

Partial-Volume Bayesian Classification of Material Mixtures in MR Volume Data using Voxel Histograms

David H. Laidlaw, Kurt W. Fleischer, Alan H. Barr

Abstract—We present a new algorithm for identifying the distribution of different material types in volumetric datasets such as those produced with Magnetic Resonance Imaging (MRI) or Computed Tomography (CT). Because we allow for mixtures of materials and treat voxels as regions, our technique reduces errors that other classification techniques can create along boundaries between materials and is particularly useful for creating accurate geometric models and renderings from volume data. It also has the potential to make volume measurements more accurately and classifies noisy, low-resolution data well.

There are two unusual aspects to our approach. First, we assume that, due to partial-volume effects, or blurring, voxels can contain more than one material, e.g., both muscle and fat; we compute the relative proportion of each material in the voxels. Second, we incorporate information from neighboring voxels into the classification process by reconstructing a continuous function, $\rho(x)$, from the samples and then looking at the distribution of values that $\rho(x)$ takes on within the region of a voxel. This distribution of values is represented by a histogram taken over the region of the voxel; the mixture of materials that those values measure is identified within the voxel using a probabilistic Bayesian approach that matches the histogram by finding the mixture of materials within each voxel most likely to have created the histogram. The size of regions that we classify is chosen to match the spacing of the samples because the spacing is intrinsically related to the minimum feature size that the reconstructed continuous function can represent.

Keywords—Bayesian probability theory, discrete signal processing, feature detection, function theory, geometric modeling, image processing, partial volume, mixture modeling and estimation, multiscale analysis, multispectral classification, multivariate segmentation, magnetic resonance imaging microscopy, scale space, tissue classification, volume measurement

I. INTRODUCTION

IDENTIFYING different materials within sampled datasets can be an important step in understanding the geometry, anatomy, or pathology of a subject. By accurately locating different materials, we can identify them as individual parts and measure their size and shape. We can also use the spatial location of materials to selectively visualize parts of the data, thus better controlling a volume-rendered image [1], a surface model [2], or a volume model created from the data, and making visible otherwise obscured or subtle features. Classification is a key step towards understanding such geometry. Figure 1 shows an example of classified MRI data; each color represents a single material identified within the data.

Applications of classified images and geometric models derived from them include surgical planning and assistance, diagnostic medical imaging, conventional computer animation,

anatomical studies, and predictive modeling of complex biological shapes and behavior.

A. A partial-volume classification approach using voxel histograms.

We use Bayesian probability theory to estimate the highest-probability combination of materials within each voxel-sized region. The estimation is based on the histogram of data values within the region. The posterior probability, which we maximize, is based on conditional and prior probabilities derived from the assumptions about what we are measuring and how the measurement process works [3]. With this information we identify the materials contained within each voxel based on the sample values for the voxel and its neighbors. We treat each voxel as a region (see Figure 2), not as a single point. The sampling theorem [4] allows us to reconstruct a continuous function, $\rho(x)$, from the samples. We then represent all of the values that $\rho(x)$ takes on within a voxel by creating a histogram of $\rho(x)$ over the voxel. Figure 3(a) shows samples, Figure 3(b) shows the function $\rho(x)$ reconstructed from the samples, and Figure 3(c) shows a continuous histogram calculated from $\rho(x)$.


We assume that each voxel is a mixture of materials, with mixtures occurring where partial-volume effects occur, i.e., where the band-limiting process blurs measurements of pure materials together. From this assumption we derive basis functions that model histograms of voxels containing a pure material and of voxels containing a mixture of two materials. Linear combinations of these basis histograms are fit to each voxel, and the most likely combination of materials is chosen probabilistically.

The regions that we classify could be smaller or larger than voxels. Smaller regions would include less information, and so the context for the classification would be reduced and accuracy would suffer. Larger regions would contain more complicated geometry because the features that could be represented would be smaller than the region. Again, accuracy would suffer. Because the spacing of sample values is intrinsically related to the minimum feature size that the reconstructed continuous function, $\rho(x)$, can represent, that spacing is a natural choice for the size of regions to be classified.


B. Related work.

Many researchers have worked on identifying the locations of materials in sampled datasets [5], [6], [7], [8]. [9] gives an extensive review of the segmentation of MRI data. However, existing algorithms still do not take full advantage of all the information in sampled images; there remains room for improvement.

D. Laidlaw is at the Biological Imaging Center, Division of Biology, Beckman Institute, California Institute of Technology, Pasadena, CA 91125, dhl@gg.caltech.edu; K. Fleischer is at Pixar Animation Studios, Richmond, CA 94804, kurt@pixar.com; A. Barr is at the Computer Graphics Laboratory, Computer Science, Division of Engineering and Applied Science, Beckman Institute, California Institute of Technology, Pasadena, CA 91125, barr@gg.caltech.edu.




(a) Original Data



(b) Results of Algorithm

Classified White Matter (white), Gray Matter (gray)
Cerebro-Spinal Fluid (blue), Muscle (red)



(c) Combined Classified Image

Fig. 1. One slice of data from a human brain. (a) The original two-valued MRI data. (b) Four of the identified materials, white matter, gray matter, cerebro-spinal fluid, and muscle, separated out into separate images. (c) Overlaid results of the new classification mapped to different colors. Note the smooth boundaries where materials meet and the much lower incidence of misclassified samples than in Figure 4.




Fig. 2. We define a sample as a scalar or vector valued element of a 2-D or 3-D dataset. A voxel is the region surrounding a sample.




Fig. 3. Continuous histograms. The scalar data in (a) and (b) represent measurements from a dataset containing two materials, A and B , as shown in Figure 5. One material has measurement values near v_A and the other near v_B . These values correspond to the Gaussian-shaped peaks centered around v_A and v_B in the histograms, which are shown on their sides to emphasize the axis that they share. This shared axis is “feature space.”

Many of these algorithms generate artifacts like those shown in Figure 4, an example of data classified with a maximum-likelihood technique based on sample values. These techniques work well in regions where a voxel contains only a single material, but tend to break down at boundaries between materials. In Figure 4 note the introduction of both stair-step artifacts, as shown between gray matter and white matter within the brain, and thin layers of misclassified voxels, as shown by the white matter between the skull and the skin. Both types of artifacts can be ascribed to the partial-volume effects ignored by the segmentation algorithms and to the assignment of discrete material types to each voxel.

[10] presents a technique that uses *a priori* information about brain anatomy to avoid the layers of misclassified voxels. However, this work still produces a classification where each voxel is assigned to a single, discrete material; results continue to exhibit stair-step artifacts. It is also very dependent on brain anatomy information for its accuracy; broader applicability is not clear.

[11] demonstrates that accounting for mixtures of materials within a voxel can reduce both types of artifacts, and approximates the relative volume of each material represented by a sample as the probability that the sample is that material. Their technique works well for differentiating air, soft tissue, and bone in CT data, but not for differentiating materials in MR data, where the measured data value for one material is often identical to the measured value for a mixture of two other materials.

[12] and [13] avoid partial-volume artifacts by taking linear combinations of components of vector measurements. An advantage of their techniques is that the linear operations they perform preserve the partial-volume mixtures within each sample value, and so partial-volume artifacts are not created. A disadvantage is that the linear operations are not as flexible as non-




Fig. 4. Discrete, maximum-likelihood (DML) classification of the same brain data shown in Figure 1. This existing method assigns each voxel to a single material class. The class is identified here by its color: gray for gray matter, blue for CSF/fluid, white for white matter, red for muscle. Note the jagged boundaries between materials within the brain and the layer of misclassified white matter outside of the skull. See Section VII for more detail.

linear operations, and so either more data must be acquired or classification results will not be as accurate.

[14] and [15] address the partial-volume issue by identifying combinations of materials for each sample value. As with many other approaches to identifying mixtures, these techniques use only a single measurement taken within a voxel to represent its contents. Without the additional information available within each voxel region, these classification algorithms are limited in their accuracy.

[16] derives a distribution of data values taken on for partial volume mixtures of two materials. We share the distribution that they derive. Their application of the distribution, however, fits a histogram of an entire dataset and then quantifies material amounts over the entire volume. In contrast with our work, they represent each voxel with a single measurement for classification purposes, and do not calculate histograms over single voxels.

[17] presents an interesting approach to partial-volume imaging that makes assumptions similar to ours about the underlying geometry being measured and about the measurement process. The results of their algorithm are a material assignment for each sub-voxel of the dataset. Taken collectively, these multiple sub-voxel results provide a measure of the mixtures of materials within a voxel but arrive at it in a very different manner than we do. This work has been applied to satellite imaging data, and so their results are difficult to compare with ours, but aspects of both may combine well.

[18] gives an overview of the technique presented below in the context of the Human Brain Project, and [19] gives a complete description. [20] describes an imaging protocol for acquiring MRI data from solids and applies our classification technique to the extraction of a geometric model from MRI data of a human tooth (see Figure 10).

II. OVERVIEW

In this section we describe the classification problem that we solve, define terms, state assumptions we make about the data we classify, and sketch the algorithm and its derivation. Sections III–VI give more information on each part of the process, with detailed derivations in Appendices A and B. Section VII shows results of the application of the algorithm to simulated MR data and to real MR data of a human brain, hand, and tooth. We discuss some limitations and future extensions in Section VIII and conclude in Section IX.

A. Problem statement.

The input to our process is sampled measurement data, from which we can reconstruct a continuous, band-limited function, $\rho(x)$, that measures distinguishing properties of the underlying materials. The output is sampled data measuring the relative volume of each material in each voxel.

B. Definitions.

We refer to the coordinate system of the space containing the object we are measuring as “spatial coordinates,” and generally use $x \in X$ to refer to points. This space is n -dimensional, where n is three for volume data, can be two for slices, and is one for the example in Figure 3. Each measurement, which may be a scalar or vector, lies in “feature space” (see Figure 3), with points frequently denoted as $v \in V$. Feature space is n_v -dimensional, where n_v is one for scalar-valued data, two for two-valued vector data, etc. Tables IV and V in Appendix B contain these and other definitions.

C. Assumptions.

We make a set of assumptions about the objects that we are measuring and the measurement process.

1: Discrete materials. The first assumption is that materials within the objects that we measure are discrete at the resolution that we are sampling. Boundaries need not be aligned with the sampling grid. Figure 5(a) shows an object with two materials. We make this assumption because we are generally looking for boundaries between materials, and because we start from sampled data, where information about detail finer than the sampling rate is blurred.

This assumption does not preclude homogeneous combinations of sub-materials that can be treated as a single material at our sampling resolution. For example, muscle may contain some water, and yet be treated as a separate material from water. This assumption is not satisfied where materials gradually transition from one to another over many samples or are not relatively uniformly mixed; however, our algorithm appears to degrade gracefully even in these cases.

2: Normally-distributed noise. We assume that noise from the measurement process is added to each discrete sample and that the noise is normally distributed. We assume a different variance in the noise for each material. This assumption is not strictly satisfied for MRI data, but seems to be satisfied sufficiently to classify data well. Note that the sample values with noise added are interpolated to reconstruct the continuous func-

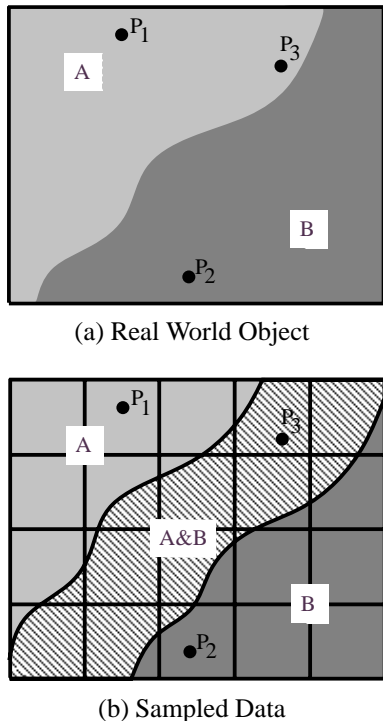


Fig. 5. Partial-volume effects. We start from the assumption that in a real-world object each point is exactly one material, as in (a). The measurement process creates samples that mix materials together; from the samples we reconstruct a continuous, band-limited measurement function, $\rho(x)$. Points P_1 and P_2 lie inside regions of a single material. Point P_3 lies near a boundary between materials, and so in (b) lies in the A&B region where materials A and B are mixed. The grid lines show sample spacing and illustrate how the regions may span voxels.

tion, $\rho(x)$. The effect of this band-limited noise is discussed further in Section VI.

3: Sampling theorem is satisfied. The third assumption we make is that the sampled datasets we classify satisfy the sampling theorem [4]. The sampling theorem states that if we sample a sufficiently band-limited function, we can exactly reconstruct that function from the samples, as demonstrated in Figure 3(b). The band-limiting creates smooth transitions in $\rho(x)$ as it traverses boundaries where otherwise $\rho(x)$ would be discontinuous. The intermediate region of Figure 5(b) shows a sampling grid and the effect of sampling that satisfies the sampling theorem. Partial-volume mixing of measurements occurs in the region labeled “A & B.” Multi-slice MRI acquisitions do not satisfy this assumption in the through-plane direction. For these datasets we interpolate the data only within each plane.

4: Linear mixtures. Each voxel measurement is a linear combination of pure material measurements and measurements of their pair-wise mixtures created by band limiting (see Figure 5).

5: Uniform tissue measurements. Measurements of the same material have the same expected value and variance throughout a dataset.

6: Box filtering for voxel histograms. The spatial measurement kernel, or point-spread function, can be approximated by a box filter for the purpose of deriving histogram basis functions.

7: Materials identifiable in histogram of entire dataset. The signatures for each material and mixture must be identifiable in a histogram of the entire dataset.

For many types of medical imaging data, including MRI and CT, these assumptions hold reasonably well, or can be satisfied sufficiently with preprocessing [21]. Other types of sampled data, e.g., ultrasound, and video or film images with lighting and shading, violate these assumptions, thus our technique does not apply directly to them.

D. Sketch of derivation.

Histograms represent the values taken on by $\rho(x)$ over various spatial regions. In section III we describe the histogram equation for a normalized histogram of data values within a region. In Section IV we use the histogram equation to create basis functions that model histograms taken over small, voxel-sized regions. These basis functions model histograms for regions consisting of single materials and for regions consisting of mixtures of two materials. Using Bayes’ Theorem, the histogram of an entire dataset, our histogram model basis functions, and a series of approximations, we derive an estimate of the most likely set of materials within an entire dataset in Section V. Similarly, given the histogram of a voxel-sized region, we derive, in Section VI, an estimate of the most likely density for each material in that voxel. The classification process is illustrated in Figure 6.

III. NORMALIZED HISTOGRAMS

In this section we present the equation for a normalized histogram of a sampled dataset over a region. We will use this equation as a building block in several later sections, with regions that vary from the size of a single voxel to the size of the entire dataset. We will also use this equation to derive basis functions that model histograms over regions containing single materials and regions containing mixtures of materials.

For a given region in spatial coordinates, specified by \mathcal{R} , the histogram $h^{\mathcal{R}}(v)$ specifies the relative portion of that region where $\rho(x) = v$, as shown in Figure 3. Because we treat a dataset as a continuous function over space, histograms, $h^{\mathcal{R}}(v) : \mathbf{R}^{nv} \rightarrow \mathbf{R}$, are also continuous functions:

$$h^{\mathcal{R}}(v) = \int \mathcal{R}(x) \delta(\rho(x) - v) dx \quad (1)$$

Equation 1 is the continuous analogue of a discrete histogram. $\mathcal{R}(x)$ is non-zero within the region of interest and integrates to one. We set $\mathcal{R}(x)$ constant in the region of interest, making every spatial point contribute equally to the histogram $h^{\mathcal{R}}(v)$, but $\mathcal{R}(x)$ can be considered a weighting function that takes on values other than zero and one to more smoothly transition between adjacent regions. Note also that $h^{\mathcal{R}}(v)$ integrates to one, which means that it can be treated as a probability density function, or PDF. δ is the Dirac-delta function.

A. Computing voxel histograms.

We calculate histograms in constant-sized rectangular “bins,” sized such that the width of a bin is smaller than the standard deviation of the noise within the dataset. This ensures that we do not lose significant features in the histogram.

We first initialize the bins to zero. We subdivide each voxel into sub-voxels, usually 4 for 2-D data or 8 for 3-D data, and evaluate $\rho(x)$ and its derivative at the center of each sub-voxel.




Fig. 6. The classification process. We collect MR data, calculate a histogram of the entire dataset, $h^{all}(v)$, and use that to determine parameters of histogram-fitting basis functions, one for each pure material and one for each mixture in the dataset. We then calculate histograms of each voxel-sized region, $h^{vox}(v)$, and identify the most likely mixture of materials for that region. The result is a sampled dataset of material densities within each voxel.

$\rho(x)$ is interpolated from the discrete data using a tri-cubic B-spline basis [22] that approximates a Gaussian. Thus, function and derivative evaluations can be made not only at sample locations, but anywhere between samples as well. From the function value and the derivative we use Equation 1 to calculate the contribution of a linear approximation of $\rho(x)$ over the sub-voxel to each histogram bin, accumulating the contributions from all sub-voxels. This gives us a more-accurate histogram than we would obtain by evaluating only the function values at the same number of points.

IV. HISTOGRAM BASIS FUNCTIONS FOR PURE MATERIALS AND MIXTURES

In this section we describe basis functions that model histograms of regions consisting of pure materials and of regions




Fig. 7. Parameters for histogram basis function. (a) Single-material histogram parameters include c , the mean value for the material, and s , which measures the standard deviation of measurements (see Equation 2). (b) Corresponding parameters for a two-material mixture basis function. s_0 and s_1 affect the slopes of the two-material histogram basis function at either end. For vector-valued data, c and s are vectors and are the mean values and standard deviations of the noise for the two constituent materials (see Equation 3).

consisting of pairwise mixtures of materials. Other voxel contents are also possible and are discussed in Section VIII. The parameters of the basis functions specify the expected value, c , and standard deviation, s , of each material's measurements (see Figure 7).

We use Equation 1 to derive these basis functions which we fit to histograms of the data. We then verify that the equations provide reasonable fits to typical MR data, which gives us confidence that our assumptions about the measurement function, $\rho(x)$, are reasonable. The details of the derivations are in Appendix A.

For a single material, the histogram basis function is a Gaussian distribution:

$$f_{\text{single}}(v; c, s) = \left(\prod_{i=1}^{n_v} \frac{1}{s_i \sqrt{2\pi}} \right) \exp \left(-\frac{1}{2} \sum_{i=1}^{n_v} \left(\frac{v_i - c_i}{s_i} \right)^2 \right), \quad (2)$$

where c is the vector-valued mean, s the vector-valued standard deviation, and v_i, c_i , and s_i scalar components of v, c , and s , respectively. We derive this equation by manipulating Equation 1 evaluated over a region of constant material, where the measurement function, $\rho(x)$, is a constant value plus additive, normally-distributed noise. Because the noise in different channels of multi-valued MRI images is not correlated, the general vector-valued normal distribution reduces to this equation with zero co-variances.

For mixtures along a boundary between two materials, we derive another equation similarly:

$$f_{\text{double}}(v; c, s) = \int_0^1 k_n((1-t)c_1 + tc_2 - v; s) dt \quad (3)$$

As with the single material case, this derivation follows from Equation 1 evaluated over a region where two materials mix. In this case, we approximate the band-limiting filter that causes partial-volume effects with a box filter and make the assumption that the variance of the additive noise is constant across the region. This basis function is a superposition of normal distributions representing different amounts of the two constituent pure materials. k_n is the normal distribution, centered at zero, t the relative quantity of the second material, c (comprised of c_1 and




Fig. 8. Basis functions fit to histogram of entire dataset. This figure illustrates the results of fitting basis functions to the histogram of the hand dataset. The five labeled circular regions represent the distribution of data values for pure materials, while the colored regions connecting them represent the distribution of data values for mixtures. The mixture between muscle (red) and fat (white), for example, is a salmon-colored streak. The green streak between the red and yellow dots is a mixture of skin and muscle. These fitted basis functions were used to produce the classified data used in Figure 11.

c_2) the expected values of the two materials, and s the standard deviation of measurements.

The assumption of a box filter affects the shape of the resulting histogram basis function. We derived similar equations for different filters (triangle, Gaussian, and Hamming), but chose the box filter derivation because we found it sufficiently accurate in practice and because its numerical tractability saves significant computation.

V. ESTIMATING HISTOGRAM BASIS FUNCTION PARAMETERS

In this section we describe parameter-estimation procedures for fitting histogram basis functions to a histogram of an entire dataset. For a given dataset we first calculate the histogram, $h^{\text{all}}(v)$, of the entire dataset. We then combine an interactive process of specifying the number of materials and approximate feature-space locations for them with an automated optimization [21] to refine the parameter estimates. Under some circumstances, users may wish to group materials with similar measurements into a single “material,” whereas in other cases they may wish the materials to be separate. The result of this process is a set of parameterized histogram basis functions, together with values for their parameters. The parameters describe the various materials and mixtures of interest in the dataset. Figure 8 shows the results of fitting a histogram. Each colored region represents one distribution, with the labeled spot-shaped regions representing pure materials and connecting shapes representing mixtures.

To fit a group of histogram basis functions to a histogram, as in Figure 8, the optimization process estimates the relative volume of each pure material or mixture (vector α^{all}), and the mean value (vector c) and standard deviation (vector s) of measurements of each material. The process is derived from the assumption that all values were produced by pure materials and

two-material mixtures. We define n_m as the number of pure materials in a dataset, and n_f as the number of histogram basis functions. Note that $n_f \geq n_m$, since n_f includes any basis functions for mixtures, as well as those for pure materials.

The optimization minimizes the function

$$\mathcal{E}(\alpha^{\text{all}}, c, s) = \frac{1}{2} \int \left(\frac{q(v; \alpha^{\text{all}}, c, s)}{w(v)} \right)^2 dv, \quad (4)$$

with respect to α^{all} , c , and s , where:

$$q(v; \alpha^{\text{all}}, c, s) = h^{\text{all}}(v) - \sum_{j=1}^{n_f} \alpha_j^{\text{all}} f_j(v; c_j, s_j) \quad (5)$$

Note that f_j may be a pure or a mixture basis function and that its parameter c_j will be a single feature-space point for a pure material or a pair for a mixture. The function $w(v)$ is analogous to a standard deviation at each point, v , in feature space, and gives the expected value of $|q(v)|$. We approximate $w(v)$ as a constant, and discuss it further in Section VIII.

Equations 4 and 5 are derived in Appendix B using Bayesian probability theory with estimates of prior and conditional probabilities.

VI. CLASSIFICATION

In this section we describe the process of classifying each voxel. This process is similar to that described in Section V for fitting the histogram basis functions to the entire dataset histogram, but now we are fitting histograms taken over small, voxel-sized regions. We use the previously computed histogram basis functions calculated from the entire dataset histogram and no longer vary the mean vector, c , or standard deviation, s . The only parameters allowed to vary are the relative material volumes (vector α^{vox}), and an estimate of the local noise in the local region (vector \bar{N}) (see Equations 6 and 7).

Over large regions including many voxels, the noise in $\rho(x)$ is normally distributed, with zero mean; however, for voxel regions the noise mean is generally non-zero. This is because normally distributed noise is added to each sample value, not to each point of $\rho(x)$. When the samples are used to reconstruct $\rho(x)$, the values $\rho(x)$ takes on near a particular sample tend to be similar, and so have a non-zero mean. We label the local mean voxel noise value \bar{N} . As derived in Appendix B the equation that we minimize, with respect to α^{vox} and \bar{N} , is:

$$\mathcal{E}(\alpha^{\text{vox}}, \bar{N}) = \frac{1}{2} \sum_{i=1}^{n_v} \left(\frac{\bar{N}_i}{\sigma_i} \right)^2 + \frac{1}{2} \int \left(\frac{q(v; \alpha^{\text{vox}}, \bar{N})}{w(v)} \right)^2 dv \quad (6)$$

where

$$q(v; \alpha^{\text{vox}}, \bar{N}) = h^{\text{vox}}(v - \bar{N}) - \sum_{j=1}^{n_f} \alpha_j^{\text{vox}} f_j(v), \quad (7)$$

the minimization is subject to the constraints

$$0 \leq \alpha_j^{\text{vox}} \leq 1, \text{ and } \sum_{j=1}^{n_f} \alpha_j^{\text{vox}} = 1,$$

and vector σ is the standard deviation of the noise over the entire dataset. For MR data the standard deviations in the signals for different materials are reasonably similar, and so we estimate σ to be an average of the standard deviations of the histogram basis functions.

With optimal vector α^{vox} for a given voxel-sized region and the mean value, vector \bar{v} , within that region, we solve for the amount of each pure material contributed by each mixture to the voxel. This is our output, the estimates of the amount of each pure material in the voxel-sized region.

$$\bar{v} = \int \left(h^{\text{vox}}(v) - \sum_{i=1}^{n_m} \alpha_i f_{\text{single}}(v) \right) dv \quad (8)$$

\bar{v} contains the mean signature of the portion of the histogram that arises only from regions with partial-volume effects. We determine how much of each pure component of pairwise mixture materials would be needed to generate \bar{v} , given the amount of each mixture that α^{vox} indicates is in the voxel. t_k represents this relative amount for mixture k , with $t_k = 0$ indicating that the mixture is comprised of only the first pure component, $t_k = 1$ indicating that it is comprised of only its second component, and intermediate values of t_k indicating intermediate mixtures. The t_k values are calculated by minimizing the following equation with respect to t , subject to the constraint $0 \leq t_k \leq 1$.

$$\mathcal{E}_{\bar{v}}(t) = \left(\bar{v} - \sum_{k=n_m+1}^{n_f} \alpha_k (t_k c_{k_a} + (1 - t_k) c_{k_b}) \right)^2 \quad (9)$$

Vector c_{k_a} is the mean value for the first pure material component of mixture k , and vector c_{k_b} the mean value for the second component. The total amount of each material is the amount of pure material added to the t_k -weighted portion of each mixture.

VII. RESULTS

We have applied our new technique to both simulated and collected MRI datasets. When results can be verified and conditions are controlled, as shown with the classification of simulated data, the algorithm comes very close to ‘‘ground truth,’’ or perfect classification. The results based on collected data illustrate that the algorithm works well on real data, with a geometric model of a tooth showing boundaries between materials, a section of a human brain showing classification results mapped on to colors, and a volume-rendered image of a human hand showing complex geometric relationships between different tissues.

We compare our Partial Volume Bayesian algorithm (PVB) with three other algorithms. The first, DML (Discrete Maximum Likelihood), assigns each voxel or sample to a single material using a Maximum Likelihood algorithm. The second, PPVC (Probabilistic Partial Volume Classifier), is described in [23], and the third is a Mixel classifier [14].

PVB significantly reduces artifacts introduced by the other techniques at boundaries between materials. In Figure 9 we compare performance of PVB, DML and PPVC on simulated data. PVB produces many fewer misclassified voxels, particularly in regions where materials are mixed due to partial-volume effects. In Figure 9(b) and (d) the difference is particularly noticeable where an incorrect layer of dark background material

TABLE I

COMPARATIVE RMS ERROR FOR THREE ALGORITHMS: PVB, PPVC, AND MIXEL. THE PPVC/PVB COMPARISON IS FROM A SIMULATED DATA TEST CASE ILLUSTRATED IN FIGURE 9, SNR=14.2. THE PPVC/MIXEL COMPARISON IS TAKEN FROM FIGURES 7 AND 8 IN [14], SNR=21.6. PVB, IN THE PRESENCE OF MORE NOISE, REDUCES THE PPVC RMS ERROR TO APPROXIMATELY HALF THAT OF THE MIXEL ALGORITHM.

	Improvement Ratio		
	PPVC	PVB	PPVC/PVB
Background	20%	6.5%	3.09
Outer	25%	4.3%	5.79
Inner	20%	6.5%	3.04
	PPVC	Mixel	PPVC/Mixel
Background	16%	9.5%	1.68
Tumor	21%	13.5%	1.56
White matter	37%	16.0%	2.31
Gray matter	36%	17.0%	2.11
CSF	18%	13.0%	1.38
All other	20%	10.0%	2.00

TABLE II

COMPARATIVE VOLUME MEASUREMENT ERROR FOR THREE ALGORITHMS (PVB, PPVC, AND MIXEL). THE PPVC/PVB COMPARISON IS FROM THE SIMULATED DATA TEST CASE ILLUSTRATED IN FIGURE 9, SNR=14.2. THE PPVC/MIXEL COMPARISON IS TAKEN FROM FIGURE 9 AND TABLE V IN [14], SNR=21.6.

PPVC	PVB	PPVC	Mixel
2.2%	0.01%	5.6%	1.6%
-5.3%	-0.45%	44.1%	7.0%
0.3%	0.15%		

has been introduced between the two lighter regions, and where jagged boundaries occur between each pair of materials. In both cases this is caused by partial-volume effects, where multiple materials are present in the same voxel.

Table I shows comparative RMS error results for the PPVC and PVB simulated data results, and also compares PPVC with the Mixel algorithm. Signal-to-noise ratio (SNR) for the data used in PPVC/PVB comparison was 14.2. SNR for the data used in PPVC/Mixel comparison was 21.6. Despite lower SNR, PPVC/PVB RMS error improvement is approximately double that of the PPVC/Mixel improvement. RMS error is defined as $\sqrt{\sum_x (\alpha(x) - p(x))^2}$, where $\alpha(x)$ is classified data and $p(x)$ is ground truth. The sum is made only over voxels that contain multiple materials.

Table II shows similar comparative results for volume measurements made between PPVC and PVB on simulated data, and between PPVC and Mixel on real data. Volume measurements made with PVB are significantly more accurate than those made with PPVC, and the PPVC to PVB improvement is better than the PPVC to Mixel improvement.

Figures 1 and 4 also show comparative results between PVB and DML. Note that the same artifacts shown in Figure 9 occur with real data and are reduced by our technique.

Models and volume-rendered images, as shown in Figures 10




Fig. 9. Comparison of DML classification (b), our new PVB classification (c), and PPVC classification (d). (a) is a reference for what “ideal” classification should produce. Note the band of dark background material in (b) and (d) between the two curved regions. This band is incorrectly classified, and could lead to errors in models or images produced from the classified data. The original dataset is simulated, two-valued data of two concentric shells, as shown in (e), with SNR of 14.2.

TABLE III

MRI DATASET SOURCES, ACQUISITION PARAMETERS, AND FIGURE REFERENCES.

Object	Machine	Voxel Size mm	$T_R/T_{E_1}/T_{E_2}$ s/ms/ms	Figs.
shells	simulated	$1.9^2 \times 3$	N/A	9
brain	GE	$0.94^2 \times 3$	2/25/50	1, 4
hand	GE	$0.7^2 \times 3$	2/23/50	11
tooth	Bruker	0.312^3	15/0.080	10




Fig. 10. A geometric model of tooth dentine and enamel created by collecting MRI data samples using a technique that images hard solid materials [20] and classifying dentine, enamel, and air in the volume data with our new PVB algorithm. Polygonal isosurfaces define the bounding surfaces of the dentine and enamel. The enamel-dentine boundary, shown in the left images, is difficult to examine non-invasively using any other technique.



Fig. 11. A volume-rendering image of a human hand dataset. The opacity of different materials is decreased above cutting planes to show details of the classification process within the hand.

and 11, benefit from our new techniques because less incorrect information is introduced into the classified datasets, thus the images and models more accurately depict the objects they are representing. Models and images such as these are particularly sensitive to errors at geometric boundaries because they illustrate the underlying geometries.

Table III lists the datasets, the MRI machine they were collected on, some collection parameters, the voxel size, and the figures in which each dataset appears. The GE machine was a 1.5T Signa. The Bruker machine was an 11.7T AMX500. Acquired data were collected with a spin-echo or fast spin-echo protocol, with one proton-weighted and one T_2 -weighted acquisition. The tooth was acquired with a technique described in [20]. Preprocessing was only performed on data used for the hand example (Figure 11). For this case each axial slice was multiplied by a constant and then offset by another to compensate for intensity falloff as a function of the distance from the center of the RF coil. The constants were chosen to make the mean values of user-identified material regions consistent from slice to slice.

VIII. DISCUSSION

We have made several assumptions and approximations while developing and implementing this algorithm. This section will discuss some of the tradeoffs, suggest some possible directions for further work, and consider some related issues.

A. Mixtures of three or more materials.

We assume that each measurement contains values from at most two materials. We chose two-material mixtures based on a dimensionality argument. In an object that consists of regions of pure materials, as shown in Figure 5, voxels containing one material will be most prevalent because they correspond to volumes. Voxels containing two materials will be next most prevalent, because they correspond to surfaces where two materials meet. As such, they are the first choice to model after those containing a single material. Our approach can be extended in a straightforward manner to handle the three-material case as well as cases with other less-frequent geometries, such as skin, tubes, or points where four materials meet. This extension could be useful for identifying sub-voxel-sized geometry within sampled data, thus extending the resolution.

B. Mixtures of materials within an object.

Based on our assumptions, voxels only contain mixtures of materials when those mixtures are caused by partial-volume effects. These assumptions are not true in many cases. By relaxing them and then introducing varying concentrations of given materials within an object, one could derive histogram basis functions parameterized by the concentrations and could fit them to measured data. The derivation would be substantially similar to that presented here.

C. Benefits of vector-valued data.

As with many other techniques, ours works on vector-valued volume data, in which each material has a characteristic vector value rather than a characteristic scalar value. Vector-valued datasets have a number of advantages and generally give better classification results. Such datasets have improved SNR and frequently distinguish similar materials more effectively (see Figure 12).

D. Partial mixtures.

We note that the histograms, $h^{\text{vox}}(v)$, for some voxel-sized regions are not ideally matched by a linear sum of basis functions. We discuss two possible sources of this mismatch.

The first source is the assumption that within a small region we still have normally distributed noise. \bar{N} models the fact that the noise no longer averages to zero, but we do not attempt to model the change in shape of the distribution as the region size shrinks.

The second source is related. A small region may not contain the full range of values that the mixture of materials can produce. The range of values is dependent on the bandwidth of the sampling kernel function. As a result, the histogram over that small region is not modeled ideally by a linear combination of pure material and mixture distributions. We are investigating model histogram basis functions with additional parameters




Fig. 12. Benefits of histograms of vector-valued data. We show histograms of an object with three materials. (a) This histogram of scalar data shows that material mean values are collinear. Distinguishing among more than two materials is often ambiguous. (b) and (c) are two representations of histograms of vector-valued data and show that mean values often move away from collinearity in higher dimensions, and so materials are easier to distinguish. High/bright locations indicate more-common (v_0, v_1) data values. While less likely, (d) shows that the collinearity problem can exist with vector-valued data.

to better match histograms [18], [19]. Modeling the histogram shape as a function of the distance of a voxel from a boundary between materials is likely to address both of these effects and give a result with a physical interpretation that will make geometric model extraction more justifiable and the resulting models more accurate.

We postulate that these two effects weight the optimization process such that it tends to make \bar{N} much larger than we expect. As a result, we have found that setting $w(v)$ to approximately 30 times the maximum value in $h^{\text{vox}}(v)$ gives good classification results. Smaller values tend to allow \bar{N} to move too much, and larger values hold it constant. Without these problems we would expect $w(v)$ to take on values equal to some small percentage of the maximum of $h^{\text{vox}}(v)$.

E. *Non-uniform spatial intensities.*

Spatial intensity in MRI datasets can vary due to inhomogeneities in the RF or gradient fields. We assume that they are small enough to be negligible for our algorithm, but it would be possible to incorporate them into the histogram basis functions by making the parameter c vary spatially.

F. *Quantitative comparison with other algorithms*

Because of the lack of a “gold standard” against which classification algorithms can be measured, it is difficult to compare our technique with others. Each technique presents a set of results from some application area, and so anecdotal comparisons can be made, but quantitative comparisons require reimplementing other algorithms. Work in generating a standard would greatly assist in the search for effective and accurate classification techniques. Our technique appears to achieve a given level of accuracy with fewer vector elements than the eigenimages of [12] or the classification results of [14], which use 3-valued data. Their results are visually similar to ours, and underscore the need for quantitative comparison. Because we interpolate neighboring sample values, we are able to achieve a given accuracy with two-valued or even scalar data, while their technique is likely to require more vector components. [13] shows good results for a human brain dataset, but we believe their technique will be less robust in the presence of material mixture signatures that overlap, a situation their examples do not include.

G. *Implementation.*

Our implementation is written in C and C++ on Unix workstations. We use a sequential-quadratic-programming (SQP) constrained-optimization algorithm [24] to fit h^{vox} for each voxel-sized region, and a quasi-Newton optimization algorithm for fitting h^{all} . The algorithm classifies approximately 10 voxels per second on a single HP9000/730, IBM RS6000/550E, or DEC Alpha AXP 3000 Model 500 workstation. We have implemented this algorithm in parallel on these machines and get a corresponding speedup on multiple machines. The performance is slower than many other methods and the current implementation would not be practical for a clinical situation. However, we believe that the algorithm can be sped up significantly through a more efficient implementation, which we have not attempted, and through the inevitable speedups in computer technology.

IX. CONCLUSIONS

Our new algorithm for classifying scalar- and vector-valued volume data produces more-accurate results than existing techniques in many cases, particularly at boundaries between materials. The improvements arise because: 1) we reconstruct a continuous function from the samples, 2) we use histograms taken over voxel-sized regions to represent the contents of the voxels, 3) we model of sub-voxel partial-volume effects caused by the band-limiting nature of the acquisition process, and 4) we use a Bayesian classification approach. We have demonstrated the technique on both simulated and real data, and it correctly classifies many voxels containing multiple materials. It also enables the creation of more-accurate geometric models and images. Because the technique correctly classifies voxels

containing multiple materials, it works well on low-resolution data, where such voxels are more prevalent. The examples also illustrate that it works well on noisy data ($\text{SNR} < 15$).

The construction of a continuous function is based on the sampling theorem, and while it does not introduce new information, it provides classification algorithms such as ours a richer context for the information. It incorporates neighbor information into the classification process for a voxel in a natural and mathematically rigorous way and thereby greatly increases classification accuracy. In addition, because the operations that can be safely performed directly on sampled data are so limited, treating the data as a continuous function helps to avoid introducing artifacts.

Histograms are a natural choice for representing voxel contents for a number of reasons. First, they generalize single measurements to measurements over a region, allowing classification concepts that apply to single measurements to be generalized. Second, the histograms can be calculated easily. Third, the histograms capture information about neighboring voxels; this increases the information content over single measurements and improves classification results. Fourth, histograms are orientation independent; orientation independence reduces the number of parameters in the classification process hence simplifying and accelerating it.

Partial-volume effects are a nemesis of classification algorithms, which traditionally have drawn from techniques that classify isolated measurements. These techniques do not take into account the related nature of spatially-correlated measurements. Many attempts have been made to model partial-volume effects, and ours continues that trend, with results that suggest that continued study is warranted.

We believe that the Bayesian approach we describe is a useful formalism for capturing the assumptions and information gleaned from the continuous representation of the sample values, the histograms calculated from them, and the partial-volume effects of imaging. Together, these allow a generalization of many sample-based classification techniques, one of which we have demonstrated.

X. ACKNOWLEDGMENTS

Many thanks to Matthew Avalos, who has been instrumental in implementation. Thanks to Barbara Meier, David Kirk, John Snyder, Bena Currin, and Mark Montague for reviewing early drafts and making suggestions. Our data was collected in collaboration with the Huntington Magnetic Resonance Center with the cooperation of Brian Ross and Jose Jimenez, and at the Caltech Biological Imaging Center jointly with Pratik Ghosh and Russell Jacobs.

This work was supported in part by grants from Apple, DEC, Hewlett Packard, and IBM. Additional support was provided by NSF (ASC-89-20219) as part of the NSF/ARPA STC for Computer Graphics and Scientific Visualization, by the DOE (DE-FG03-92ER25134) as part of the Center for Research in Computational Biology, by the National Institute on Drug Abuse, the National Institute of Mental Health, and the National Science Foundation as part of the Human Brain Project, and by the Beckman Institute Foundation. All opinions, findings, conclusions, or recommendations expressed in this document are those

of the authors and do not necessarily reflect the views of the sponsoring agencies.

APPENDIX

I. DERIVATION OF HISTOGRAM BASIS FUNCTIONS

In this appendix we derive parameterized model histograms that we use as basis functions, f_i , for fitting histograms of data. We derive two forms of basis functions: one for single, pure materials; another for two-material mixtures that arise due to partial-volume effects in sampling. Equation 1, the histogram equation, is:

$$h^{\mathcal{R}}(v) = \int \mathcal{R}(x)\delta(\rho(x) - v)dx$$

and measures a histogram of the function $\rho(x)$ over a region defined by $\mathcal{R}(x)$. x ranges over spatial locations, and v over feature space. Note that if $\rho(x)$ contains additive noise, $n(x; s)$, with a particular distribution, $k_n(v; s)$, then the histogram of ρ with noise is the convolution of $k_n(v; s)$ with $\rho(x) - n(x; s)$ (i.e., $\rho(x)$ without noise). $k_n(v; s)$ is, in general, a normal distribution. Thus,

$$h^{\mathcal{R}}(v) = k_n(v; s) * \int \mathcal{R}(x)\delta((\rho(x) - n(x; s)) - v)dx$$

A. Pure Materials

For a single pure material we assume that the measurement function has the form:

$$\rho_{\text{single}}(x) = c + n(x; s) \quad (11)$$

where c is the constant expected value of a measurement of the pure material, and s is the standard deviation of additive, normally-distributed noise.

The basis function we use to fit the histogram of the measurements of a pure material is

$$\begin{aligned} f_{\text{single}}(v; c, s) &= \int \mathcal{R}(x)\delta(\rho_{\text{single}}(x) - v)dx \\ &= \int \mathcal{R}(x)\delta(c + n(x; s) - v)dx \\ &= k_n(v; s) * \int \mathcal{R}(x)\delta(c - v)dx \\ &= k_n(v; s) * \left(\delta(c - v) \int \mathcal{R}(x)dx \right) \\ &= k_n(v; s) * \delta(c - v) \\ &= k_n(v - c; s) \\ &= \left(\prod_{i=1}^{n_v} \frac{1}{s_i \sqrt{2\pi}} \right) \exp \left(-\frac{1}{2} \sum_{i=1}^{n_v} \left(\frac{v_i - c_i}{s_i} \right)^2 \right) \quad (12) \end{aligned}$$

Thus, $f_{\text{single}}(v; c, s)$ is a Gaussian distribution with mean c and standard deviation s . v_i , c_i , and s_i are scalar components of v , c , and s . We assume the noise is independent in each element of vector-valued data, which for MRI appears to be reasonable.

B. Mixtures

For a mixture of two pure materials, we assume the measurement function has the form:

$$\rho_{\text{double}}(x) = \ell_{\text{double}}(x; c_1, c_2) + n(x; s) \quad (13)$$

where ℓ_{double} approximates the band-limiting filtering process, a convolution with a box filter, by interpolating the values within the region of mixtures linearly between c_1 and c_2 , the mean values for the two materials.

$$\ell_{\text{double}} = (1 - t)c_1 + tc_2 \quad (14)$$

$$\begin{aligned} f_{\text{double}}(v; c, s) &= \int \mathcal{R}(x)\delta(\rho_{\text{double}}(x) - v)dx \\ &= \int \mathcal{R}(x)\delta(\ell_{\text{double}}(x; c_1, c_2) + n(x; s) - v)dx \\ &= k_n(v; s) * \int \mathcal{R}(x)\delta(\ell_{\text{double}}(x; c_1, c_2) - v)dx \\ &= \int_0^1 k_n(v; s) * \delta((1 - t)c_1 + tc_2 - v)dt \\ &= \int_0^1 k_n((1 - t)c_1 + tc_2 - v; s)dt \quad (15) \end{aligned}$$

II. DERIVATION OF CLASSIFICATION PARAMETER ESTIMATION

In this appendix we derive the equations that we optimize to find model histogram parameters and to classify voxel-sized regions. We use Bayesian probability theory [3] to derive an expression for the probability that a given histogram was produced by a particular set of parameter values in our model. We maximize an approximation to this ‘‘posterior probability’’ to estimate the best-fit parameters.

$$\text{maximize } P(\text{parameters} \mid \text{histogram}) \quad (16)$$

We use this optimization procedure for two purposes:

- **Find model histogram parameters.** Initially, we find parameters of basis functions to fit histograms of the entire dataset h^{all} . This gives us a set of basis functions that describes histograms of voxels containing pure materials or pairwise mixtures.
- **Classify voxel-sized regions.** We fit a weighted sum of the basis functions to the histogram of a voxel-sized region h^{vox} . This gives us our classification (in terms of the weights α). The posterior probabilities P^{all} and P^{vox} share many common terms. In the following derivation we distinguish them only where necessary, using P where their definitions coincide.

A. Definitions

Table IV lists Bayesian probability terminology as used in [3] and in our derivations. Table V defines additional terms used in this section.

TABLE IV
PROBABILITIES, USING BAYESIAN TERMINOLOGY FROM [3].

$P(\alpha, c, s, \bar{N} h)$	posterior probability (we maximize this)
$P(\alpha, c, s, \bar{N})$	prior probability
$P(h \alpha, c, s, \bar{N})$	likelihood
$P(h)$	global likelihood

TABLE V
DEFINITIONS OF TERMS USED IN THE DERIVATIONS.

Term	Dim.	Definition
n_m	scalar	number of pure materials
n_f	scalar	number of pure materials & mixtures
n_v	scalar	dim. of measurement (feature space)
α	n_f	relative volume of each mixture and material within the region
c	$n_f \times n_v$	mean of material measurements for each material
s	$n_f \times n_v$	standard deviation of material measurements (chosen by procedure discussed in Section V) for each material
\bar{N}	n_v	mean value of noise over the region
p_{1-6}	scalars	arbitrary constants
$h^{\text{all}}(v)$	$\mathbf{R}^{n_v} \rightarrow \mathbf{R}$	histogram of an entire dataset
$h^{\text{vox}}(v)$	$\mathbf{R}^{n_v} \rightarrow \mathbf{R}$	histogram of a tiny, voxel-sized region

B. Optimization

We perform the following optimization to find the best-fit parameters:

$$\text{maximize } P(\alpha, c, s, \bar{N}|h) \quad (17)$$

With $P \equiv P^{\text{all}}$, we fit histogram basis function parameters $c, s, \alpha^{\text{all}}$ to the histogram of an entire dataset, $h^{\text{all}}(v)$. With $P \equiv P^{\text{vox}}$, we fit $\alpha^{\text{vox}}, \bar{N}$ to classify the histogram of a voxel-sized region, $h^{\text{vox}}(v)$.

C. Derivation of the posterior probability, $P(\alpha, c, s, \bar{N}|h)$

We start with Bayes' Theorem, expressing the posterior probability in term of the likelihood, the prior probability, and the global likelihood.

$$P(\alpha, c, s, \bar{N}|h) = \frac{P(\alpha, c, s, \bar{N})P(h|\alpha, c, s, \bar{N})}{P(h)} \quad (18)$$

Each of the terms on the right side is approximated below, using p_{1-6} to denote positive constants (which can be ignored during the optimization process).

C.1 Prior probabilities.

We assume that α, c, s , and \bar{N} are independent, so

$$P(\alpha, c, s, \bar{N}) = P(\alpha)P(c, s)P(\bar{N}) \quad (19)$$

Because the elements of α represent relative volumes, we require that they sum to 1 and are positive.

$$P(\alpha) = \begin{cases} 0 & \text{if } \sum_{j=1}^{n_f} \alpha_j \neq 1 \\ 0 & \text{if } \alpha_j < 0 \text{ or } \alpha_j > 1 \\ p_1 & \text{(constant) otherwise} \end{cases} \quad (20)$$

We use a different assumption for $P(c, s)$ depending on which fit we are doing (h^{all} or h^{vox}). For fitting $h^{\text{all}}(v)$, we consider all values of c, s equally likely:

$$P^{\text{all}}(c, s) = p_6 \quad (21)$$

For fitting h^{vox} , the means and standard deviations, c, s , are fixed at c^0, s^0 (the values determined by the earlier fit to the entire data set):

$$P^{\text{vox}}(c, s) = \delta(c - c^0, s - s^0) \quad (22)$$

For a small region, we assume that the mean noise vector, \bar{N} , has normal distribution with standard deviation σ :

$$P^{\text{vox}}(\bar{N}) = p_2 e^{-\frac{1}{2} \sum_{i=1}^{n_v} \left(\frac{\bar{N}_i}{\sigma_i}\right)^2} \quad (23)$$

For a large region, the mean noise vector, \bar{N} , should be very close to zero; hence, $P^{\text{all}}(\bar{N})$ will be a delta function centered at $\bar{N} = 0$.

C.2 Likelihood.

We approximate the likelihood, $P(h|\alpha, c, s, \bar{N})$, by analogy to a discrete normal distribution. We define $q(v)$ to measure the difference between the “expected” or “mean” histogram for particular α, c, s, \bar{N} and a given histogram $h(v)$:

$$q(v; \alpha, c, s, \bar{N}) = h(v - \bar{N}) - \sum_{j=1}^{n_f} \alpha_j f_j(v; c, s) \quad (24)$$

Now we create a normal-distribution-like function. $w(v)$ is analogous to the standard deviation of q at each point of feature space:

$$P(h|\alpha, c, s, \bar{N}) = p_3 e^{-\frac{1}{2} \int \left(\frac{q(v; \alpha, c, s, \bar{N})}{w(v)}\right)^2 dv} \quad (25)$$

C.3 Global likelihood.

Note that the denominator of Equation 18 is a constant normalization of the numerator:

$$\begin{aligned} P(h) &= \int P(\hat{\alpha}, \hat{c}, \hat{s}, \hat{N}) P(h|\hat{\alpha}, \hat{c}, \hat{s}, \hat{N}) d\hat{\alpha} d\hat{c} d\hat{s} d\hat{N} \\ &= p_4 \end{aligned} \quad (26)$$

C.4 Assembly

Using the approximations discussed above, we arrive at the following expression for the posterior probability:

$$\begin{aligned} P(\alpha, c, s, \bar{N}|h) &= \\ & p_5 P(\alpha) P(c, s) \exp\left(-\frac{1}{2} \sum_{i=1}^{n_v} \left(\frac{\bar{N}_i}{\sigma_i}\right)^2\right) \\ & \exp\left(-\frac{1}{2} \int \left(\frac{q(v; \alpha, c, s, \bar{N})}{w(v)}\right)^2 dv\right) \end{aligned} \quad (28)$$

For fitting h^{all} , the mean noise is assumed to be zero, so maximizing equation 28 is equivalent to minimizing \mathcal{E}^{all} to find the free parameters $(\alpha^{\text{all}}, c, s)$:

$$\mathcal{E}^{\text{all}}(\alpha^{\text{all}}, c, s) = \frac{1}{2} \int \left(\frac{q(v; \alpha^{\text{all}}, c, s)}{w(v)} \right)^2 dv \quad (29)$$

subject to $P(\alpha^{\text{all}}) \neq 0$. Because both $P(\alpha^{\text{all}})$ and $P^{\text{all}}(c, s)$ are constant valued in that region, they are not included.

For fitting h^{vox} , the parameters c and s are fixed, so maximizing equation 28 is equivalent to minimizing \mathcal{E}^{vox} to find the free parameters $(\alpha^{\text{vox}}, \bar{N})$:

$$\mathcal{E}^{\text{vox}}(\alpha^{\text{vox}}, \bar{N}) = \frac{1}{2} \sum_{i=1}^{n_v} \left(\frac{\bar{N}_i}{\sigma_i} \right)^2 + \frac{1}{2} \int \left(\frac{q(v; \alpha^{\text{vox}}, \bar{N})}{w(v)} \right)^2 dv \quad (30)$$

subject to $P(\alpha^{\text{vox}}) \neq 0$.

As stated in Equation 6, Section VI, Equation 30 is minimized to estimate relative material volumes, α^{vox} , and the mean noise vector, \bar{N} .

REFERENCES

- [1] Marc Levoy, "Display of surfaces from volume data," *IEEE Computer Graphics and Applications*, vol. 8, no. 3, pp. 29–37, May 1988.
- [2] William E. Lorensen and Harvey E. Cline, "Marching cubes: A high resolution 3D surface construction algorithm," in *Computer Graphics (SIGGRAPH '87 Proceedings)*, Maureen C. Stone, Ed., July 1987, vol. 21, pp. 163–169.
- [3] T. J. Loredo, "From Laplace to supernova SN1987A: Bayesian inference in astrophysics," in *Maximum Entropy and Bayesian Methods*, P. Fougere, Ed. Kluwer Academic Publishers, Denmark, 1989.
- [4] Alan V. Oppenheim, Alan S. Willsky, and Ian T. Young, *Signals and Systems*, Prentice-Hall, Inc., New Jersey, 1983.
- [5] Michael W. Vannier, Robert L. Butterfield, Douglas Jordan, William A. Murphy, Robert G. Levitt, and Mokhtar Gado, "Multispectral analysis of magnetic resonance images," *Radiology*, vol. 154, pp. 221–224, 1985.
- [6] Michael W. Vannier, Christopher M. Speidel, and Douglas L. Rickman, "Magnetic resonance imaging multispectral tissue classification," in *Proc. Neural Information Processing Systems (NIPS)*, Aug. 1988.
- [7] Harvey E. Cline, William E. Lorensen, Ron Kikinis, and Ferenc Jolesz, "Three-dimensional segmentation of MR images of the head using probability and connectivity," *Journal of Computer Assisted Tomography*, vol. 14, no. 6, pp. 1037–1045, Nov., Dec. 1990.
- [8] Richard P. Duda and Peter E. Hart, *Pattern Classification and Scene Analysis*, John Wiley and Sons, New York, 1973.
- [9] L. P. Clarke, R. P. Velthuizen, M. A. Camacho, J. J. Neine, M. Vaidyanathan, L. O. Hall, R. W. Thatcher, and M. L. Silbiger, "MRI segmentation: Methods and applications," *Magnetic Resonance Imaging*, vol. 13, no. 3, pp. 343–368, 1995.
- [10] Marc Joliot and Bernard M. Mazoyer, "Three-dimensional segmentation and interpolation of magnetic resonance brain images," *IEEE Transactions on Medical Imaging*, vol. 12, no. 2, pp. 269–277, 1993.
- [11] Robert A. Drebin, Loren Carpenter, and Pat Hanrahan, "Volume rendering," in *Computer Graphics (SIGGRAPH '88 Proceedings)*, John Dill, Ed., Aug. 1988, vol. 22, pp. 65–74.
- [12] Joe P. Windham, Mahmoud A. Abd-Allah, David A. Reimann, Jerry W. Froelich, and Allan M. Hagggar, "Eigenimage filtering in MR imaging," *Journal of Computer Assisted Tomography*, vol. 12, no. 1, pp. 1–9, 1988.
- [13] Yi-Hsuan Kao, James A. Sorenson, and Stefan S. Winkler, "MR image segmentation using vector decomposition and probability techniques: A general model and its application to dual-echo images," *Magnetic Resonance in Medicine*, vol. 35, pp. 114–125, 1996.
- [14] H. S. Choi, D. R. Haynor, and Y. M. Kim, "Partial volume tissue classification of multichannel magnetic resonance images — a mixel model," *IEEE Transactions on Medical Imaging*, vol. 10, no. 3, pp. 395–407, 1991.
- [15] Derek R. Ney, Elliot K. Fishman, Donna Magid, and Robert A. Drebin, "Volumetric rendering of computed tomography data: Principles and techniques," *IEEE Computer Graphics and Applications*, vol. 10, no. 2, pp. 24–32, Mar. 1990.
- [16] Peter Santiago and H. Donald Gage, "Quantification of MR brain images by mixture density and partial volume modeling," *IEEE Transactions on Medical Imaging*, vol. 12, no. 3, pp. 566–574, september 1993.
- [17] Zhenyu Wu, Hsiao-Wen Chung, and Felix W. Wehrli, "A bayesian approach to subvoxel tissue classification in NMR microscopic images of trabecular bone," *Journal of Computer Assisted Tomography*, vol. 12, no. 1, pp. 1–9, 1988.
- [18] David H. Laidlaw, Alan H. Barr, and Russell E. Jacobs, "Goal-directed brain micro-imaging," in *Neuroinformatics: An Overview of the Human Brain Project*, Steven H. Koslow and Michael F. Huerta, Eds., vol. 1, chapter 6. Kluwer, Feb. 1997.
- [19] David H. Laidlaw, *Geometric Model Extraction from Magnetic Resonance Volume Data*, Ph.D. thesis, California Institute of Technology, 1995.
- [20] Pratik R. Ghosh, David H. Laidlaw, Kurt W. Fleischer, Alan H. Barr, and Russell E. Jacobs, "Pure phase-encoded MRI and classification of solids," *IEEE Transactions on Medical Imaging*, vol. 14, no. 3, pp. 616–620, 1995.
- [21] David H. Laidlaw, "Tissue classification of magnetic resonance volume data," Master's thesis, California Institute of Technology, 1992.
- [22] Richard Bartels, John Beatty, and Brian Barsky, *An Introduction to Splines for Use in Computer Graphics and Geometric Modeling*, Morgan Kaufmann Publishers, Palo Alto, CA, 1987.
- [23] H. S. Choi, D. R. Haynor, and Y. M. Kim, "Multivariate tissue classification of MRI images for 3-d volume reconstruction — a statistical approach," *SPIE Medical Imaging III: Image Processing*, vol. 1092, pp. 183–193, 1989.
- [24] NAG, *NAG Fortran Library*, Numerical Algorithms Group, 1400 Opus Place, Suite 200, Downers Grove, Illinois 60515, 1993.