

A Nonparametric Method for Automatic Correction of Intensity Nonuniformity in MRI Data

John G. Sled,* Alex P. Zijdenbos, *Member, IEEE*, and Alan C. Evans

Abstract—A novel approach to correcting for intensity nonuniformity in magnetic resonance (MR) data is described that achieves high performance without requiring a model of the tissue classes present. The method has the advantage that it can be applied at an early stage in an automated data analysis, before a tissue model is available. Described as nonparametric nonuniform intensity normalization (N3), the method is independent of pulse sequence and insensitive to pathological data that might otherwise violate model assumptions. To eliminate the dependence of the field estimate on anatomy, an iterative approach is employed to estimate both the multiplicative bias field and the distribution of the true tissue intensities. The performance of this method is evaluated using both real and simulated MR data.

Index Terms—Intensity nonuniformity, magnetic resonance imaging, RF field inhomogeneity, shading artifact.

I. INTRODUCTION

MAGNETIC RESONANCE (MR) signal intensity measured from homogeneous tissue is seldom uniform; rather it varies smoothly across an image. This intensity nonuniformity is usually attributed to poor radio frequency (RF) coil uniformity, gradient-driven eddy currents, and patient anatomy both inside and outside the field of view. Although these 10%–20% intensity variations have little impact on visual diagnosis, the performance of automatic segmentation techniques which assume homogeneity of intensity within each class can be significantly degraded. A robust, automatic, and inexpensive means of correcting for this artifact is essential for such automatic processing techniques to be accurate in labeling each voxel with a tissue type. Furthermore, correcting for intensity nonuniformity may benefit quantitative measurements such as those used in tissue metabolite studies.

In considering MR intensity nonuniformity it is important to distinguish between the rapid interslice variations sometimes observed with two-dimensional (2-D) multislice sequences and the smooth intensity variations present in most acquisitions including those using three-dimensional (3-D) sequences. The former can be dealt with by methods which normalize the

intensities of individual slices [1]–[5]. Correction for the latter is the subject of this paper.

Besides being smoothly varying within homogeneous regions, the model for intensity nonuniformity is not easily characterized. Nonuniformity due to the frequency response of the receiver and spatial sensitivity of the unloaded RF coils is systematic and can in principle be corrected for by regular calibration or theoretical modeling [6], [7]. However, nonuniformity due to induced currents and the spatial inhomogeneity of the excitation field depend on the geometry and electrical properties of the subject as well as the pulse sequence and coil polarization [8].

Stollberger and Wach [9] have described a method for measuring the RF excitation field *in vivo* based on the change in nonuniformity with excitation flip angle. A complementary method for measuring the static uniformity of the reception coil *in vivo*, by imaging with a second reception coil, is described in [10]. Given an *in vivo* map of the RF excitation and reception fields one can correct for intensity nonuniformity using the Bloch equations [11]. Although the extended scan time of the former methods makes them impractical for clinical use, faster techniques such as echo-planar imaging (EPI) can be used to more rapidly acquire RF field maps [12]. However, EPI hardware is not always available and the need to collect this additional data precludes retrospective analysis.

As an alternative to improving intensity uniformity during data acquisition, a number of post-processing methods to compensate for nonuniformity have been proposed [13]–[18]. However, the need for expert supervision, in choosing a set of sample voxels expected to have similar intensities for example, has prevented their widespread use. Robust and automatic, the strength of the nonuniformity correction method described here is that it requires neither extended scan time nor expert supervision. In this report we describe the N3 method and present validation results using both simulated and real volumetric magnetic resonance imaging (MRI) data.

II. THEORY

A. Nonuniformity Model

The problem of correcting for intensity nonuniformity is greatly simplified if it is modeled as a smooth multiplicative field. This model has been widely used [14]–[16], [18] and is consistent with the multiplicative nonuniformity arising from variations in the sensitivity of the reception coil and to a lesser extent with the nonuniformity due to induced currents and nonuniform excitation.

Manuscript received December 27, 1996; revised October 28, 1997. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was Z.-P. Liang. Asterisk indicates corresponding author.

J. G. Sled is with the McConnell Brain Imaging Centre, Montréal Neurological Institute and McGill University, 3801 University Street, Montréal, P.Q., H3A 2B4 Canada (e-mail: jgsled@bic.mni.mcgill.ca).

A. P. Zijdenbos and A. C. Evans are with the McConnell Brain Imaging Centre, Montréal Neurological Institute and McGill University, Montréal, P.Q., H3A 2B4 Canada

Publisher Item Identifier S 0278-0062(98)03151-6.




Fig. 1. Probability densities for fields estimated from white matter on MR scans of 12 individuals taken with 12 different MR scanners. Note that these are the distributions of the intensities f rather than log intensities \hat{f} . The T1 scan is a gradient-echo 3-D acquisition with TE = 11 ms and TR = 35 ms, while the proton density (PD) and T2 scans are two echoes TE = 30 ms and TE = 80 ms of a multislice spin echo acquisition with TR = 3 s. (a) T1-weighted, (b) T2-weighted, and (c) PD-weighted.

Consider the following model of image formation in MR:

$$v(\mathbf{x}) = u(\mathbf{x})f(\mathbf{x}) + n(\mathbf{x}) \quad (1)$$

where at location \mathbf{x} , v is the measured signal, u is the true signal emitted by the tissue, f is an unknown smoothly varying bias field, and n is white Gaussian noise assumed to be independent of u . The problem of compensating for intensity nonuniformity is the task of estimating f . The combination of additive and multiplicative interference makes this task difficult.

Consider a noise-free case in which the true intensities u at each voxel location \mathbf{x} are independent identically distributed random variables. Using the notation $\hat{u}(\mathbf{x}) = \log[u(\mathbf{x})]$ the image formation model becomes additive

$$\hat{v}(\mathbf{x}) = \hat{u}(\mathbf{x}) + \hat{f}(\mathbf{x}). \quad (2)$$

Consider the distribution of values that \hat{f} takes over the region of interest (ROI) to be the probability distribution of a random variable. For example, if \hat{f} is a linearly increasing field aligned on a rectangular region then \hat{f} will have a uniform distribution.

Let U , V , and F be the probability densities of \hat{u} , \hat{v} , and \hat{f} , respectively. Making the approximation that \hat{u} and \hat{f} are independent or uncorrelated random variables, the distribution of their sum is found by convolution as follows (for details see [19]):

$$V(\hat{v}) = F(\hat{v}) * U(\hat{v}) = \int F(\hat{v} - \hat{f})U(\hat{f})d\hat{f}. \quad (3)$$

The nonuniformity distribution F can be viewed as blurring the intensity distribution U .

B. Correction Strategy

From a signal processing perspective, the blurring due to the field reduces the high frequency components of U . The task of correcting for intensity nonuniformity is that of restoring the frequency content of U . Since the shape of the blurring kernel F is not known, it is not clear what frequency components of U need to be restored to get from the observed distribution V to the true distribution U . However, since the nonuniformity field \hat{f} is restricted to be smooth and slowly varying, the space of possible distributions U corresponding to a given

distribution V is small enough that the problem becomes tractable. Our approach to correcting for nonuniformity is to find the smooth, slowly varying, multiplicative field that maximizes the frequency content of U .

As evidence of the simple form of the distribution F , consider the distributions shown in Fig. 1. These have been derived from fields fitted to manually labeled regions of white matter on 12 individuals. Each individual was scanned using the same pulse sequences but on a different MR machine (data collected as part of multicenter clinical trial). Included in the 12 are machines made by Philips, Siemens, and GE. As shown, the large scale features of F vary little between scans. In particular, F is typically unimodal or at least well approximated by a unimodal distribution. These results also suggest that the full width at half maximum (FWHM) of the distribution F lies between 0.1 and 0.4 for typical brain scans.

Returning to the optimization criterion, one could in principle search through all possible fields \hat{f} to find the one that maximizes the high frequency content of U . However, there are two problems with this approach: the search space of all 3-D fields \hat{f} is extremely large; and spectral estimates and related measures such as entropy are difficult to compute with sufficient accuracy to detect subtle changes in U .

Our approach is to propose a distribution for U by sharpening the distribution V , and then to estimate a corresponding smooth field which produces a distribution U close to the one proposed. While searching through the space of all distributions U may seem no more tractable than searching through the space of all fields \hat{f} , there is an important difference in that we can take advantage of the simple form of the distribution F . Suppose that the distribution of F is Gaussian. Then we need only search the space of all distributions U corresponding to Gaussian distributed F having zero mean and given variance. In this way the space of all distributions U is collapsed down to a single dimension, the width of the F distribution.

In practice, the field distribution F is only approximately Gaussian and some of our assumptions, such as zero noise, are violated. To contend with these difficulties, we take an incremental approach to estimating the true distribution of intensities U and corresponding field \hat{f} . Since any Gaussian distribution can be decomposed into a convolution of narrower Gaussian distributions, the space of all U distributions

corresponding to Gaussian distributed F can be searched incrementally by deconvolving narrow Gaussian distributions from subsequent estimates of U . The benefit of this approach is that between subsequent estimates of U , a corresponding smooth field \hat{f} is estimated. The constraint that the field be smooth changes the shape of the proposed distribution U to one that is consistent with the field. These perturbations of U perturb F from its Gaussian shape and compensate for the distortion of V caused by noise and other factors. The iterative process can be viewed as traveling in the space of all U distributions along a path corresponding to smooth fields \hat{f} with increasingly wider distributions. These iterations proceed until no further changes in \hat{f} or U result from deconvolving narrow Gaussian distributions from V .

C. Field Estimation

Further theory is presented here to explain the process of proposing distributions for U and estimating corresponding fields. For notational simplicity, we will assume that the true distribution of intensities U can be arrived at in a single iteration by deconvolving a distribution F , which is Gaussian, from V . The full iterative description of the method is left for a subsequent section.

Given the distribution U , the method of estimating the corresponding field is as follows. For a measurement \hat{v} at some location \mathbf{x} , \hat{u} is estimated using the distributions U and F . Since the choice of the location \mathbf{x} is arbitrary, the measurement \hat{v} can be treated as a random sample from the distribution V . The expected value of \hat{u} given a measurement \hat{v} is as follows:

$$E[\hat{u}|\hat{v}] = \int_{-\infty}^{\infty} \hat{u}p(\hat{u}|\hat{v}) d\hat{u} \quad (4)$$

$$= \int_{-\infty}^{\infty} \hat{u} \frac{p(\hat{u}, \hat{v})}{p(\hat{v})} d\hat{u}. \quad (5)$$

Writing $p(\hat{v})$ as $V(\hat{v})$ and using (3), we have

$$E[\hat{u}|\hat{v}] = \frac{1}{V(\hat{v})} \int_{-\infty}^{\infty} \hat{u}p(\hat{u}, \hat{v}) d\hat{u} \quad (6)$$

$$= \frac{1}{V(\hat{v})} \int_{-\infty}^{\infty} \hat{u}p(\hat{u}, \hat{f}) d\hat{u} \quad (7)$$

$$= \frac{1}{V(\hat{v})} \int_{-\infty}^{\infty} \hat{u}p(\hat{u})p(\hat{f}) d\hat{u} \quad (8)$$

$$= \frac{\int_{-\infty}^{\infty} \hat{u}F(\hat{v} - \hat{u})U(\hat{u}) d\hat{u}}{\int_{-\infty}^{\infty} F(\hat{v} - \hat{u})U(\hat{u}) d\hat{u}}. \quad (9)$$

An estimate of \hat{f} can be obtained using the estimate of \hat{u} from (9) as follows:

$$\hat{f}_e(\hat{v}) = E[\hat{f}|\hat{v}] = \hat{v} - E[\hat{u}|\hat{v}] \quad (10)$$

where \hat{f}_e is an estimate of \hat{f} at location \mathbf{x} based on the single measurement of \hat{v} at \mathbf{x} . This estimate can be smoothed by the operator S to produce

$$\hat{f}_s(\hat{v}) = S\{\hat{f}_e(\hat{v})\} = S\{\hat{v} - E[\hat{u}|\hat{v}]\} \quad (11)$$

an estimate of \hat{f} based on all of the measurements in a neighborhood of \mathbf{x} . Smoothing is described in a subsequent section.

Given a distribution F and the measured distribution of intensities V , the distribution U can be estimated using a deconvolution filter as follows:

$$\tilde{G} = \frac{\tilde{F}^*}{|\tilde{F}|^2 + Z^2} \quad (12)$$

$$\tilde{U} \approx \tilde{G}\tilde{V} \quad (13)$$

where $*$ denotes complex conjugate, \tilde{F} is the Fourier transform of F , and Z is a constant term to limit the magnitude of \tilde{G} . This estimate of U is then used to estimate a corresponding field \hat{f} .

D. An Example in One Dimension

This process of estimating \hat{f} is illustrated in Fig. 2 for the one dimensional (1-D) case with \hat{u} as a square wave. The square wave can be viewed as two tissue classes corresponding to the low and high parts of the cycle. The log intensities \hat{u} , \hat{f} , and \hat{v} are shown in Fig. 2(a). The magnitude of the nonuniformity field is just large enough that the low intensity tissue on the left overlaps by roughly 10% with the intensity of the high-intensity tissue on the right.

We obtain the probability densities U , V , and F by computing the histogram of the signals \hat{u} , \hat{f} , and \hat{v} . These distributions are shown in Fig. 2(b)–(d), respectively. These histograms are interpreted as the probability distributions of the given signal. In practice, only V will be known while U and F have to be estimated. Shown in Fig. 2(e) is the Gaussian kernel used in place of the actual distribution of F in (9). Although its exact width is arbitrary, note that this distribution is narrower than F .

The distributions U and the Gaussian kernel F are used to compute from (10) the mapping, shown in Fig. 2(f), which maps measured intensities \hat{v} to field estimates \hat{f}_e . This mapping when applied to the measured signal produced the field estimate shown in Fig. 2(g). Note the sharp jumps in this estimate in the regions where the intensities of the upper and lower tissue classes overlap. This field estimate is smoothed to produce \hat{f}_s shown in Fig. 2(h). While the degree of smoothing is arbitrary, in filtering terms the smoothing filter should be chosen to have as small a bandwidth as possible yet still pass the nonuniformity field undistorted. In particular, it is not necessary to have a smoothing filter able to remove the sharp jumps in \hat{f}_e shown in Fig. 2(g) since these can be removed in subsequent iterations.

While the smooth field estimate after one iteration is not sufficient to completely remove the nonuniformity, it is enough of a correction that the tissue distributions no longer overlap. If we use the smooth estimate of the field \hat{f}_s to correct the measured signal \hat{v} and repeat the entire estimation process again, we arrive at a second estimate of the nonuniformity field that is indistinguishable from the true field. The second estimates of \hat{f}_e and \hat{f}_s are shown in Fig. 2(g) and (h).

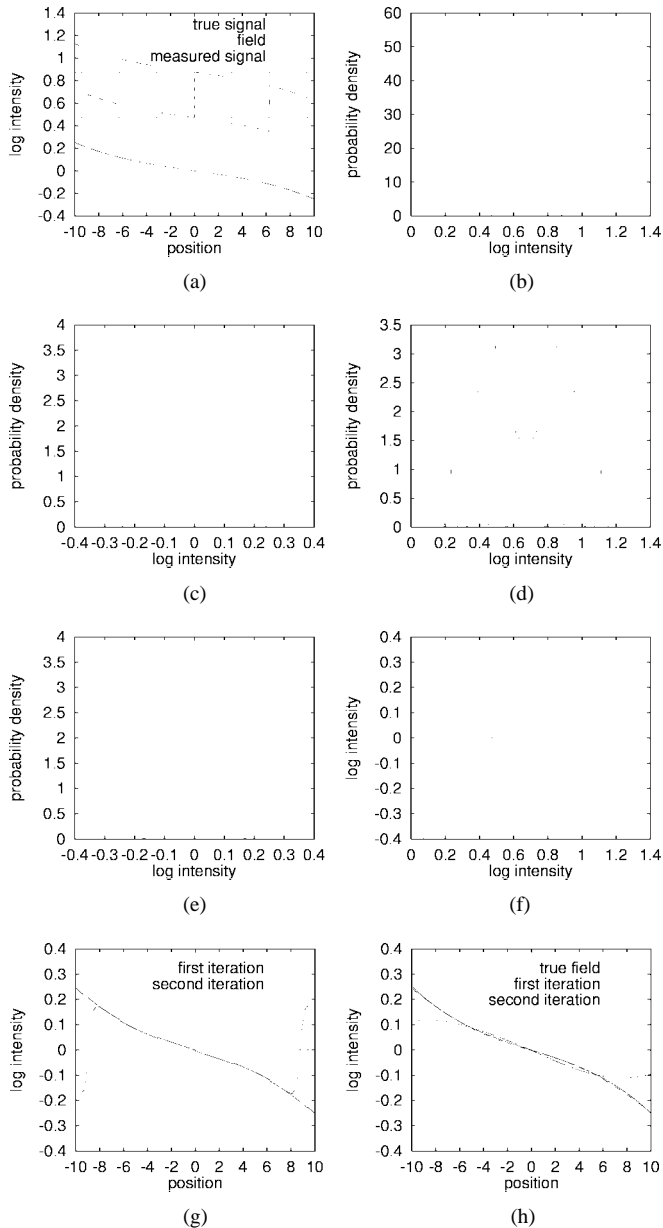


Fig. 2. An example of nonparametric correction of a 1-D signal. See Section II-D for detailed explanation: (a) $\hat{u}(x)$, $\hat{f}(x)$, and $\hat{v}(x)$, (b) $U(\hat{u})$, (c) $F(\hat{f})$, (d) $V(\hat{v})$, (e) a Gaussian kernel, (f) $\hat{f}_e(\hat{v}) = \hat{v} - E[\hat{u}|\hat{v}]$, (g) $\hat{f}_e[\hat{v}(x)]$, and (h) $\hat{f}(x)$ and $\hat{f}_s(x)$.

E. An Example Using a Simulated MR Volume

As an illustration of the field estimation process in 3-D, consider the simulation of a T1-weighted MR scan shown in Fig. 3(a). This simulation, described in more detail in Section III-B, is based on a three tissue model of the brain and incorporates intensity nonuniformity, noise, and partial volume effects. The measured intensity distribution V is shown in Fig. 3(b). The mapping between measured intensity and field estimate for this volume based on the distributions of V and a Gaussian distribution F with full-width-half-maximum (FWHM) of 0.15 is shown in Fig. 3(d). This mapping is applied to produce the volume in Fig. 3(c). Measurement noise causes this estimate to be noisy. Smoothing this estimate produces the field shown in Fig. 3(e). For comparison, the




Fig. 3. Field estimates for a simulated MR volume. (a) A slice from a simulated MR volume. (b) The histogram of the volume shown in (a). This is considered an estimate of the distribution V . (c) Bias field estimates f_e created by applying the mapping in (d) to the image in (a). (d) The mapping between image intensity and bias field estimate, derived from the histogram in (b). (e) A smoothed estimate f_s of the bias field created by smoothing the volume shown in (c). (f) The actual bias field present in the volume shown in (a). See Section II-E for further details.

actual field imposed on the data during simulation is shown in Fig. 3(f). Although the estimated field is much smaller in magnitude than the true field, the shape is similar. Experiments, described later, show that with subsequent iterations the field estimate will grow to narrow this difference.

F. Implementation Details

A flow chart describing the nonparametric correction process is shown in Fig. 4. Besides the processing steps described previously, there are a number of steps needed for practical implementation of the algorithm.

The first step shown in Fig. 4, “identify foreground,” is to segment and remove empty regions from the volume. Besides the numerical problems associated with transforming values near zero to the log domain, these background regions provide no information about the nonuniformity field. Since the accuracy of this segmentation is not critical, the foreground is determined using a simple threshold chosen automatically by analyzing the histogram of the volume [20].

Another consideration in implementing the N3 approach is measuring the distribution V from the unprocessed MR data. For simplicity, a histogram with equal-size bins and a

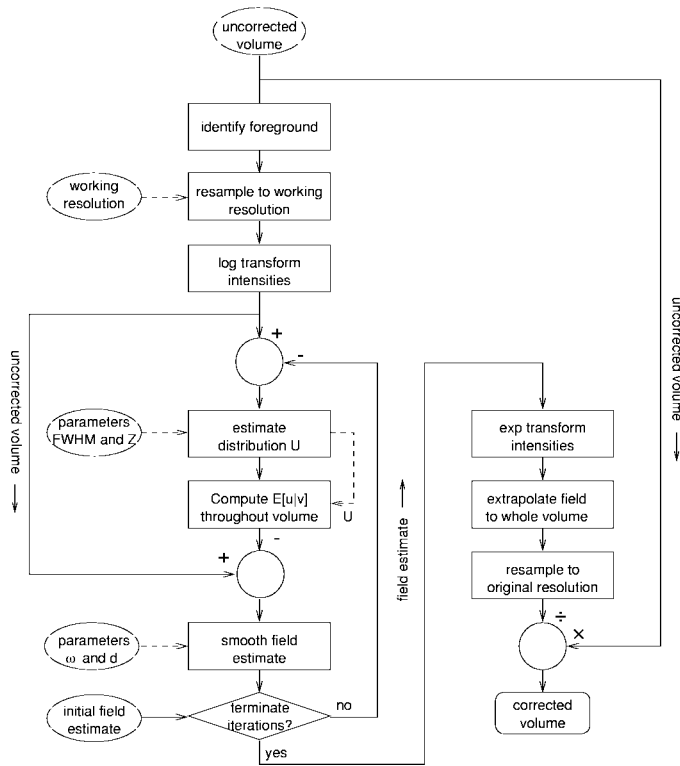


Fig. 4. A flow chart describing the N3 method. Ellipses represent user-selected parameters and priors. Rectangles are processing steps. The flow of volumetric data is represented by solid lines and of other data by dashed lines. Circles perform arithmetic operations on a voxel by voxel basis. The result of the process is a corrected volume.

triangular Parzen window [21] is used to estimate V . Given a set of N measurements $\hat{v}(\mathbf{x}_i)$ and locations \mathbf{x}_i , V is estimated as follows:

$$V(\hat{v}_j) = \frac{1}{N} \sum_{i=1}^N \frac{1}{h} \varphi \left[\frac{\hat{v}_j - \hat{v}(\mathbf{x}_i)}{h} \right] \quad (14)$$

$$\varphi(s) = \begin{cases} 1 - |s| & |s| < 1 \\ 0, & \text{elsewhere} \end{cases} \quad (15)$$

where \hat{v}_j are the centers of the bins and h is the distance between them. For a typical 20% bias field, the scale factor f ranges from 0.9 to 1.1 which corresponds to \hat{f} between 0.1 and 0.1. MR volumes generally have sufficient data to estimate V at a resolution h better than a tenth of this range.

Smoothing the nonuniformity field at full resolution is computationally expensive, so the MR data is subsampled without averaging to a lower resolution. Since the nonuniformity field is slowly varying, reducing a 1-mm isotropically sampled volume to 3 mm has a negligible effect on the field estimate and substantially accelerates computation. The processing of a volume with ten iterations of the N3 method is reduced from 4.5 h to 7 min of CPU time on an 200-MHz Intel Pentium workstation running Linux (floating point performance: 5.0 SPECfp95) by resampling to the coarser resolution. The final field estimate is resampled to the original resolution and used to correct the original volume.

The measure used to terminate the iterations is the coefficient of variation in the ratio between subsequent field

estimates, computed as follows:

$$e = \frac{\sigma\{r_n\}}{\mu\{r_n\}}, \quad n = 1 \dots N \quad (16)$$

where r_n is the ratio between subsequent field estimates at the n th location, σ denotes standard deviation, and μ denotes mean. This measure is chosen so as to be insensitive to global scale factors that may accumulate with iterations. Iteration is stopped when e drops below 0.001, typically after ten iterations.

G. Smoothing

The manner in which the field estimate is smoothed has a significant impact on the performance of the correction method. Smoothing is particularly challenging for this problem because the scale over which the field varies is comparable to the size of the region being smoothed. Conventional filtering techniques proved unsatisfactory for this application since boundary effects significantly degraded overall performance.

A computationally tractable approach to smoothing that performs well on bounded domains is to approximate the data by a linear combination of smooth basis functions. An attractive basis for this is the compactly supported spline known as a B spline. Spline approximation incorporating smoothness constraints is superior to filtering techniques in dealing with missing data since the behavior of a spline curve can be constrained even if there is insufficient data to support a basis function.

Details of computing a B spline approximation can be found in the Appendix. The smoothness of the approximation is determined by two parameters: ω , referred to as the smoothing parameter, and d , the distance between basis functions. Since splines are being used as a filter for this application, the smoothness of the approximation must be chosen rather than derived from the data [22], [23]. The relationship between the smoothness of the approximation and the smoothing parameter is nonlinear. However, since the normalization of the B spline has been chosen to eliminate the dependence of ω on scale and number of data points, ω can be fixed and the distance between basis functions varied instead.

III. EXPERIMENTS

A. Correcting Random Fields

As a first step in validating the N3 method, consider a case in which the measurements are independent and identically distributed random variables. Suppose the distribution of the true intensities U is as shown in Fig. 5(a). While this distribution is taken from a T1-weighted MR scan, the choice of U for the purpose of this example is arbitrary. The volume is a cube with 32 voxels on an edge and the nonuniformity field to be removed is a parabolic function aligned with the center of the volume. The corresponding distribution F of the nonuniformity field is shown in Fig. 5(b) along with the distribution V in Fig. 5(c).

A slice through the center of this idealized volume is shown in Fig. 5(d) and (f) before and after the nonuniformity field




Fig. 5. (a) Intensity distribution U for an idealized volume. Random samples from this distribution are used to create the volume in (d). (b) Distribution F for a parabolic nonuniformity field. (c) Distribution V corresponding to volume shown in (f). (d) A slice through a random volume. (e) and (g) A parabolic nonuniformity field. (f) The volume in (d) multiplied by the parabolic field in (e). (h) Estimated field. (i) Ratio of actual field to estimated field.

is applied. Although the volume itself appears noise-like, measurement noise n is taken to be zero for this experiment. We compared convergence rate and accuracy of the N3 method in removing the nonuniformity field for different widths of the estimated field distribution F . In addition, the N3 method is compared to simple filtering. The results of these and following experiments are described in the next section. The significance of this experiment is to demonstrate how the N3 method is able to take advantage of the structure of the distribution U as compared to conventional filtering.

The performance of a correction method is best evaluated by comparing the field estimate to the true field if it is known. For example, a field estimate made for the random volume of Fig. 5(f) is shown in Fig. 5(h) along with the actual field in Fig. 5(g). Besides differing subtly in shape, the two surfaces also differ by a multiplicative factor. This scale factor has no impact on the quality of the correction since MR intensity is relative. However, in quantifying the performance of the method this factor needs to be removed. Equation (16) can also be used for this purpose if r_n is taken as the ratio of estimated to actual field intensity at location \mathbf{x}_n . This is the coefficient of variation of the field shown in Fig. 5(i).

B. Simulated MR Volumes

Nonuniformity correction methods such as N3 can be validated with real data using subjective measures of image quality and by assessing the reduction of variability in tasks such




Fig. 6. Simulated T1-, T2-, and PD-weighted volumes. Nonuniform sensitivity of the reception coil has been simulated causing the intensity to drop off at the bottom right. The noise distribution has a standard deviation 3% of the mean intensity of white matter.

as segmentation. However, the large number of uncontrolled factors in such experiments confound attempts to evaluate and optimize performance. In particular, partial volume effects, true anatomical variability, and an unknown nonuniformity field prevent a sensitive analysis of methodological parameters. These technical issues are circumvented in our analysis by the use of an MR simulator which incorporates realistic models for noise and partial volume. Slices through simulated T1, T2, and PD volumes are shown in Fig. 6.

The anatomical model for the simulations is derived from high-quality T1-, T2-, and proton density (PD)-weighted scans formed from the average of 27, 12, and 12 scans, respectively [24], [25] of a normal individual. The three volumes were resampled into stereotaxic space with 1-mm isotropic sampling and corrected for intensity nonuniformity using the N3 method. The brain region of the T1-weighted scan was then segmented into the three tissue classes using a manually trained minimum distance [26] classifier. Among the results of several different classifiers, this segmentation was chosen and manually edited by a neuroanatomist to improve the classification of deep structures and brain stem. Partial volume regions were formed by eroding each of the tissue regions using a six neighbor structuring element. The proportions of each tissue class in the partial volume regions were taken from a segmentation of the T1 volume using a Bayesian classifier [21], [27].

The MR simulator is designed to produce volumes with an intensity histogram similar to that of real data, based on a discrete labeling of each voxel as either cerebrospinal fluid (CSF), grey matter, or white matter. However, representing an MR volume in terms of a few classes of homogeneous tissue is unsatisfactory for simulation since it does not reflect the variability seen in real data. Often the properties of a tissue class vary for different structures [28], which leads to a broadening of the peaks in the histogram. The approach taken to produce a simulation with a convincing histogram is to color each region with random noise having medium to low spatial frequency and an intensity distribution appropriate for the given tissue type. This noise can be viewed as “pseudo-anatomy,” representing spatial variations within tissue due to anatomy. By validating the correction algorithm on a number of realizations of this pseudo-anatomy, we avoid biasing the results toward any preferred shape of the bias field.

To complete the simulation, a multiplicative intensity nonuniformity field is imposed on the volume and Rician distributed noise is added throughout. Rician noise, typical




Fig. 7. Intrinsic intensity distributions for each tissue class in T1-, T2-, and PD-weighted simulations.

of that found in real MR images, is simulated by adding a complex Gaussian distributed random intensity to a voxel, then computing its absolute value.

The intensity distributions used to generate the pseudo-anatomy were created by computing the histogram of intensities within the eroded or pure tissue regions for each of the T1, T2, and PD averages. These distributions are shown in Fig. 7. We assume that these distributions reflect the range of intensities intrinsic to the tissue class uncorrupted by noise, intensity nonuniformity, and partial volume effects. Random fields having these intensity distributions were created by rejection sampling [19]. Each random field is first constructed at 10-mm resolution and subsequently resampled to the desired 1-mm resolution using tricubic interpolation.

C. Correcting Simulated Data

Although N3 does not require a model of the expected tissue distributions, there are still a number of model-independent parameters that need to be selected (see Fig. 4). These are the smoothing parameter ω , the distance between basis functions d , the noise term in the deconvolution filter Z , and the FWHM of the deconvolution kernel F . Experiments with simulated data show little dependence on Z , which was fixed arbitrarily to be 0.1 for the purpose of this analysis. The parameters ω and d both control smoothness, but the relationship between smoothness and ω is more complex than the one between smoothness and d . As a result, ω was fixed arbitrarily at 1.0. The remaining two parameters, FWHM and d , will be considered further.

Simulated MR volumes like those shown in Fig. 6 were used to evaluate the effect of FWHM on correction performance. For this analysis, two different nonuniformity fields were generated from combinations of linear, quadratic, and Gaussian terms. Both fields vary in magnitude by 20% within the brain volume; the second field has more curvature. Slices from these two fields are shown in Fig. 8. The first field is barely visible in the simulations shown in Fig. 6. The performance of the correction method is evaluated using (16) to compute the difference expressed as a percentage between the estimated field and that imposed explicitly during simulation.

In addition, an analysis of the effect of basis function distance d on correction performance, for a fixed FWHM, was done using the same simulated data. These results were




Fig. 8. Slices through two nonuniformity fields used to construct simulated MR scans. (a) Field #1. (b) Field #2.

compared against the results of fitting B splines to the imposed field directly.

D. Correcting Experimental Data

Correction of real MR scans is the final step in validating the N3 method. For this experiment, regions of pure grey and white matter were manually labeled on 12 sets of T1-, T2-, and PD-weighted scans acquired at 12 different sites. The volumes were transformed into stereotaxic space so that the same labeling may be applied to each of the three modalities. All 12 individuals are multiple sclerosis patients having a moderate number of white matter lesions.

For each volume, the coefficient of variation in white and grey matter intensity is computed before and after correction. Since this measure cannot distinguish between intensity nonuniformity, noise, and anatomical intensity variations, it is only suitable for showing that intensity nonuniformity is qualitatively reduced.

IV. RESULTS

A. Correcting Random Fields

Correcting for intensity nonuniformity in a random field illustrates that N3 is able to correct volumes that lack regions of contiguous tissue. The improvement in the field estimate with iterations for four different widths of the deconvolution kernel F is shown in Fig. 9. In all four cases the method is converging in the sense that the distance between subsequent field estimates, computed using (16), drops to zero. Note that the FWHM of the deconvolution kernel has an impact on both convergence rate and performance. There is a subtle improve-

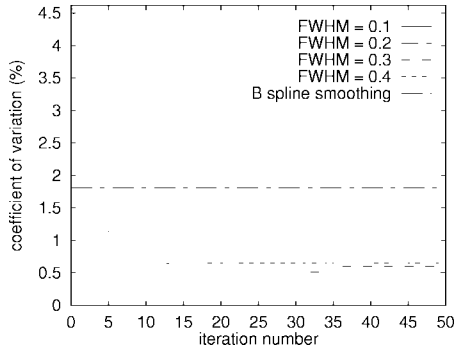


Fig. 9. Error measure versus iterations for four different estimates of the field distribution F . The FWHM of the field distributions F used in deconvolution are 0.1, 0.2, 0.3, and 0.4. The maximum coefficient of variation 4.5% corresponds to no correction. The horizontal line is the result of filtering the uncorrected data directly.

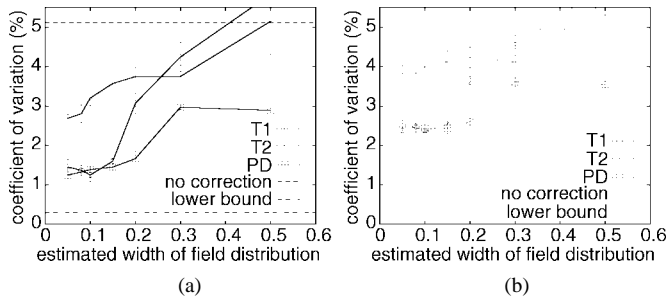


Fig. 10. Field estimation error on simulated data for a range of the FWHM parameter and two different nonuniformity fields. Lines are shown for the coefficients of variation corresponding to no correction for the two cases. Also shown is the lower bound on estimation error given by fitting splines to the nonuniformity field directly. The error bars are at plus and minus two standard deviations, based on ten realizations of the noise and pseudo-anatomy. (a) Field #1. (b) Field #2.

ment in accuracy and a significant decrease in convergence rate as the FWHM is decreased.

For comparison, the same B spline approximation that is used in the N3 method is applied to the uncorrected data directly. The smoothness of the two is comparable since the N3 method, unlike conventional iterative filtering, has been designed such that field smoothness does not accumulate with iterations. As expected, since the N3 method takes advantage of the structure of the distribution U , it outperformed simple smoothing in all four cases.

B. Correcting Simulated Data

The performance curves for ten simulated brains and a range of the FWHM parameter are plotted in Fig. 10 for the two different nonuniformity fields shown in Fig. 8 and fixed basis function distance $d = 200$ mm. A coefficient of variation below the lines corresponding to no correction in Fig. 10 indicates a reduction in intensity nonuniformity. Within most of the range of the FWHM shown, the N3 method substantially improves uniformity. In general, the performance improves as the FWHM parameter is decreased, a trend that is consistent with the assumption that the deconvolution kernel is narrow. However, for small values of the FWHM parameter convergence is slow, requiring up to 50 iterations to converge.

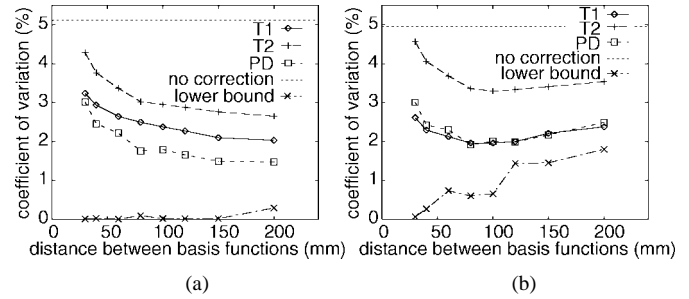


Fig. 11. Correction performance as a function of basis function distance on simulated T1, T2, and PD scans. The FWHM parameter for this analysis is 0.15. Also shown is the approximation error for fitting splines to the nonuniformity field directly. The discontinuities in this curve are caused by changes in the number of basis functions, which for efficiency is the minimum needed for full support at the given distance between basis functions. (a) Field #1. (b) Field #2.

TABLE I
MEAN COEFFICIENT OF VARIATION IN TISSUE
INTENSITY BEFORE AND AFTER N3 CORRECTION

| | | before | after |
|-------------|--------------|--------|-------|
| T1-weighted | white matter | 5.8% | 5.1% |
| | grey matter | 10.1% | 9.8% |
| T2-weighted | white matter | 10.0% | 9.0% |
| | grey matter | 15.4% | 14.3% |
| PD-weighted | white matter | 6.4% | 4.9% |
| | grey matter | 8.8% | 6.8% |

In choosing the FWHM parameter, one makes a tradeoff between accuracy and computation time.

It should be noted that the ideal result of zero estimation error is unattainable. In particular, the noise and pseudo-anatomy present in the simulations will have low frequency components indistinguishable from intensity nonuniformity. Furthermore, smoothness constraints imposed on the spline fitting operation preclude an exact match to the field. The lower bound imposed by spline approximation is also shown in Fig. 10. Noise and pseudo-anatomy make the true lower bound somewhat higher.

The results of correcting simulated brains for a range of basis function distances d are shown in Fig. 11. The greater curvature of the second field produces broad minima in the error curve. As in Fig. 10, the lower bound on the estimation error given by directly fitting splines to the nonuniformity field is also shown. Basis function distances of 100–200 mm are nearly optimal for these two cases. It should be noted that due to the need to suppress noise in the simulations, the optimal basis function distance is greater than that which can exactly match the nonuniformity field.

C. Correcting Experimental Data

The coefficient of variation was computed for white and grey matter intensity in each patient in each modality. The mean across patients is shown in Table I before and after correction for intensity nonuniformity. The statistical significance of these results was computed using a one sided Wilcoxon signed rank test for a paired difference experiment. All categories except grey matter in T1 showed a significant

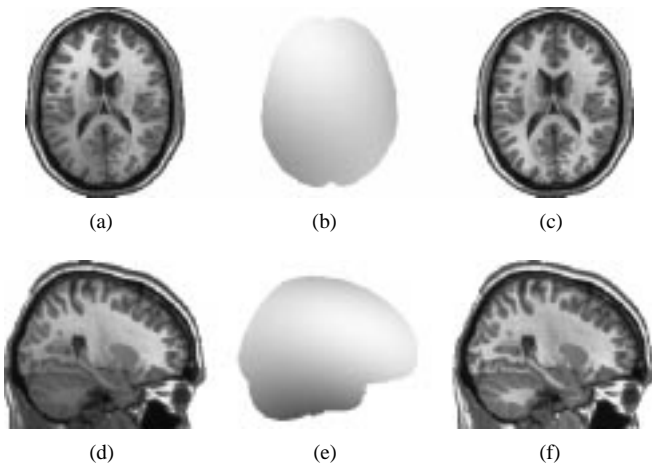


Fig. 12 Intensity nonuniformity correction of a T1 weighted 27-scan averaged gradient-echo MR scan: (a) and (d) transaxial and sagittal views of uncorrected data; (b) and (e) nonuniformity field estimated by the N3 method; (c) and (f) corrected data.




Fig. 13. Intensity nonuniformity correction of a surface coil MR scan: (a) and (d) transaxial and sagittal views of uncorrected data; (b) and (e) nonuniformity field estimated by the N3 method; (c) and (f) corrected data.

reduction in variability at the 99% confidence level. The basis function distance and FWHM parameters have been fixed for this experiment at 200 mm and 0.15, respectively.

As a final result, two volumes are shown before and after correction. The first, the T1-weighted 27-scan averaged MRI shown in Fig. 12, is useful for illustration since intensity nonuniformity that would normally be obscured by noise is clearly visible. Note how the intensity of white matter in the cerebellum is raised to that of the rest of the volume.

The second example, the surface coil scan shown in Fig. 13, illustrates the performance of the algorithm under extreme conditions, where nonuniformity is clearly visible. Since the signal intensity drops below the level of noise near the front of the head in this image, the foreground was segmented manually rather than by the automated method described previously. It should be noted that in regions where the signal is enhanced, the noise is enhanced as well. Consequently the noise level will vary substantially in a surface coil image after correction. Computing as before the coefficient of variation in manually labeled regions of white matter shows a large

reduction in white matter variability from 48% to 9.7% after correction.

V. DISCUSSION

We have described a new method (N3) of correcting for intensity nonuniformity that does not rely on a parametric model of tissue intensities, nor on segmenting a volume into contiguous regions. The former has been demonstrated by the N3 method's ability to correct for intensity nonuniformity in MR data irrespective of the pulse sequence and without prior training. The latter has been demonstrated by the method's ability to correct for nonuniformity using the structure of a random field, even when the nonuniformity is not clearly visible. N3 is a fully automatic iterative method that operates on 3-D volumetric data sets.

The algorithm requires that two parameters be selected: d , the smoothness of the estimated field; and FWHM, the width of the deconvolution kernel. While a scheme could be devised to choose the smoothness parameter d automatically, our experiments show little dependence on this parameter (see Fig. 11) and it was fixed at 200 mm. The choice of the FWHM parameter is more sensitive as it determines the tradeoff between accuracy and convergence rate. The results of Fig. 10 suggest that a single value of this parameter provides uniform performance across pulse sequences and between subjects, a result that is consistent with our qualitative assessment of several real MR scans. The FWHM parameter has been fixed at 0.15 throughout.

In addition to choosing these parameters, the foreground needs to be segmented from background. While this is trivial and automatic for volumetric scans, the surface coil image in Fig. 13 is more difficult to segment as the signal drops below the noise level at the front of the head. While manual segmentation is satisfactory, perhaps a more elegant solution would be to extend the method to weight the contribution of each region based on its signal to noise ratio.

A difficulty with iterative optimization methods such as N3 is that one is never certain whether the solution found is the global or a local minimum in the objective function. In our experience with random fields, other kinds of simulated data, and real data, the method converges to the same solution regardless of the initial field estimate. This behavior suggests that the optimization is not prone to local minima.

The optimization criterion for the N3 method is to find the smooth, slowly varying, multiplicative field that maximizes the frequency content of the distribution of tissue intensities. We assume that the field that satisfies this criterion is a good approximation of the true nonuniformity field. While all of the results presented support this hypothesis, in principle it need not be true. In particular, if the object is itself a smooth and slowly varying field then the correction field that maximizes the frequency content of the intensity distribution will also remove the natural variations from the object. For example, if the intensity of brain tissue varies slowly in intensity from anterior to posterior then these variations would be removed by the algorithm along with variations due to intensity nonuniformity. This is a limitation of all retrospective nonuniformity

correction methods as the two sources of variation cannot be distinguished based on the image data alone.

For comparison, consider the approach described by Zijdenbos *et al.* [5] and Dawant, *et al.* [14]. This method estimates intensity nonuniformity by fitting splines to the white matter of the brain. It is heavily model dependent in that it assumes that regions of pure white matter, unaffected by partial volume, can be correctly identified. Similarly, consider the expectation maximization (EM) approach proposed by Wells *et al.* [18]. In this method, tissue is modeled as belonging to one of several tissues classes. For example, a brain could be modeled as white matter, grey matter, and CSF. At each iteration the true intensity is computed at each voxel based on the proportion of each tissue class it is estimated to contain. The residual nonuniformity field is then computed and smoothed over the ROI.

The classifiers upon which these methods rely need to be trained. Wells argues that once trained, the same classifier can be used across subjects. However, the classifier still needs to be retrained for different scan parameters. Another difficulty with the EM method is that intensities that fall outside the tissue model of the classifier receive excessively large corrections that may distort the field estimate. Guillemaud and Brady [29] address this by introducing an additional class for tissues outside the model.

A different approach described by Meyer *et al.*, [16] avoids the issue of building a tissue intensity model by making use of contiguous regions irrespective of their tissue type. A smooth field is globally fit to the log intensities of all regions allowing for a constant offset in each regions to reflect its tissue type. Since this method can only make use of contiguous regions, the correction may suffer in areas of the brain such as the cerebellum where such regions are not present. A second difficulty arises if contiguous regions are mistakenly broken into smaller regions creating additional undesired degrees of freedom. If the relative sizes of these regions differ significantly the problem of fitting a smooth field can become ill conditioned. In addition, the smaller the regions are, the less information they contribute to the field estimate. In the limiting case an isolated voxel contributes no information. We expect this method to perform best on images with large scale structure. In this respect, Meyer's method and the N3 method are complementary since the N3 method is well suited to random fields.

All of these methods are based on the assumption of a smooth multiplicative field model. While this model is accurate for nonuniform sensitivity of a reception coil, it is only an approximation for excitation field nonuniformity since the multiplicative nonuniformity field corresponding to a smooth excitation field may have small discontinuities caused by differing relaxation times between regions of different tissue. Many pulse sequences, including standard spin echo sequences, are designed to be insensitive to low levels of excitation field inhomogeneity. In addition, excitation is often done with a scanner's body coil, which is larger and more uniform than the head coil used for reception. There are, however, additional field variations, beyond those measured for an unloaded coil, caused by interaction with the subject in the form of induced currents and standing wave effects. Although these in-

teractions affect both reception and excitation uniformity, there is some evidence to suggest these interactions are negligible in the human head at 1.5-T using quadrature coils [3], [30].

VI. CONCLUSIONS

An iterative method for correction of intensity nonuniformity in MR volumes has been described that avoids some of the restrictive model assumptions that plague other methods. In particular, this N3 method does not require a model of the tissue intensities in terms of discrete tissue classes, nor does it rely on a segmentation of the volume into homogeneous regions. Instead, a nonparametric model of the tissue intensities is derived directly from the data. N3 requires only two parameters to be selected: one controlling the smoothness of the estimated nonuniformity, the other controlling the tradeoff between accuracy and convergence rate. Our experiments show that both of these parameters can be chosen to provide uniform performance independent of pulse sequence, subject, and field shape. This is a considerable advantage in automated data analysis as the method can be applied at an early stage, without prior knowledge of the data. For instance, no special precautions need be taken for pathological data that might otherwise violate tissue model assumptions.

Experiments with simulated MR volumes show that the N3 method substantially reduces intensity nonuniformity in a variety of cases. Furthermore, real data corrected by the method is visually uniform and shows a statistically significant reduction in tissue intensity variation. Robust, fully automatic, and requiring little domain specific knowledge, N3 is attractive as a preprocessing step for a variety of MRI analysis applications.

APPENDIX

The tensor cubic B spline approximation of a function is given by

$$p(\mathbf{x}) = \sum_{i=1}^{M_x} \sum_{j=1}^{M_y} \sum_{k=1}^{M_z} \theta_{ijk} B_i(x) B_j(y) B_k(z) \quad (17)$$

where $\mathbf{x} = [x \ y \ z]$ and B is a 1-D cubic B spline in the variable x , y , or z . For example, the B spline for x is given by

$$B_i(x) = \sum_{s=0}^4 \frac{(-1)^s}{d^3} \binom{4}{s} \left[x - \lambda_{i-s}^{(x)} \right]^3 \mu \left[x - \lambda_{i-s}^{(x)} \right] \quad (18)$$

$$\mu(x) = \begin{cases} x, & x > 0 \\ 0, & \text{elsewhere} \end{cases} \quad (19)$$

where $\lambda_i^{(x)}$ is referred to as a knot location and d is the distance between knots. A spline $B_i(x)$ only takes nonzero values on the interval $[\lambda_{i-4}^{(x)}, \lambda_i^{(x)}]$.

The B spline coefficients θ_{ijk} for least squares approximation of a set of data are found by minimizing as follows:

$$\min_{\theta} E(\theta) + \omega R(\theta) \quad (20)$$

where E reflects the closeness of fit to the data and R reflects the roughness of the approximate function. Choosing the parameter ω determines the tradeoff between the two.

Suppose we have N measurements P_n at locations \mathbf{x}_n . Then E and R are as follows:

$$E(\theta) = \frac{1}{N} \sum_{n=1}^N [P_n - p(\mathbf{x}_n)]^2 \quad (21)$$

$$R(\theta) = \frac{1}{V} \int_D \sum_{i=1}^3 \sum_{j=1}^3 \left[\frac{\partial^2 p(\mathbf{x})}{\partial x_i \partial x_j} \right]^2 dx \quad (22)$$

where D contains the ROI and V is the volume of D .

In matrix form, the solution of (20) for θ is given by

$$\theta = (A^T A + \omega J)^{-1} A^T Z \quad (23)$$

where an element of the N by Q matrix A is

$$A_{n,ijk} = \frac{\partial p(\mathbf{x}_n)}{\partial \theta_{ijk}} = B_i(x_n) B_j(y_n) B_k(z_n) \quad (24)$$

$$Q = M_x M_y M_z \quad (25)$$

and J is a Q by Q matrix given by

$$J = \sum_{\substack{\alpha_x, \alpha_y, \alpha_z \geq 0 \\ \alpha_x + \alpha_y + \alpha_z = 2}} \frac{2}{\alpha_x! \alpha_y! \alpha_z!} J_x^{(\alpha_x)} \otimes J_y^{(\alpha_y)} \otimes J_z^{(\alpha_z)} \quad (26)$$

where $J_x^{(l)}_{i,j}$ is an element of an M_x by M_x matrix with

$$J_x^{(l)}_{i,j} = \frac{1}{V} \int_D B_i^{(l)}(x) B_j^{(l)}(x) dx. \quad (27)$$

The symbol \otimes denotes Kronecker product. θ and Z are column vectors with elements in their natural order. D is most conveniently taken as

$$D = \left[\lambda_0^{(x)}, \lambda_{M_x-3}^{(x)} \right] \times \left[\lambda_0^{(y)}, \lambda_{M_y-3}^{(y)} \right] \times \left[\lambda_0^{(z)}, \lambda_{M_z-3}^{(z)} \right] \quad (28)$$

although this may be larger than the ROI.

Equation (23) may not be solvable directly if the smoothing parameter is very small or zero since the matrix $A^T A$ is often ill conditioned. Hayes and Halliday [31] have proposed a technique using Householder transformations to stabilize the method. However, in practice ω can be chosen large enough to avoid numerical difficulties.

ACKNOWLEDGMENT

The authors would like to thank C. Holmes, N. Kabani, V. Kollokian, and S. Dumoulin for their assistance in preparing the data for this work. The MNI N3 software package is available freely by anonymous ftp at ftp://ftp.bic.mni.mcgill.ca/pub/mni_n3/.

REFERENCES

- [1] H. S. Choi, D. R. Haynor, and Y. Kim, "Partial volume tissue classification of multichannel magnetic resonance images—A mixel model," *IEEE Trans. Med. Imag.*, vol. 10, pp. 395–407, Sept. 1991.
- [2] M. Özkan, B. M. Dawant, and R. J. Maciunas, "Neural-network-based segmentation of multimodal medical images: A comparative and prospective study," *IEEE Trans. Med. Imag.*, vol. 12, pp. 534–544, Sept. 1993.
- [3] A. Simmons, P. S. Tofts, G. J. Barker, and S. R. Arridge, "Sources of intensity nonuniformity in spin echo images," *Magn. Reson. Med.*, vol. 32, pp. 121–128, 1994.
- [4] D. A. G. Wicks, G. J. Barker, and P. S. Tofts, "Correction of intensity nonuniformity in MR images of any orientation," *Magn. Reson. Imag.*, vol. 11, no. 2, pp. 183–196, 1993.
- [5] A. P. Zijdenbos, B. M. Dawant, and R. A. Margolin, "Intensity correction and its effect on measurement variability in the computer-aided analysis of MRI," in *Proc. 9th Int. Symp., Exhibit. Computer Assisted Radiology (CAR)*, Berlin, Germany, June 1995, pp. 216–221.

- [6] E. R. McVeigh, M. J. Bronskill, and R. M. Henkelman, "Phase and sensitivity of receiver coils in magnetic resonance imaging," *Med. Phys.*, vol. 13, pp. 806–814, Nov./Dec. 1986.
- [7] S. E. Moyher, D. B. Vigneron, and S. J. Nelson, "Surface coil MR imaging of the human brain with an analytic reception profile correction," *J. Magn. Reson. Imag.*, vol. 5, pp. 139–144, Mar./Apr. 1995.
- [8] J. G. Sled, A. C. Evans, and G. B. Pike, "Standing-wave and RF penetration artifacts caused by elliptic geometry: an electrodynamic analysis of MRI," *Int. Soc. Magn. Reson. in Med.*, p. 1512, 1997.
- [9] R. Stollberger and P. Wach, "Imaging of the active B1 field *in vivo*," *Magn. Reson. Med.*, vol. 35, pp. 246–251, 1996.
- [10] P. A. Narayana, W. W. Brey, M. V. Kulkarni, and C. L. Sievenpiper, "Compensation for surface coil sensitivity variation in magnetic resonance imaging," *Magn. Reson. Imag.*, vol. 6, no. 3, pp. 271–274, 1988.
- [11] F. Bloch, W. Hanson, and M. Packard, "Nuclear induction," *Physical Rev.*, vol. 69, p. 127, Feb. 1946.
- [12] K. R. Thulborn, F. E. Boada, J. D. Christensen, F. R. Haug-Hellinger, T. G. Reese, and J. M. Kosewski, "B1 correction maps and apparent water density maps as tools for quantitative functional MRI," *Soc. Magn. Reson. Med.*, vol. 1, p. 347, 1993.
- [13] L. Axel, J. Costantini, and J. Listerud, "Intensity correction in surface-coil MR imaging," *Amer. J. Roentgenol.*, vol. 148, pp. 418–420, Feb. 1987.
- [14] B. M. Dawant, A. P. Zijdenbos, and R. A. Margolin, "Correction of intensity variations in MR images for computer-aided tissue classification," *IEEE Trans. Med. Imag.*, vol. 12, pp. 770–781, Dec. 1993.
- [15] K. O. Lim and A. Pfefferbaum, "Segmentation of MR brain images into cerebrospinal fluid spaces, white and gray matter," *J. Comput. Assist. Tomogr.*, vol. 13, pp. 588–593, July/Aug. 1989.
- [16] C. R. Meyer, P. H. Bland, and J. Pipe, "Retrospective correction of intensity inhomogeneities in MRI," *IEEE Trans. Med. Imag.*, vol. 14, pp. 36–41, Mar. 1995.
- [17] P. A. Narayana and A. Borthakur, "Effect of radio frequency inhomogeneity correction on the reproducibility of intra-cranial volumes using MR image data," *Magn. Reson. Med.*, vol. 33, pp. 396–400, Mar. 1995.
- [18] W. M. Wells III, W. E. L. Grimson, R. Kikinis, and F. A. Jolesz, "Adaptive segmentation of MRI data," *IEEE Trans. Med. Imag.*, vol. 15, no. 4, pp. 429–442, 1996.
- [19] L. Devroye, *Non-Uniform Random Variate Generation*. Englewood Cliffs, NJ: Prentice-Hall, 1986.
- [20] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Trans. Biomed. Eng.*, vol. BME-9, pp. 63–66, 1979.
- [21] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*. New York: Wiley, 1973.
- [22] L. L. Schumaker and F. I. Utreras, "On generalized cross validation for tensor smoothing splines," *SIAM J. Scientific, Statistical Computing*, vol. 11, pp. 713–731, July 1990.
- [23] A. Thompson, J. Brown, and T. Titterton, "A study of methods of choosing the smoothing parameters in image restoration by regularization," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 13, no. 4, pp. 326–339, 1991.
- [24] C. J. Holmes, R. Hoge, L. Collins, and A. C. Evans, "Enhancement of T1 MR images using registration for signal averaging," in *Proc. 2nd Int. Conf. Functional Mapping of the Human Brain*, 1996, p. S28.
- [25] C. J. Holmes, R. Hoge, R. P. Woods, A. C. Evans, and A. W. Toga, "Enhancement of T2 and proton density MR images using registration for signal averaging," in *Proc. 2nd Int. Conf. Functional Mapping of the Human Brain*, 1996, p. S28.
- [26] J. Bezdek, *Pattern Recognition with Fuzzy Objective Function Algorithms*. New York: Plenum, 1981.
- [27] M. Kamber, R. Shinghal, D. L. Collins, G. S. Francis, and A. C. Evans, "Model-based 3-D segmentation of multiple sclerosis lesions in magnetic resonance brain images," *IEEE Trans. Med. Imag.*, vol. 14, pp. 442–453, Sept. 1995.
- [28] I. Harvey, P. S. Tofts, J. K. Morris, D. A. Wicks, and M. A. Ron, "Sources of T_1 variance in normal human white matter," *Magn. Reson. Imag.*, vol. 9, no. 1, pp. 53–59, 1991.
- [29] R. Guillemaud and M. Brady, "Enhancement of MR images," in *Proc. 4th Int. Conf. Visualization in Biomedical Computing*, 1996, pp. 107–116.
- [30] G. O. Glover, C. E. Hayes, N. J. Pelc, W. A. Edelstein, O. M. Mueller, H. R. Hart, C. J. Hardy, M. O'Donnell, and W. D. Barber, "Comparison of linear and circular polarization for magnetic resonance imaging," *J. Magn. Reson.*, vol. 64, pp. 255–270, 1985.
- [31] J. G. Hayes and J. Halliday, "The least-squares fitting of cubic spline surfaces to general data sets," *Inst. Math. Applicat.*, vol. 14, pp. 89–103, 1974.