

Unsupervised Partial Volume Estimation in Single-Channel Image Data

Dzung L. Pham
Laboratory of Personality & Cognition
Gerontology Research Center, NIA/NIH
Baltimore, MD 21224
dpham@cc.nih.gov

Jerry L. Prince
Center for Imaging Science
The Johns Hopkins University
Baltimore, MD 21218
prince@jhu.edu

Abstract

Partial volume effects are present in nearly all medical imaging data. These artifacts blur the boundaries between different regions, making accurate delineation of anatomical structures difficult. In this paper, we propose a method for unsupervised estimation of partial volume effects in single-channel image data. Based on a statistical image model, an algorithm is derived for estimating both partial volumes and the means of the different tissue classes in the image. To compensate for the ill-posed nature of the estimation problem, we employ a Bayesian approach that places a prior probability model on the parameters. We demonstrate on simulated and real images that the new algorithm is superior in several respects to the fuzzy and Gaussian clustering algorithms that have previously been used for modeling partial volume effects.

1 Introduction

Segmentation of three-dimensional volumetric images is an important goal in many medical imaging applications such as in the localization of pathology, quantification, and computer integrated surgery. Because of the finite resolution of imaging devices, however, nearly all images suffer from partial volume effects. These artifacts result when different tissue classes in the image contribute to a single voxel, thereby causing boundaries to be blurred and regions difficult to localize using standard segmentation methods. In addition, standard segmentation techniques that do not address partial volume effects fail to capture fine details that may be present in the original image.

Partial volume effects have been a particular concern in medical imaging applications because preservation of fine details, accurate quantification, and precise delineation of anatomical structures can all be critical in the diagnosis and analysis of pathology. The most common method of dealing with partial volume effects is to obtain a soft or fuzzy seg-

mentation. Rather than exclusively classifying a voxel as belonging to a particular class, soft segmentation methods allow for a continuous grade of membership within different classes. A soft segmentation is illustrated in Figure 1. Figure 1a shows one slice of a simulated magnetic resonance (MR) image of the brain [3]. Figure 1b is a map of the true partial volume contribution for gray matter in the simulated image. It is a continuously-valued image ranging from zero to one.

Because of the underdetermined nature of estimating partial volumes, most methods in the medical imaging literature have focused on multi-channel images (cf. [2, 11, 5]). Algorithms for estimating partial volumes in single-channel images have been fairly limited. Rather than measure the partial volume content, some methods have attempted to segment out partial volume voxels as a separate class [10]. Other methods have employed soft clustering or classification algorithms [4, 9, 8]. These algorithms, however, do not explicitly model partial volume effects and are therefore susceptible to certain artifacts in the resulting segmentation. Figure 1c shows a gray matter membership function computed using the fuzzy c-means algorithm (FCM), which has been shown to provide reasonable estimates of tissue volumes [1]. When compared to the true image in Figure 1b, the FCM result possesses an artificial rim around the ventricles near the center of the image. This is because the partial volume averaging of white matter and cerebrospinal fluid has the same intensity as gray matter. FCM operates purely on the image intensities and cannot distinguish between the two.

In this paper, we pose the segmentation problem in a framework that explicitly considers partial volume effects in scalar images. The true image is modeled as the union of disjoint sets in continuous space. The acquired image is a blurred, sampled, and noise degraded observation of the true image. By estimating the response of the acquisition to the different tissue classes in the image, we obtain a representation which localizes the anatomy while preserving nearly all the information present in the original image. As




Figure 1. Example of a soft segmentation: (a) slice from simulated MR brain image, (b) true gray matter partial volume, (c) FCM gray matter membership, (d) gray matter partial volume estimated using proposed new algorithm.

shown in Figure 1d, an advantage of our approach is that it can potentially avoid certain artifacts that occur in most other intensity-based classification algorithms. Our method is also unsupervised and fully automatic after the specification of a few parameters. This work greatly expands and improves upon our initial results presented in [7].

2 Imaging model

In this section we describe the model that will be used to describe the observed image. We begin with a continuous model of the original image, and allow partial volume effects to be introduced by the sampling process of the image acquisition. It is shown that the standard constraints on estimated partial volumes are valid under this model if certain mild constraints are placed on the point spread function of the acquisition device. Segmentation of the image is achieved by estimating the response of the imaging system to the characteristic functions corresponding to each tissue class. Because the estimation problem is severely ill-posed, a Bayesian approach is used to incorporate prior information and constrain the solution.

2.1 Image formation

We assume that the original image consists of K disjoint sets in the continuous (Euclidean) spatial domain Ω_c . In medical images, each set corresponds to a different tissue class. Each class need not be connected, and possesses a distinct intensity c_k for $k = 1, \dots, K$, called the class mean. We only consider the case when c_k is scalar valued. The true image $I(\mathbf{x})$ maps a spatial location $\mathbf{x} \in \Omega_c$ to a

class mean and can be written as

$$I(\mathbf{x}) = \sum_{k=1}^K \chi_k(\mathbf{x}) c_k, \mathbf{x} \in \Omega_c \quad (1)$$

where $\chi_k(\mathbf{x})$ is the characteristic function of set k .

Figure 2 shows our model of the image acquisition process. The original image is blurred by a linear shift invariant filter $h(\mathbf{x})$ (called the point spread function), point sampled, and corrupted by additive noise. Let $\Omega \subset \Omega_c$ be the set of the discrete samples in the observed image. The observed image can then be written as

$$y_j = \sum_{k=1}^K (\chi_k * h)_j c_k + \eta_j, j \in \Omega \quad (2)$$

where “*” denotes a convolution, η_j is white Gaussian noise with variance σ^2 .

We define the *tissue spread function* $s_k(\mathbf{x})$ to be the convolution of the characteristic function for tissue k and the point spread function:

$$s_k(\mathbf{x}) = (\chi_k * h)(\mathbf{x}) \quad (3)$$

We define the *tissue spread coefficients* s_{jk} to be samples of the tissue spread function at $j \in \Omega$:

$$s_{jk} = (\chi_k * h)_j \quad (4)$$

We can therefore rewrite Eq. (2) as

$$y_j = \sum_{k=1}^K s_{jk} c_k + \eta_j, j \in \Omega \quad (5)$$

Note that because of the blurring caused by the point spread function, a class need not be physically present in a voxel




Figure 2. Block diagram of image acquisition process.

to possess a non-zero spread coefficient. Ideally, a more accurate representation of the actual partial volume could be obtained by deconvolving the spread coefficients with the point spread function to reconstruct the underlying characteristic functions of each tissue class. Because the characteristic functions are not bandlimited functions, however, deconvolution is not practical. In this paper, we focus on the more modest task of estimating s_{jk} given the observations y_j , yielding a soft segmentation of each tissue class. Because the class means c_k are also unknown, we will describe a procedure for estimating both the class means and spread coefficients in Section 3.

2.2 Constraints on spread coefficients

Given an unconstrained point spread function, spread coefficients can have any value. For example, if the point spread function is a sinc function (the case of an ideal low pass filter), the resulting spread coefficients can be negative or positive valued. This unconstrained nature, however, leads to difficulties in interpreting a segmentation based on the spread coefficients as well as practical difficulties in deriving a robust estimation scheme.

It is therefore beneficial to restrict the spread coefficients to behave as membership functions [12], which satisfy the following constraints:

$$\sum_{k=1}^K s_{jk} = 1, \text{ for all } j \in \Omega \quad (6)$$

$$s_{jk} \geq 0, \text{ for all } j \in \Omega, k = 1, \dots, K \quad (7)$$

These constraints have a very intuitive notion and are also satisfied by posterior probability functions computed using Gaussian classification and clustering algorithms. It can be shown that the above constraints will be satisfied by the spread coefficients given fairly mild conditions on the point spread function [6]. We provide here two conditions that ensure that Eqs. (6) and (7) will be true for any image.

Proposition 1 *The spread coefficients satisfy Eq. (6) if*

$$\int_{\Omega_c} h(\mathbf{x}) d\mathbf{x} = 1.$$

Proposition 2 *The spread coefficients satisfy Eq. (7) if*

$$h(\mathbf{x}) \geq 0, \text{ for all } \mathbf{x} \in \Omega_c$$

The first proposition states that if the point spread function integrates to unity, Eq. (6) will be satisfied. This constraint is easily met because any non-unity scale factor that is introduced can be equivalently interpreted as a scale factor on the class means. The second proposition is more restrictive because the point spread function will not be nonnegative for all images. In general, $h(\mathbf{x})$ can possess negative lobes, such as in a sinc function. However, in most situations, the total positive area will nearly always be much larger than the negative area and Proposition 2 can be assumed to be satisfied to good approximation.

3 Estimation algorithm

In this section, we derive an algorithm for estimating the spread coefficients and class means. Note that even if the constraints (6) and (7) are satisfied, this estimation problem is still ill-posed. Thus, we further constrain the spread coefficients such that only two classes contribute to a single voxel. This is a reasonable assumption for high resolution images. We also place a prior probability model on both the spread coefficients and the class means.

3.1 MAP Estimation

Let \mathbf{s} denote the vector of spread coefficients and \mathbf{c} denote the vector of class means. Then the maximum *a posteriori* (MAP) estimate is given by

$$[\mathbf{s}, \mathbf{c}]^T = \underset{[\mathbf{s}, \mathbf{c}]^T}{\operatorname{argmax}} f(\mathbf{y}|\mathbf{s}, \mathbf{c}) f(\mathbf{s}) f(\mathbf{c}) \quad (8)$$

where \mathbf{y} is the vector of observed voxel intensities. We assume \mathbf{y} is conditionally independent given the spread coefficients. Then from Eq. (5), we obtain the joint conditional probability function:

$$f(\mathbf{y}|\mathbf{s}, \mathbf{c}) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left(-\frac{1}{2\sigma^2} \sum_{j \in \Omega} \left(y_j - \sum_{k=1}^K s_{jk} c_k \right)^2 \right) \quad (9)$$

For simplicity and without loss of generality, we assume that $2\sigma^2 = 1$. To further constrain the estimation problem, we assume that up to two classes, denoted class a and b , contribute to any single voxel. Thus, Eq. (9) simplifies to

$$f(\mathbf{y}|\mathbf{s}, \mathbf{c}) = \frac{1}{\sqrt{\pi}} \exp \left(- \sum_{j \in \Omega} (y_j - s_{ja}c_a - s_{jb}c_b)^2 \right) \quad (10)$$

where $s_{jb} = 1 - s_{ja}$. Here, a and b are implicitly functions of j .

3.2 Prior on spread coefficients

A prior probability model is placed on \mathbf{s} to enforce two *a priori* assumptions: (1) spread coefficients tend to be close to one or zero and (2) a spread coefficient tends to have a value similar to its neighbors. The former assumption is reasonable in high resolution data where the spread coefficient is less likely to be influenced by tissue boundaries. Based on these two assumptions and the constraints of Eqs.(6) and (7), we use the following marginal prior probability model:

$$f(\mathbf{s}_j) \propto \begin{cases} \pi(\mathbf{s}_j) \exp \left(\lambda \sum_{i \in N_j} (\mathbf{s}_j^T \mathbf{s}_i) \right) & s_{jk} \in [0, 1] \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

where $\mathbf{s}_j = [s_{j1}, \dots, s_{jK}]^T$, N_j is the set of neighbors of voxel j , and $\pi(\mathbf{s}_j)$ is a stationary prior used to weight different configurations of s_{jk} . The parameter λ controls the strength of the prior and is selected empirically. As a simple rule of thumb, λ should be in the range of the standard deviation of the noise in the image so that the log of the prior has a similar magnitude to the log of the joint conditional probability function (10). The joint prior probability is given by

$$f(\mathbf{s}) = \frac{1}{Z} \prod_{j \in \Omega} f(\mathbf{s}_j), \quad (12)$$

where Z is a normalizing constant.

Given K classes and the fact that only two classes may contribute to a single voxel, there are $D = \binom{K}{2}$ possible combinations of classes with nonzero spread coefficients. As we will show in Section 4, situations arise where certain classes are more likely to exhibit partial volume effects than others. The stationary prior probability function $\pi(\mathbf{s}_j)$ takes this into account by assigning a higher probability to more likely configurations. We constrain the stationary prior such that $\sum_{d=1}^D \pi_d = 1$, where the index d corresponds to a unique class configuration. In the absence of this prior knowledge, one can set $\pi_d = 1/D$.

Eq. (11) is maximized when s_{jk} is similar to its neighbors and when all s_{jk} are close to one or zero. Note also

that this form of a prior possesses an adaptive smoothing quality; greater smoothing is enforced where voxels represent pure tissue than in voxels that are partial volume averaged. This can be seen in the following example. Suppose $s_{ia} = 0.5$ and $s_{ib} = 0.5$ for some $i \in N_j$. Then $\mathbf{s}_j^T \mathbf{s}_i$ is equal to 0.5 for any value of \mathbf{s}_i , as long as the contributing classes are the same for the two voxels. Conversely, if $s_{ia} = 1$, the term is maximized only if $s_{ja} = 1$.

3.3 Prior on class means

A prior probability model can also be placed on the class means to provide more robust estimation. We note that Eq. (9) is maximized only if two class means are equal to the minimum and the maximum of the observed intensities, respectively. This condition is counterbalanced somewhat by the prior on the spread coefficients. However, we have found that the estimation is more robust when the following prior probability is applied:

$$f(c_k) = \frac{1}{\sqrt{2\pi\sigma_c^2}} \exp \left(-\frac{1}{2\sigma_c^2} (c_k - \mu)^2 \right) \quad (13)$$

where $\sigma_c = \sigma_i/K$ and σ_I and μ are the overall standard deviation and mean of the observed data, respectively. This biases the estimation of the class means towards being in the center of the intensity space, and in practice, we have found this to yield good results.

3.4 Algorithm description

To maximize the joint posterior probability function, a coordinate descent method is employed. The algorithm iteratively seeks a zero gradient condition of Eq. (8) with respect to the spread coefficients and class means by maximizing the marginal posterior probability at each voxel. To enforce the assumption that only two classes may contribute to a single voxel, \mathbf{s}_j must be solved for all D possible combinations, and the configuration yielding the highest probability is selected. This procedure will allow the algorithm to avoid some of the rim artifacts found in standard clustering methods.

The estimator for the spread coefficients is derived by taking the first derivative of Eq. (8) with respect to s_{ja} and setting it to zero. This yields

$$s_{ja} = \frac{(c_b - y_j)(c_b - c_a) + \lambda \left(\sum_{i \in N_j} (s_{ia} - s_{ib}) \right)}{(c_b - c_a)^2} \quad (14)$$

Values of s_{ja} are restricted to lie between zero and one. The value of s_{jb} is computed as $1 - s_{ja}$. Similarly, the zero gradient condition on the class means leads to the following




Figure 3. SETS applied to a four class test image: (a) one slice of the test image, (b) isosurface reconstruction of gray sphere using FCM, (c) isosurface reconstruction of gray sphere using SETS.

equation:

$$c_k = \frac{\sum_{j \in \Omega} s_{jk} (y_j - \sum_{l \neq k} s_{jl} c_l) + \mu / 2\sigma_c^2}{\sum_{j \in \Omega} s_{jk}^2 + 1/2\sigma_c^2}. \quad (15)$$

To initialize our algorithm, we obtain initial class means as in [8] and then apply the simplified, voxel independent estimation algorithm as described in [7]. Our algorithm, which we call Statistical Estimation of Tissue Spread (SETS), can be stated as follows:

SETS Algorithm:

1. Obtain initial class means.
2. Estimate spread coefficients using simplified model.
3. Estimate spread coefficients using the following procedure:
 - (a) Compute (14) for D combinations of a and b .
 - (b) Select spread coefficient configuration with highest probability.
 - (c) Repeat 3(a) & 3(b) for all voxels in the image.
4. Estimate class means using (15).
5. Repeat 3 & 4 until convergence.

Convergence is assumed to be achieved when the maximum change over all spread coefficients between iterations is less than 0.01. Because the joint posterior probability function is rather shallow—different spread coefficient configurations can yield almost identical probabilities—oscillations can result during early iterations of the algorithm. Thus, we employ a simple relaxation procedure which weights the current estimate with the previous estimate. This relaxation is eventually removed as iterations are increased and the algorithm nears convergence.

4 Results

Figure 3 shows the results of applying SETS and FCM to a three-dimensional simulated image consisting of three solid spheres. Figure 3a shows one slice from this image. Each sphere has a distinct intensity and the overall image has been slightly blurred. Both the FCM and SETS algorithms were applied to this image for four classes, corresponding to background, dark gray, gray, and white. Figure 3b shows an isosurface computed from the gray sphere's FCM membership function, at a level of 0.5. The figure has been cropped to focus on the non-empty regions. Similar to the example in Section 1, FCM cannot distinguish between the gray sphere and partial volume averaging of the white sphere and background, thus forming a broken shell around the region of the white sphere. Figure 3c shows the same isosurface computed from the SETS estimate of the gray sphere's spread coefficients. SETS correctly identifies the partial volume effects and no rim artifact is present.

Figure 4 shows the results of applying SETS, FCM, and the AGEM [9] algorithms to a real T_1 -weighted, 3-D MR image of the brain. AGEM is basically a Gaussian clustering algorithm with a Markov random field prior to enforce smoothness on the segmentation. It produces posterior probability functions for each class which can also be considered as an indicator of partial volume. The AGEM algorithm was performed without gain field correction. Figure 4a shows one slice from the original image. Figures 4b-d are the FCM membership functions for cerebrospinal fluid (CSF), gray matter (GM), and white matter (WM). Figure 4e is the FCM hard segmentation, computed by assigning the voxels to the class with highest membership. Note that the FCM result has two main problems. A rim around the ventricles is present in the GM membership, even though this is partial volume averaging of WM and CSF. In addition, the ventric-

Table 1. Error measures from simulated data results computed from four different methods: first three columns show the RMS error between the truth model and the estimated partial volumes; last column shows the misclassification rate (MCR) of the hard segmentations.

Method	Error measure			
	GM	WM	CSF	MCR
FCM	0.1393	0.1262	0.1251	3.988%
AGEM	0.1900	0.1632	0.2001	3.950%
SETS1	0.1518	0.1105	0.1194	3.808%
SETS2	0.1163	0.1122	0.1003	3.863%

ular CSF is darker than the sulcal CSF. This is an example of the asymptotic behavior of FCM (see [7]): in FCM, as an observed intensity becomes more distant from all the means, the membership values for all K classes converges to $1/K$.

Figures 4f-j show the results computed using the SETS algorithm ($\lambda = 15$, $\pi_d = 0.45$ for partial volume averaging of GM and WM, $\pi_d = 0.45$ for GM and CSF, and $\pi_d = 0.1$ for the partial volume averaging of WM and CSF, denoted SETS1). Figure 4f is simply the result of reconstructing the original image using the estimated class means and spread coefficients. Nearly all the information present in the original image has been retained. The SETS results do not exhibit either the rim artifact nor the poor asymptotic behavior. However, because the prior does not favor large regions of partial volume voxels, some of the subcortical GM details has been removed. This can be alleviated somewhat by adjusting the parameters. Despite this problem, details of the cortical GM have been very well preserved. The next row shows the same results computed using SETS ($\lambda = 15$ and $\pi_d = 0$ for partial volume averaging of WM and CSF, denoted SETS2). The result is more similar to the FCM result, but the poor asymptotic behavior is not present.

The last row shows the results using the AGEM algorithm. In general, AGEM's posterior probability functions are too hard to be used for partial volume estimation, causing it to lose some of the anatomical details present in the original image. Furthermore, the AGEM result also possesses the rim artifact. Although not observed in this image, the incorrect asymptotic behavior for partial volume estimation can sometimes occur in Gaussian clustering algorithms as well.

As a preliminary validation experiment, all four algorithms were applied to a simulated 3-D MR image, taken from the McGill University Brainweb database [3], with 3% noise and no inhomogeneity. A slice from this image is shown in Figure 1a. The first three columns of Table 1

show the average root mean squared (RMS) error between the truth model and the GM, WM, and CSF partial volume estimates. For FCM, the fuzzy membership functions were used as the partial volume estimates. For AGEM, the posterior probability functions were used. The last column shows the misclassification rate (MCR), computed as the ratio of misclassified voxels over the total number of voxels. The SETS algorithms generally provides improved results over the other methods. The only exception is the GM error for SETS1 is higher than for FCM. This is because of the loss of some of the subcortical gray matter structures outweighs the removal of the rim artifacts. As indicated by the hard segmentation results, however, these gray matter structures corresponded to voxels that were mostly white matter.

5 Future work and discussion

By varying the parameters of the prior probability models, a tradeoff was found in the SETS segmentation of MR brain images: on one hand, a rim artifact around the ventricles was obtained, and on the other hand, some subcortical gray matter details were lost. In some applications where accurate recovery of these structures is not required (eg. cortex reconstruction), this loss may be permissible. In future work, we plan to investigate modification of the prior models to limit this loss. Incorporation of atlas information might be one approach to alleviate this problem. In addition, extensions of the algorithm to simultaneously adapt to intensity inhomogeneities would be helpful for the processing of magnetic resonance images that possess shading artifacts.

Acknowledgments

Funding for this research was partially provided by NIH/NINDS Grant RO1NS37747.

References

- [1] M. Brandt, T. Bohan, L. Kramer, and J. Fletcher. Estimation of CSF, white and gray matter volumes in hydrocephalic children using fuzzy clustering of MR images. *Computerized Medical Imaging and Graphics*, 18:25–34, 1994.
- [2] H. Choi, D. Haynor, and Y. Kim. Partial volume tissue classification of multichannel magnetic resonance images—a mixel model. *IEEE Trans. Medical Imaging*, 10:395–407, 1991.
- [3] D. Collins, A. Zijdenbos, V. Kollokian, J. Sled, N. Kabani, et al. Design and construction of a realistic digital brain phantom. *IEEE Trans. Med. Imag.*, 17:463–468, 1998.
- [4] R. Herndon, J. Lancaster, J. Giedd, and P. Fox. Quantification of white matter and gray matter volumes using T1 parametric images using fuzzy classifiers. *J. Magnetic Resonance Imaging*, 6:425–435, 1996.




Figure 4. Comparison of soft segmentation algorithms: (a) one slice of original MR image, (b)-(d) FCM membership functions for gray matter, white matter, and CSF, (e) FCM hard segmentation, (f) SETS1 reconstructed image, (g)-(i) SETS1 spread coefficients, (j) SETS1 hard segmentation, (k) SETS2 reconstructed image, (l)-(n) SETS2 spread coefficients, (o) SETS2 hard segmentation, (p)-(r) AGEM posterior probabilities, (s) AGEM hard segmentation.

- [5] D. Laidlaw, K. Fleischer, and A. Barr. Partial-volume bayesian classification of material mixtures in MR volume data using voxel histograms. *IEEE Trans. Med. Imag.*, 17:98–107, 1998.
- [6] D. Pham. *Statistical estimation and pattern recognition methods for robust segmentation of magnetic resonance images*. PhD thesis, Johns Hopkins University, 1999.
- [7] D. Pham and J. Prince. Partial volume estimation and the fuzzy c-means algorithm. In *IEEE International Conference on Image Processing (ICIP98)*, volume III, pages 819–822, 1998.
- [8] D. Pham and J. Prince. Adaptive fuzzy segmentation of magnetic resonance. *IEEE Trans. Med. Imag.*, 18:737–752, 1999.
- [9] D. Pham and J. Prince. A generalized EM algorithm for robust segmentation of magnetic resonance images. In *Proceedings of The 33rd Annual Conference on Information Sciences and Systems (CISS99)*, pages 558–563, 1999.
- [10] P. Santago and H. Gage. Quantification of MR brain images by mixture density and partial volume modeling. *IEEE Trans. Medical Imaging*, 12:566–574, 1993.
- [11] H. Soltanian-Zadeh, J. Windham, D. Peck, and A. Yagle. Optimal transformation for correcting partial volume averaging effects in magnetic resonance imaging. *IEEE Trans. Nucl. Sci.*, 11:302–318, 1992.
- [12] L. Zadeh. Fuzzy sets. *Information and Control*, 8:338–353, 1965.