuming the intensity of each class to be normally distributed; (c) downsampled image $\tilde{\mathbf{Y}}$ containing partial voluming; (d) tissue fraction $t_{i1} = 1 - t_{i2}$ in each voxel i of $\tilde{\mathbf{Y}}$; (e) and (f): histograms of \mathbf{Y} and $\tilde{\mathbf{Y}}$ respectively and their underlying models.

3 Model Parameter Estimation

Given the observed data $\tilde{\mathbf{Y}}$, the aim is to reconstruct the label image \mathbf{L} or, more modestly, to estimate the tissue fractions \mathbf{t}_i in each voxel. Before these issues can be addressed, the model parameters $\Phi = \{\Phi_Y, \Phi_L\}$ need to be estimated somehow from $\tilde{\mathbf{Y}}$. In this paper, we estimate the parameters that maximize the likelihood of the data $f(\tilde{\mathbf{Y}} | \Phi) = \sum_{\mathbf{L}} f(\tilde{\mathbf{Y}} | \mathbf{L}, \Phi_Y) f(\mathbf{L} | \Phi_L)$ with $f(\tilde{\mathbf{Y}} | \mathbf{L}, \Phi_Y) = \prod_i f(\tilde{y}_i | \mathbf{t}_i, \Phi_Y)$ given by (1). To this end, we use an EM algorithm that iteratively maximizes the expected value of the log-likelihood $\log f(\mathbf{Y}, \mathbf{L} | \Phi)$ of the unknown data $\{\mathbf{Y}, \mathbf{L}\}$, where the expectation is based on the observed data $\tilde{\mathbf{Y}}$ and the parameters $\Phi^{(m-1)}$ estimated in the previous iteration ($m - 1$) (see [3] for details):

Expectation step : find the function

$$Q(\Phi | \Phi^{(m-1)}) = E_{\mathbf{L}, \mathbf{Y}}[\log f(\mathbf{Y}, \mathbf{L} | \Phi) | \tilde{\mathbf{Y}}, \Phi^{(m-1)}] \tag{2}$$

Maximization step: find

$$\Phi^{(m)} = \arg \max_{\Phi} Q(\Phi | \Phi^{(m-1)})$$

Because $f(\mathbf{Y}, \mathbf{L} | \Phi) = f(\mathbf{Y} | \mathbf{L}, \Phi_Y) \cdot f(\mathbf{L} | \Phi_L)$, the expectation function can be written as $Q(\Phi | \Phi^{(m-1)}) = Q_Y(\Phi_Y | \Phi^{(m-1)}) + Q_L(\Phi_L | \Phi^{(m-1)})$. Maximization of Q_Y yields closed-form expressions for estimating the class intensity parameters Φ_Y that depend on the posterior probabilities $f(\mathbf{t}_i | \tilde{\mathbf{Y}}, \Phi^{(m-1)})$. These depend on the prior spatial model $f(\mathbf{L} | \Phi_L)$, for which three different models were investigated [3]:

Model A: no spatial correlation Every mixing combination \mathbf{t} has a spatially invariant prior probability $\pi_{\mathbf{t}}$.

Model B: no spatial correlation and uniform prior Every mixing combination \mathbf{t} has a spatially invariant prior probability $\pi_{\mathbf{t}}$, that is the same for all non-pure \mathbf{t} .

Model C: Markov random field The spatial distribution of labels l_j in L is governed by an Ising/Potts MRF whose parameters regulate how much of each tissue is present and how voxels of a particular tissue type are clustered.

For models A and B, $f(\mathbf{t}_i | \tilde{\mathbf{Y}}, \Phi^{(m-1)}) \propto f(\tilde{\mathbf{y}}_i | \mathbf{t}_i, \Phi_Y^{(m-1)}) \cdot \pi_{\mathbf{t}}^{(m-1)}$ with the prior probabilities $\pi_{\mathbf{t}}$ estimated by maximization of $Q_L(\Phi_L | \Phi^{(m-1)})$. For model C, $f(\mathbf{t}_i | \tilde{\mathbf{Y}}, \Phi^{(m-1)})$ can not be calculated analytically. We therefore use the MCEM algorithm [4] to approximate the expectation over the labels \mathbf{L} in (2) by sampling the distribution $f(\mathbf{L} | \tilde{\mathbf{Y}}, \Phi^{(m-1)})$. The MRF parameters are calculated using the pseudo-likelihood approximation [5].

4 Results

The performance of the 3 spatial models was validated on simulated 2-D data generated with different parameter sets Φ according to model C, i.e. with the underlying label image \mathbf{L} modeled as an MRF sample. The three spatial models A, B and C were fitted to the data starting from the same randomized initial parameters. In all cases, the MRF model C resulted in more accurate classifications than model A and B [3]. Moreover, models A and B, which are entirely histogram-based, often failed to correctly estimate the underlying model parameters. In many cases the use of prior spatial knowledge as provided by model C showed indispensable for robust estimation of the model parameters.

Figure 2 shows an example with 2 classes and $M = 3^2$ subvoxels per voxel. Model A fits the histogram well, but the underlying model is not correctly estimated. Since there is no restriction on the weights $\pi_{\mathbf{t}}$ for the mixing proportions \mathbf{t} , the algorithm has simply adjusted these to get a good histogram fit, thereby setting the prior probability for pure tissue to zero. With model B, all the mixing fractions corresponding to non-pure tissues are forced to have the same weight, resulting in the typical flat shape of the total intensity model for PV voxels that is commonly used in the literature [6,7,8,9]. However, the true mixing fractions in this example are not equal at all and therefore model B is condemned to fail. Only model C succeeds to retrieve the correct model parameters.

Figure 3 shows an example for 3 classes and $M = 2^2$ subvoxels per voxel. Because the intensity of voxels mixing the tissues with the lowest and highest




Fig. 2. Simulated data, two tissue types: (a) observed data \tilde{Y} ; (b) histogram of \tilde{Y} with the underlying model overlaid; (c) model initialization; (d,e,f): resulting histogram fit obtained with models A, B and C, respectively.

intensity is similar to the intensity of pure voxels of the tissue type with intermediate intensity, model A is severely underconstrained in this case and was therefore not considered. Model B and C fit the global histogram equally well, but only model C retrieves the correct model parameters, while model B yields incorrect parameters that are different with different initializations [3].

Our current implementation of the MRF of model C is only 2-D, implying that tissue boundaries are assumed to be orthogonal to the image slice. Figure 4a shows an axial slice of a high-resolution 1 mm isotropic T1-weighted image of the head through the central part of the brain where this assumption is more or less valid. Since different combinations of mixing tissues have overlapping intensities, model A was not considered. Comparing the results obtained with models B and C, it is clear that the MRF of model C reduces the noise in the segmentations considerably and forces partial volume voxels to lie on the border between the constituent tissues, in contrast to model B.

5 Discussion

Our statistical framework for PV segmentation is an extension of the methods of Choi *et al.* [10], Pham and Prince [11] and Nocera and Gee [12] when model C is used. Defining an MRF prior that imposes similar tissue combinations t_i over neighboring voxels, these approaches iteratively assign a mixing fraction t_i to each voxel and update the mean intensity of every pure tissue type. However,




Fig. 3. Simulated data, three tissue types: (a) label image L ; (b) observed partial volumed image \tilde{Y} ; (c) histogram of \tilde{Y} with the true model overlaid; (d) model initialization; (e,f) histogram fit obtained with model B and C, respectively.

these methods assume that all tissues have the same diagonal covariance structure, which does not need to be the case in real MR data, and the estimation of this covariance is not addressed. Furthermore, the methods in [10,12] use an MRF that simply imposes similar t_i over neighboring voxels, which erroneously encourages voxels to contain equal amounts of every tissue type everywhere in the image, in combination with extreme values for the mean intensity of pure tissues [3]. In contrast, our method additionally estimates tissue-specific covariance matrices as well, and defines an MRF model on subvoxels rather than on the voxels directly, thereby naturally imposing homogeneous regions of pure tissues bordered by PV voxels.

Our method is also an extension of the PV techniques of Santago and Gage [8], Wu *et al.* [9], Laidlaw *et al.* [6] and Ruan *et al.* [7]. These methods estimate their parameters by fitting the model to the histogram, thereby discarding all spatial information, and assume that all mixing proportions are equally alike when two tissues mix in a voxel, similar to our model B. However, the assumption of equally probable mixing fractions lacks any basis, as can be seen from figures 1f and 2b. Our model A does not make any prior assumption about the mixing proportions at all, but this introduces so many degrees of freedom that the model fitting is severely underconstrained. However, in contrast to [6,7,8,9], our approach allows incorporating a prior spatial model to guide the model fitting, which allows the mixing proportions to be non-uniform without making the estimation problem underconstrained.

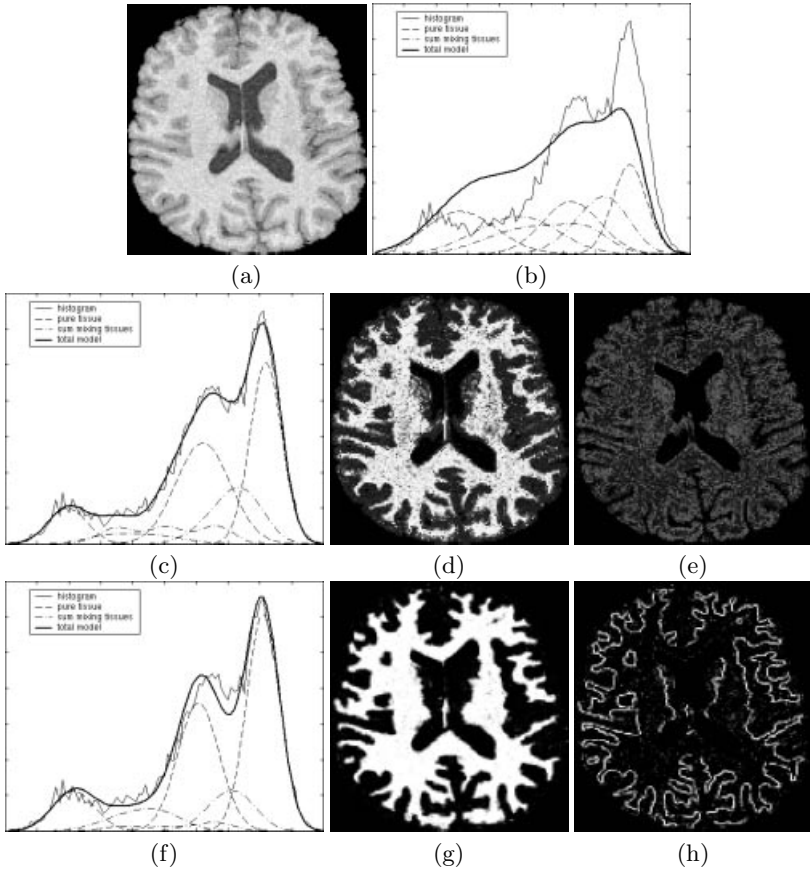


Fig. 4. 2-D MR image: (a) image data \tilde{Y} ; (b) histogram with the model initialization overlaid; (c,d,e) resulting histogram fit, expected fraction of white matter and estimated probability for partial voluming between white and gray matter, obtained using model B; (f,g,h) idem, but using model C.

While there may be many parameter sets that provide a close histogram fit, what discriminates the true solution is that it provides a meaningful classification in the images in agreement with the spatial model. With model C, our method iteratively interleaves a statistical classification of the image voxels using spatial information and an according update of the model parameters. Our results demonstrate that the use of such spatial information during the model fitting is often required in order to obtain reliable results.

6 Conclusion

In this paper, we presented a statistical framework for PV segmentation that combines and extends existing techniques. A partial volumed image is consid-

ered a downsampled version of a fictive higher-resolution image that does not contain partial voluming, and the model parameters of this underlying image are estimated using an EM algorithm. This leads to an iterative approach that interleaves a statistical classification of the image voxels using spatial information and an according update of the model parameters. We have demonstrated on simulated data that the use of appropriate prior spatial models not only improves the classification, but is often indispensable for robust parameter estimation as well. Future work will focus on developing such models that accurately describe the shape of the brain.

Acknowledgments

This work was supported by the EC-funded BIOMED-2 program under Grant BMH4-CT96-0845 (BIOMORPH) and Grant BMH4-CT98-6048 (QAMRIC), and by the Research Fund KULeuven under Grant GOA/99/05 (VHS+). The work of K. Van Leemput was supported by a grant from the IWT, Belgium. F. Maes is a Postdoctoral Fellow of the FWO-Vlaanderen, Belgium. The work of D. Vandermeulen was supported in part by the F.W.O.-Vlaanderen under Grant 1.5.397.97.

References

1. K. Van Leemput, F. Maes, D. Vandermeulen, and P. Suetens. Automated model-based tissue classification of MR images of the brain. *IEEE Transactions on Medical Imaging*, 18(10):897–908, october 1999.
2. A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, 39:1–38, 1977.
3. K. Van Leemput, F. Maes, D. Vandermeulen, and P. Suetens. A statistical framework for partial volume segmentation. Technical Report KUL/ESAT/PSI/0102, K.U.Leuven, ESAT, Leuven, Belgium, April 2001. Available: <http://bilbo.esat.kuleuven.ac.be>.
4. G. C. G. Wei and M. Tanner. A monte carlo implementation of the EM algorithm and the poor man's data augmentation algorithm. *J. Amer. Stat. Assoc.*, 85:699–704, 1990.
5. J. Besag. Efficiency of pseudo-likelihood estimation for simple gaussian fields. *Biometrika*, 64:616–618, 1977.
6. D. H. Laidlaw, K. W. Fleischer, and A. H. Barr. Partial-volume bayesian classification of material mixtures in MR volume data using voxel histograms. *IEEE Transactions on Medical Imaging*, 17(1):74–86, february 1998.
7. S. Ruan, C. Jaggi, J. Xue, J. Fadili, and D. Bloyet. Brain tissue classification of magnetic resonance images using partial volume modeling. *IEEE Transactions on Medical Imaging*, 19(12):1179–1187, December 2000.
8. P. Santago and H.D. Gage. Quantification of MR brain images by mixture density and partial volume modeling. *IEEE Transactions on Medical Imaging*, 12(3):566–574, September 1993.

9. Z. Wu, H.-W. Chung, and F.W. Wehrli. A bayesian approach to subvoxel tissue classification in NMR microscopic images of trabecular bone. *MRM*, 31:302–308, 1994.
10. H. S. Choi, D. R. Haynor, and Y. Kim. Partial volume tissue classification of multichannel magnetic resonance images—a mixel model. *IEEE Transactions on Medical Imaging*, 10(3):395–407, september 1991.
11. D. L. Pham and J. L. Prince. Unsupervised partial volume estimation in single-channel image data. In *Proceedings of IEEE Workshop on Mathematical Methods in Biomedical Image Analysis – MMBIA '00*, pages 170–177, 2000.
12. L. Nocera and J. C. Gee. Robust partial volume tissue classification of cerebral MRI scans. In K. M. Hanson, editor, *Proceedings of SPIE Medical Imaging 1997: Image Processing*, volume 3034 of *SPIE Proceedings*, pages 312–322. Bellingham, WA:SPIE, 1997.