Rapid and Effective Correction of RF Inhomogeneity for High Field Magnetic Resonance Imaging

Mark S. Cohen*, Richard M. DuBois, and Michael M. Zeineh

UCLA Brain Mapping Division, UCLA School of Medicine, Los Angeles, California

◆ ═══════════════════════════════════ ◆

Abstract: The well-known variability in the distribution of high frequency electromagnetic fields in the human body causes problems in the analysis of structural information in high field magnetic resonance images. We describe a method of compensating for the purely intensity-based effects. In our simple and rapid correction algorithm, we first use statistical means to determine the background image noise level and the edges of the image features. We next populate all "noise" pixels with the mean signal intensity of the image features. These data are then smoothed by convolution with a gaussian filter using Fourier methods. Finally, the original data that are above the noise level are normalized to the smoothed images, thereby eliminating the lowest spatial frequencies in the final, corrected data. Processing of a 124 slice, 256×256 volume dataset requires under 70 sec on a laptop personal computer. Overall, the method is less prone to artifacts from edges or from sensitivity to absolute head position than are other correction techniques. Following intensity correction, the images demonstrated obvious qualitative improvement and, when subjected to automated segmentation tools, the accuracy of segmentation improved, in one example, from 35.3% to 84.7% correct, as compared to a manually-constructed gold standard. *Hum. Brain Mapping 10:204–211, 2000.*   © 2000 Wiley-Liss, Inc.

Key words: human brain/anatomy and histology; magnetic resonance imaging; theoretical models; artifacts; segmentation; radio frequency

◆ ═══════════════════════════════════ ◆

INTRODUCTION

Although highly uniform radio frequency (RF) fields can be produced [Hayes et al., 1985; Li et al., 1997] in nonconducting materials (e.g., silicon oil), the dielectric properties of biological materials result in RF field distributions that are frequency dependent [Hetherington et al., 1994; Simmons, 1994; Vaughan et al., 1998]. In magnetic resonance (MR) imaging, this results in image intensity nonuniformities that vary with pulse sequence, MR field strength and body tissues. Such intensity variation becomes significant when its magnitude is large compared to the image contrast to noise ratio, as it can easily color the interpretation of viewed images, and decreases the useable image contrast. In most cases, the good RF homogeneity and the moderate signal-to-noise ratio (SNR) of imaging at field strengths of 0.5 Tesla or less, combine to make this a minimal problem. At higher fields, however, the images take on pronounced intensity distortion. Body coil images at 1.5 Tesla, for example,

often seem to have a position dependent intensity (shading). At field strengths of 3 Tesla or more, the image intensity uniformity, even in head imaging, often is problematic.

Increasingly, high-resolution brain images are submitted to automated segmentation programs to define important structural boundaries. An example is dividing the images, first for tissue classification into gray matter, white matter, cerebrospinal fluid, and extracranial tissue compartments, and subsequently into finer parcellations, such as subcortical structures and specific disease entities [Jernigan et al., 1993; Kennedy, 1989]. The preferred acquisition methods, in many cases, are 3D (volume) scans, using low flip angle gradient echo FLASH [Frahm et al., 1986] and its derivatives. These can produce extremely high resolution and good SNR ratios in a reasonably short time (as compared, for example, to slice-selective spin echo or fast spin echo). Such images can often be acquired with near isotropic spatial resolution. In our experience, however, even with inversion prepared spoiled gradient echo studies, the relatively weak contrast between brain compartments may be overwhelmed by the in-plane (and through-plane) brightness variation present in the images. Consequently, although the SNR and resolution of the images are good, their usefulness for morphometric analysis is compromised.

We have developed a correction method that avoids the problems of intensity artifacts and other contaminants and that executes very rapidly without the need for expensive computer hardware or computation time. Here we present the method, the basic results on image quality, and examples of quantitative improvements in MR segmentation.

MATERIALS AND METHODS

Images

The algorithm is completely blind to the scan parameters used for data collection. In our experience, however, the most severe intensity variations are as-

Abbreviations	
tr	repetition time
te	echo time
T1	longitudinal relaxation time
T2*	transverse relaxation time
NEX	number of excitations
SPGR	Spoiled gradient recalled acquisition
FFT	Fast Fourier transform

sociated with low flip angle (FLASH-type) [Frahm et al., 1986] imaging for the reasons discussed previously. We compared the image intensity uniformity in 2D spin echo and gradient echo images on our scanner, a 3.0 Tesla General Electric Signa® instrument with gradient systems and image processing modified by Advanced NMR systems. The gradient echo scans were performed using the SPGR sequence using a repetition time (tr)/echo time (te) of 30/6 msec, 16 number of excitations (NEX), 20° flip angle, a field of view of 24 cm and a 3 mm slice thickness. The spin echo images were acquired using a tr of 500 msec, a te of 11 msec, and 2 NEX, with the other parameters the same. For the 3D volume scans of our human subjects, we used the spoiled gradient recalled acquisition (SPGR) sequence [Darrasse et al., 1988], with a tr of 14.5 ms, te of 3.3 msec, and a flip angle of 20°. The images were acquired with a field of view of 24 cm, a 256×256 in-plane image matrix and 60 or 124 scan planes of 1.2 mm thickness and 1 or 2 NEX. For our phantom 3D images, we used the same sequence with 60 slice locations and 3 mm thickness and 1 NEX. All subjects signed standard informed consents approved by the UCLA institutional review boards. We have evaluated the program on a wide variety of input datasets, but have selected the T1-weighted volume series as a worst case example, and one that is highly relevant to the problems of tissue segmentation which motivated the development of this approach.

Algorithm

Our method involves five steps (shown in Fig. 1). (1) Determine (automatically) the threshold intensity below which pixels can be classified as "noise" (Fig. 1A). (2) Determine the average signal intensity in the nonnoise locations. (3) Fill the "noise" locations with the average image intensity of the nonnoise locations (Fig. 1B). (4) Smooth the image data (Fig. 1C). (5) Normalize the signal intensity of the raw image by the smooth image and correct the intensity so that the average pixel intensity (in the volume) is retained after correction. Specifically, the image intensity at location i(SI[i]) becomes:

$$SI[i] = SI_{average} \frac{SI_{raw}[i]}{SI_{smooth}[i]}, \qquad (1)$$

where $SI_{average}$ is the average intensity over the volume, $SI_{raw}[i]$ is the original pixel intensity, and $SI_{smooth}[i]$ is the pixel intensity at the ith location following smoothing (Fig. 1D).
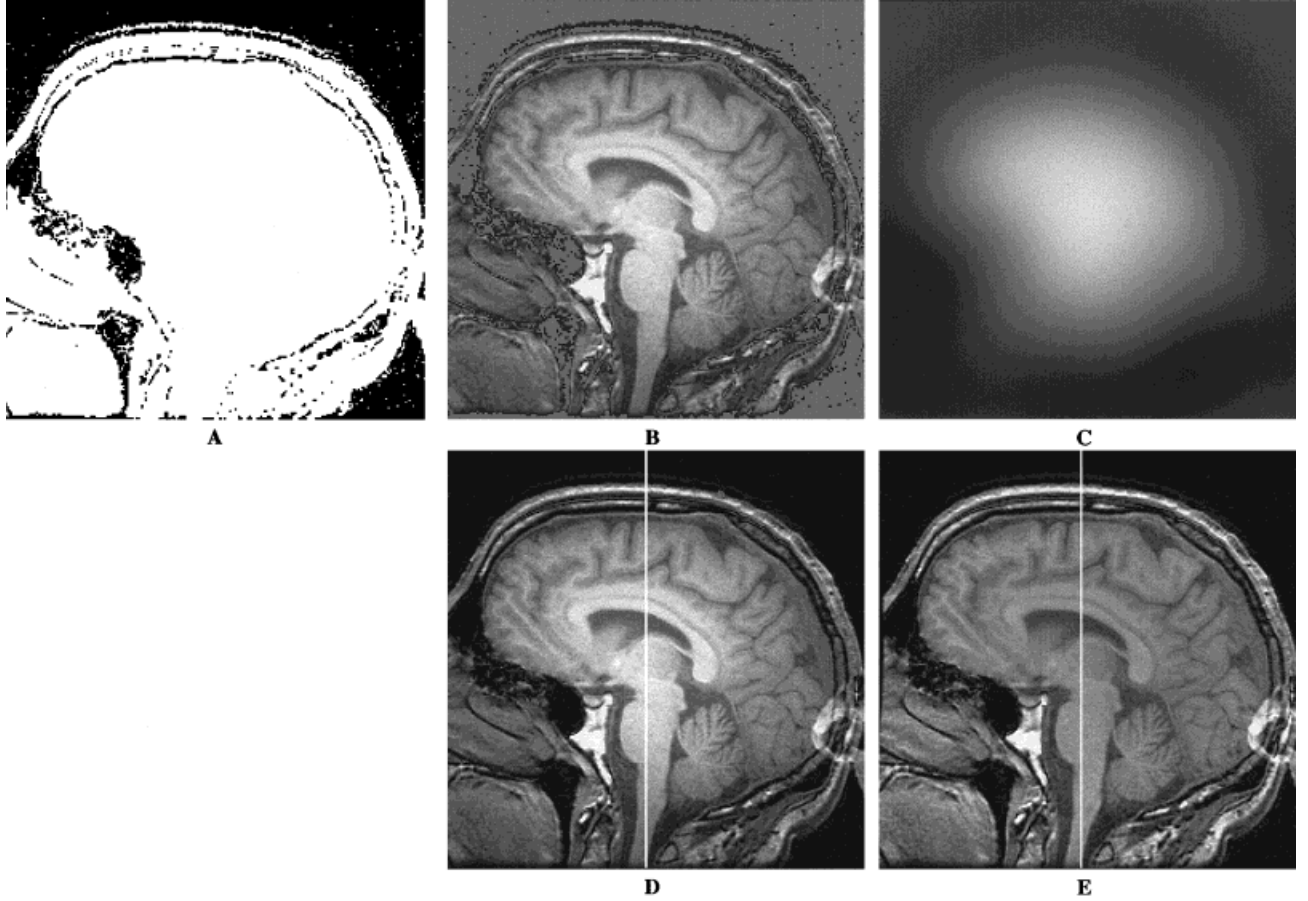
Figure 1.

Processes involved in intensity equalization (1 of 124 slices shown). (**A**) Automatically determined segmentation of the input image into data (white) and background (black). (**B**) The background pixels are first filled with the average signal found in the nonnoise pixels. (**C**) Three-dimensional smoothing is performed by convolution with a gaussian kernel. (**D**) The original input image shows marked brightness irregularity. (**E**) Following correction, the image features are much more easily resolved (N.B., images D and E are presented with identical window settings).

Threshold determination is performed by first creating a histogram of all pixel intensities in the entire volume and selecting as a threshold the first histogram minimum in signal intensities (Fig. 1B). The histogram minimum is determined by first smoothing the intensity histogram with a Hanning kernel. We next determine the first histogram maximum at the lowest signal intensities, and we then calculate the histogram minimum in the lowest 15% of the signal intensities. This step is crucial to the elimination of edge blooming; therefore, the user is given a variety of options for overriding the automatically determined levels.

Image smoothing is performed by convolving the image volume with a gaussian kernel having a default half-width of 3/8 the image dimensions. For example, given a typical $256 \times 256 \times 124$-volume dataset, we use a smoothing kernel of $96 \times 96 \times 46$ pixels. To keep computing time to a practical level, the convolutions are performed in k-space:

$$SI_{smooth} = SI_{raw}(x, y, z) \otimes G(x, y, z)$$

$$= \mathcal{F}^{-1}[\mathcal{F}(SI_{raw}(x, y, z))\mathcal{F}(G(x, y, z))], \quad (2)$$

where $\mathcal{F}$ is the Fourier transform operation, $\mathcal{F}^{-1}$ is its inverse, $\otimes$ is the convolution operator, $SI_{raw}(x,y,z)$ is the image intensity of all pixels in the 3D volume and $G(x,y,z)$ is a three-dimensional gaussian function with the dimensions discussed previously.

Computers

We have implemented the RF correction algorithm on a Digital Equipment Corporation Alphastation
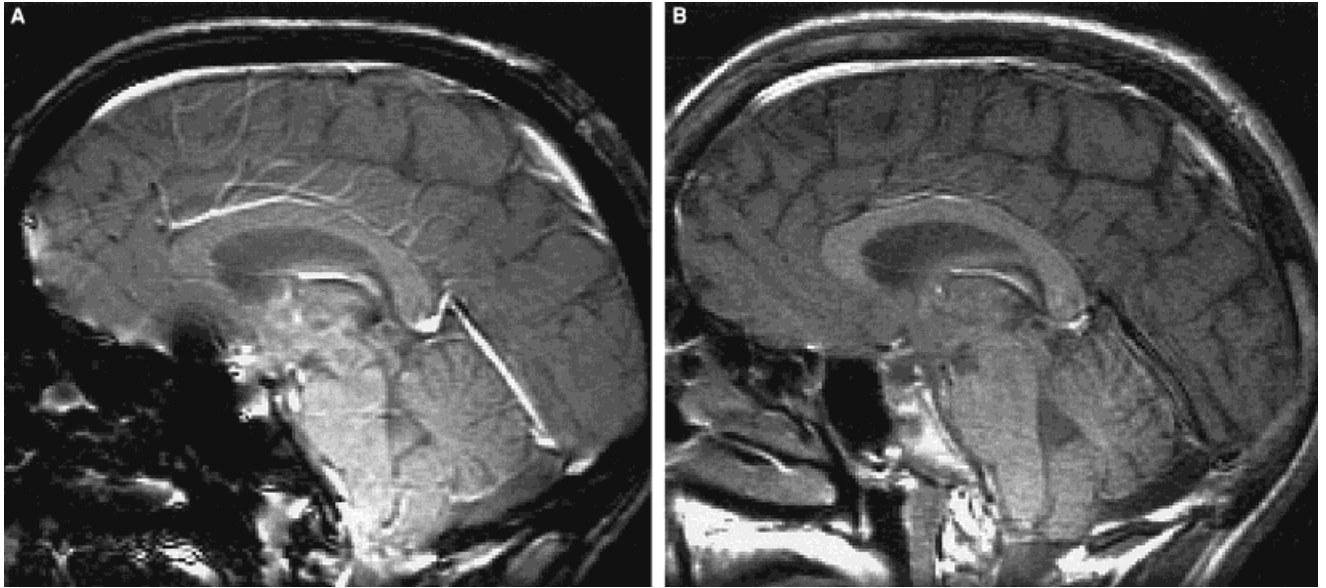
Figure 2.
Comparison of SPGR gradient echo (**A**) and spin echo (**B**) MRI of the head at 3 Tesla. The scan on the left shows much more pronounced brightness nonuniformity (as a hyperintensity of the brainstem) resulting from the sensitivity of small flip angle imaging to variations in flip angle.

500/266 workstation with 128 Mbytes of RAM (Compaq Corp., Houston, TX, USA) and on an Apple Macintosh G3 333 MHz Powerbook (laptop) computer having 192 Mbytes of RAM (Apple Computer, Inc., Cupertino, CA, USA) the results of which are reported here. We have successfully compiled and run the program, without changes, on a Pentium III-based computer under Linux and on a Silicon Graphics Reality Monster under IRIX 6.5. The source code for all machines is identical, but the Alphastation requires a byte- swapping step to convert our MR images to its internal byte order; the byte order is detected automatically. Because the 3D Fast Fourier transform (FFT) is performed in memory the program requires a 64 Mbyte data buffer for the 124-slice 3D volumes discussed in the next section, and an additional 16 Mbyte working buffer. We typically allocate a 96 Mbyte partition for the program.

Validation

One of the main targets of this technology is the automated, or semi-automated, tissue segmentation of MR images. To test our intensity correction tool we compared the efficacy of a commonly used and powerful segmentation tool [Teo et al., 1997] using manual segmentation of the same scan planes as a "gold standard." We specifically segmented the white matter in selected sagittal sections near the midline of the brain and within the occipital lobe. We calculated the difference images between the automated segmentation and the manually guided gold standard and expressed the quality of the correction as one minus the fraction of pixels that differed between manual and automated segmentation, expressed as a percentage. A value of 100% would therefore indicate that all pixels were classified identically using manual and automated method; a value of zero would indicate that no pixels are held in common.

RESULTS

To understand better the reasons for problems in brightness homogeneity in high field imaging, we compared the results obtained by scanning with small flip angle gradient echo, and conventional (90°–180°) spin echo imaging, as shown in Figure 2. The gradient echo scans show a pronounced hot spot in the middle of the images, where the RF intensity is presumably weakest. This is the predicted result of the relationship between flip angle, T1, tr, and signal intensity (see discussion).

Using a 124-slice 3D volume of 256 × 256 pixel images, the correction process requires 90 sec on our Alphastation. Running the same dataset on the Apple laptop took 68 sec. Of this, some 42 sec were spent in
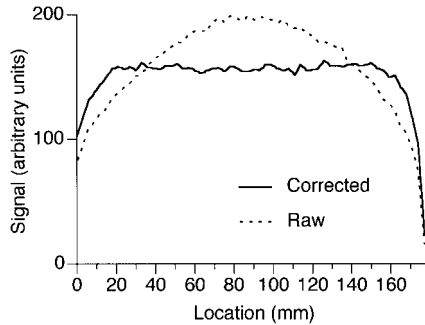
Figure 3.

Signal intensity profile along the Z axis in a uniform phantom prior to and following image correction.

the still poorly optimized 3D FFT functions; we anticipate much better performance as G4 processors become available on the Macintosh, as they include highly optimized vector calls to Fourier transform routines. The output image files were, of course, identical for the two computers.

The results of the sequential processing steps used in our method are illustrated in Figure 1, as discussed in the Methods section. The subjective results (using the automatically calculated default parameters) of the correction are extremely encouraging. For comparison, Figure 1D and E compare the image prior to, and following, intensity correction. The line appearing in those figures indicates the profile that appears in Figure 3.

We assessed the correction method quantitatively on the phantom images as shown in Figure 3. This compares an intensity profile through the Z axis in our 60 slice volume prior to and following image correction. The method does not completely correct the edges, although they are improved. The local intensity variations, which presumably result from Gibbs ringing in the 3D sequence, are well preserved following correction. As a figure of merit, we calculated the r.m.s. deviation in the signal, prior to and following correction, from the mean signal intensity in the region of the image that contained the phantom. The deviation was 20.1% and 7.9% for the uncorrected and corrected images, respectively.

A quantitative comparison in human imaging is shown in Figure 5, which shows the signal intensity profile along the vertical line that appears in Figure 1D and E. In this case, the uncorrected signal is nearly 70% higher in the center of the images than it is following correction. Note also, that there is little evidence of artifactual "blooming" of the signal near the image edges.

Figures 4 and 6 compare the segmentation results prior to and following correction in a midline sagittal image and in a slice within the occipital lobe. Using manual segmentation (of the uncorrected image) as the gold standard, the automated results for the midline sagittal data were 51.2% accurate (5901 pixels were classified manually as white matter, 2877 pixels had a different classification with the automated segmentation) without intensity correction and 70% with intensity correction. Figure 4C shows the segmentation results for the sagittal data. A similar segmenta-
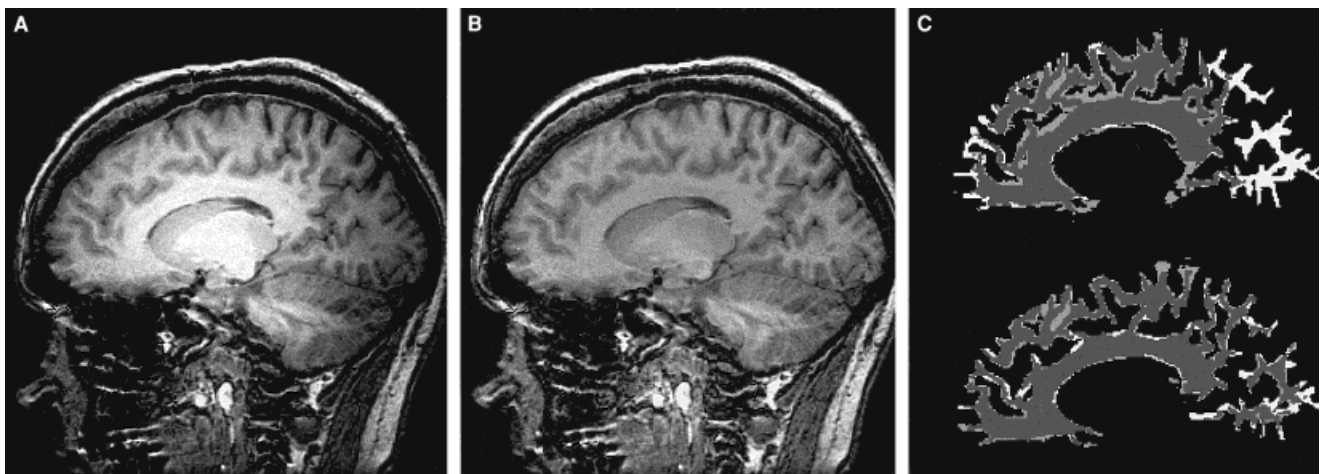
Figure 4.

(**A**) Uncorrected midline sagittal image. (**B**) The same image following correction. (**C**) Segmentation results—top: prior to correction, bottom: following correction. Dark gray pixels are those that were identified as white matter both manually and automatically. White pixels were "missed" by automatic segmentation; light gray pixels were included by automatic by automatic segmentation but not manually.
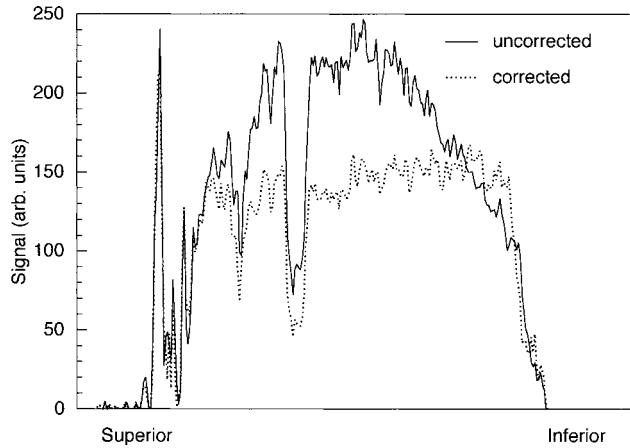
Figure 5.
Quantitative comparison of the signal intensities in a 3D gradient echo scan prior to (solid line) and following (dashed line) the application of our new RF intensity correction. Superior and inferior labels refer to the locations on Figure 1C and D. The prominent dark feature in the intensity profile is the medial aspect of the lateral ventricle.

tion analysis for a portion of a sagittal slice in the occipital lobe yielded 35.3% correct classification without intensity correction and 84.7% with it, as shown in Figure 5.

DISCUSSION

As more users move to high field MR instruments for both clinical and research applications, the intrinsic problems of image intensity nonuniformity become increasingly troublesome. For intensity-based image segmentation these lead to large quantitative errors and have forced users to resort to highly supervised segmentation. Even for qualitative viewing, the intensity variations result in a loss of useful contrast, as the images must be viewed with relatively large window ranges (low contrast) to represent the dynamic range of the pixel intensities. This leads to the impression that the high field images are of relatively low contrast.

A wide variety of intensity correction algorithms exist already, as reviewed briefly in the following section, and are effective to varying degrees. The simplest approaches normalize the input images by smoothed versions of themselves, so that the low spatial frequencies are removed. This approach, in our hands, has been reasonably effective, except for the fact that it causes prominent artifacts at the edges of structures. This takes place because the large intensity discontinuities at the edges appear, in the smoothed images, as slowly varying signal changes. The result, following normalization, is a large artificial increase in MR signal intensity in image contents at the edges of objects, and a less disturbing decrease in the noise intensity outside of the tissue volume.

More sophisticated approaches model the RF field inside the receiver coil, preferably under (simulated) loaded conditions [Wicks et al., 1993]. This strategy is appealing because it is largely independent of the image contents and their intrinsic signal differences. Although this may seem an advantage, it is, in fact, a liability when the scan repetition times (trs) are less than the tissue T1. The nonuniform RF transmission (excitation) profile may lead to signal intensity variations in homogeneous tissue volumes, as the signal
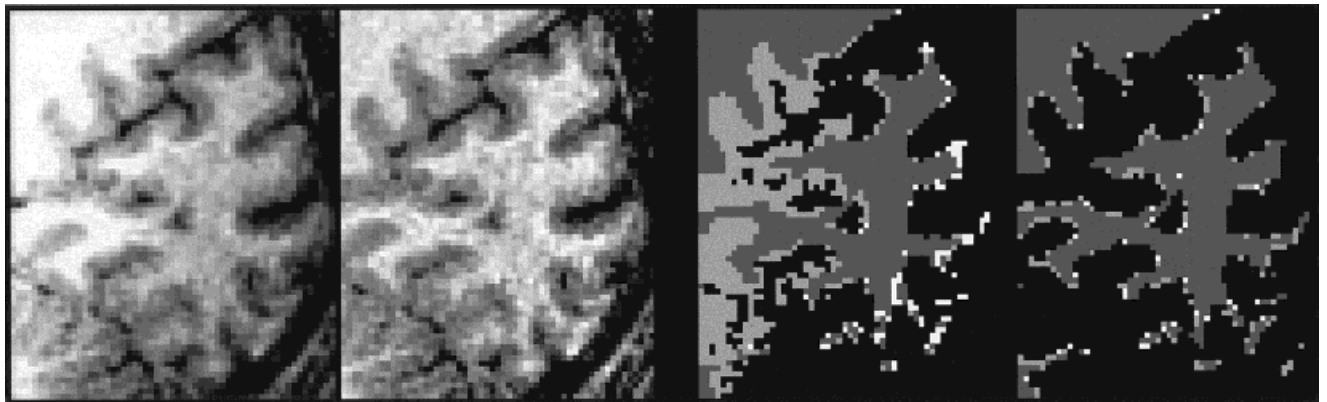
Figure 6.
(Left) Uncorrected sagittal image in the occipital lobe. (Middle left) The same image following correction. (Middle right) Segmentation results prior to correction. (Right) Segmentation results following correction. Color values as in Figure 4.
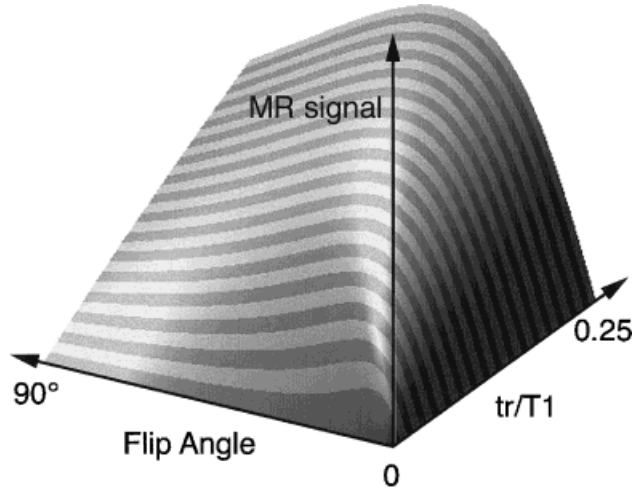
Figure 7.
The FLASH equation (3) predicts that when tr is short compared to T1, the MR signal derived using small flip angles will be greater than that using large flip angles.

intensity (SI) depends on T1, flip angle, and tr, according to the well-known FLASH equation:

$$SI = kM_0 \sin(\alpha)\exp(-te/T2^*)$$

$$\times \frac{1 - \exp(-tr/T1)}{1 - \cos(\alpha)\exp(-tr/T1)}. \quad (3)$$

Due to the inhomogeneity of the RF fields within the body, the effective flip angle will vary with location. It can be shown readily [Ernst and Anderson, 1966] that the MR signal is at its maximum when:

$$\alpha = \cos^{-1}\exp(-tr/T1). \quad (4)$$

This maximum may occur at very small flip angles when longitudinal relaxation time (T1) is long and tr is short. For example, because the flip angle is typically lower toward the center of the sample, this frequently results in increased brightness toward the center of the images. This RF-dependent intensity profile is particularly egregious when the trs are very short, as is common in 3D volume imaging; an example of this problem occurs as Figure 2 in this paper. Figure 7 shows the predicted relationship between tr, T1, flip angle, and MR signal (from Eq. 3) and indicates clearly the large relative signal intensity increases to be expected from small flip angles at very short trs. In fact, due to dielectric resonances within the rather inhomogeneous conducting volume of the head, which may account for the majority of the observed RF nonuniformity [Sled and Pike, 1998], it is difficult, or impos-

sible to determine a priori RF field distributions to use in intensity correction models.

Another class of methods utilizes a priori information about the sample contents to inform the correction algorithm [DeCarli et al., 1996; Koivula et al., 1997]. An extremely interesting and powerful feature-based approach has been described by Sled and co-workers [Sled et al., 1998]. In their method, the image intensity is adjusted iteratively at varying levels of anatomic detail, to remove only global variations in image intensity that are independent of key anatomical features. The results are quite impressive. However, iterative solutions such as this are computationally inefficient; even on very powerful computers they are extremely time intensive and, at least in current implementations, place significant requirements on the computers themselves. Further, they may not generalize easily to imaging of more varied samples. A related approach corrects the images following a knowledge-based segmentation of the images. Guillemand and Brady [1997], describe a method, based on prior work of Wells [Wells 3d et al., 1996] in which the images are segmented into tissue subtypes, and the bias field is then amended on the segmentation results. The method is iterative, first performing a crude segmentation, and then using the assumption that the intensity should be uniform in the classified tissue types to calculate a bias field. Although such an approach can lead to excellent results, the complexity is relatively high and, as noted by the authors, the quality of the correction can be expected to depend on the accuracy of the segmentation.

The qualitative and quantitative results shown here underscore the importance of performing some sort of intensity correction, especially on high field datasets. It is particularly noteworthy, however, that the intensity variations have a complex dependence on position and T1, as emphasized in Figure 2. One consequence is that following virtually any intensity-based correction the quantitative values of T1 will be lost.

Regarding our segmentation results, these data point to, but do not in themselves confirm, a substantial advantage from intensity correction. Unfortunately, there is no ground truth for the absolute differentiation of, for example, white from gray matter. Further, as performed here, this could be seen as much as a study of the segmentation algorithm as the intensity correction method. Nevertheless, a cursory examination of Figures 4 and 5 reveals that the segmentation performance is very much improved by prior intensity correction. The algorithm we used by Teo et al. [1997] is in some ways particularly revealing; this tool has been optimized for brain matter segmenta-

tion, and *should* be relatively resistant to slow variations in image intensity. In practice, when the magnitude of these variations is large, as it is in our 3 Tesla unit, even this segmentation scheme is ultimately limited.

CONCLUSIONS

Although there are a wide variety of intensity correction methods already available, we feel that a rapid, unsupervised, and general-purpose tool set has until now been needed. Our new algorithm is in many ways unsophisticated, even simplistic. These considerations notwithstanding, we believe that we have developed a highly useful approach to image intensity correction that offers the advantages of rapid operation, ease of setup and that is essentially independent of the MR instrument, the scan parameters, or the object being imaged. Our analysis of the signal intensity behavior of FLASH-type scanning underscores this point: even accurate quantitative determination of the actual RF field strength within the loaded coil will still leave the T1-dependent intensity variations uncorrected. Although it would probably be possible separately to measure the tissue T1 and fit the pixel intensities to a MR sequence dependent correction, it would not likely be practical to do so.

Quantitative comparisons with other intensity corrections are now in progress; the gains for our intended application in image segmentation are substantial, however, both in improved accuracy and in reducing the time required for segmentation. The full source code for both unix and Macintosh computers is available for anonymous download at: http://www.brainmapping.org/sharedCode.

ACKNOWLEDGMENTS

REFERENCES

Darrasse IL, Mao et al. (1988): Spoiling techniques in very fast TR imaging. Society for Magnetic Resonance in Medicine seventh annual meeting, San Francisco, CA.

DeCarli C, Murphy DG, et al. (1996): Local histogram correction of MRI spatially dependent image pixel intensity nonuniformity. J Magn Reson Imaging 6:519–528.

Ernst R, Anderson W (1966): Application of Fourier transform spectroscopy to magnetic resonance. Rev Sci Instr 37:93–102.

Frahm JA, Haase, et al. (1986): Rapid NMR imaging of dynamic processes using the FLASH technique. Magn Reson Med 3:321–327.

Guillemaud R, Brady M (1997): Estimating the bias field of MR images. IEEE Trans Med Imaging 16:238–251.

Hayes C, Edelstein W, et al. (1985): An efficient, highly homogeneous radiofrequency coil for whole-body NMR imaging at 1.5 T. J Magn Reson 63:622–628.

Jernigan TL, Archibald S, et al. (1993): Magnetic resonance imaging morphometric analysis of cerebral volume loss in human immunodeficiency virus infection. The HNRC Group. Arch Neurol 50:250–255.

Kennedy DK (1989): Anatomic segmentation and volumetric calculation in cross-sectional magnetic resonance imaging. IEEE TMI 8:1–7.

Koivula A, Alakuijala J, et al. (1997): Image feature based automatic correction of low-frequency spatial intensity variations in MR images. Magn Reson Imaging 15:1167–1175.

Li S, Collins CM, et al. (1997): A method to create an optimum current distribution and homogeneous B1 field for elliptical birdcage coils. Magn Reson Med 37:600–608.

Simmons A. (1994): Sources of intensity nonuniformity in spin echo images at 1.5T. Magn Reson Med 32:121–128.

Sled JG, Pike GB (1998): Standing-wave and RF penetration artifacts caused by elliptic geometry: an electrodynamic analysis of MRI. IEEE Trans Med Imaging 17:653–662.

Sled JG, Zijdenbos AP, et al. (1998): A nonparametric method for automatic correction of intensity nonuniformity in MRI data. IEEE Trans Med Imaging 17:87–97.

Teo PC, Sapiro G, et al. (1997): Creating connected representations of cortical gray matter for functional MRI visualization. IEEE Trans Med Imaging 16:852–863.

Vaughan JT, Hetherington HP, et al. (1994): High frequency volume coils for clinical NMR imaging and spectroscopy. Magn Reson Med 32:206–218.

Wells 3d WM, Grimson WEL, et al. (1996): Adaptive segmentation of MRI data. IEEE Trans Med Imaging 15:429–442.

Wicks DA, Barker GJ, et al. (1993): Correction of intensity nonuniformity in MR images of any orientation. Magn Reson Imaging 11:183–196.