

Sequence-independent segmentation of magnetic resonance images

Bruce Fischl,^{a,b,*} David H. Salat,^a André J.W. van der Kouwe,^a Nikos Makris,^c
Florent Ségonne,^b Brian T. Quinn,^a and Anders M. Dale^{a,d}

^aDepartment of Radiology, MGH, Athinoula A Martinos Center, Harvard Medical School, Charlestown, MA 02129, United States

^bMIT AI Lab, Cambridge, MA 02139, United States

^cDepartment of Neurology, Center for Morphometric Analysis, MGH, Harvard Medical School, Charlestown, MA 02129, United States

^dDepartment of Neurosciences, University of California at San Diego, La Jolla, CA 92093, United States

Available online 22 September 2004

We present a set of techniques for embedding the physics of the imaging process that generates a class of magnetic resonance images (MRIs) into a segmentation or registration algorithm. This results in substantial invariance to acquisition parameters, as the effect of these parameters on the contrast properties of various brain structures is explicitly modeled in the segmentation. In addition, the integration of image acquisition with tissue classification allows the derivation of sequences that are optimal for segmentation purposes. Another benefit of these procedures is the generation of probabilistic models of the intrinsic tissue parameters that cause MR contrast (e.g., T1, proton density, T2*), allowing access to these physiologically relevant parameters that may change with disease or demographic, resulting in nonmorphometric alterations in MR images that are otherwise difficult to detect. Finally, we also present a high band width multiecho FLASH pulse sequence that results in high signal-to-noise ratio with minimal image distortion due to B0 effects. This sequence has the added benefit of allowing the explicit estimation of T2* and of reducing test–retest intensity variability.

© 2004 Elsevier Inc. All rights reserved.

Keywords: Morphometry; MRI; Alzheimer's disease

Introduction

Neurodegenerative disorders, psychiatric disorders, and healthy aging are all frequently associated with structural changes in the brain. These changes can cause alterations in the imaging properties of brain tissue, as well as changes in morphometric properties of brain structures. Morphometric changes may include variations in the volume or shape of subcortical regions, as well as alterations in the thickness, area, and folding pattern of the cortex. Surface-based analyses can provide an accurate assessment of cortical variability,

while volumetric techniques can detect changes in noncortical structures. For example, changes in ventricular or hippocampal volume are frequently associated with a variety of diseases (e.g., Killiany et al., 2000; Puri et al., 1999; Wolf et al., 2001).

Unfortunately, these changes are frequently difficult to detect due to technical reasons. For example, combining scans from different manufacturers would be an attractive option, as it would allow large-scale studies to be carried out across different scan sites, potentially resulting in significantly increased statistical power to detect small changes. However, merging data across scanners is problematic, as the contrast properties of the scans may vary due to subtle differences in the way the images are acquired (e.g., fat suppression, gradient and RF spoiling schemes, slab selectivity, etc.). An additional problem is that distortions in the imaging process can be scanner specific (Schmitt, 1985). Probably the largest factor in preventing the combination of scans across studies and/or sites is that no set of standard pulse parameters exists for acquiring structural data. The variation in the imaging parameters results in changes in the contrast properties of the resulting images that are independent of the underlying tissue but rather reflect the physics of the imaging process. These changes in contrast properties can introduce variability in morphometric measures that do not reflect biological effects but rather technological ones and thus should be minimized.

Here we present a set of techniques for embedding forward models of the imaging process into a segmentation algorithm. This has a number of substantial benefits. First, it allows the explicit estimation of the intrinsic tissue properties that are the main source of contrast in magnetic resonance images (MRIs), including T1, T2*, and proton density (P). The segmentation is then phrased in terms of these parameters, which are properties of the brain and independent of scanner and pulse sequence, rendering it insensitive to variation in the pulse parameters.¹

* Corresponding author. Nuclear Magnetic Resonance Center, Massachusetts General Hospital/Harvard Medical School, Building 149, 13th Street, Charlestown, MA 02129. Fax: +1 617 726 7422.

E-mail address: fischl@nmr.mgh.harvard.edu (B. Fischl).

Available online on ScienceDirect (www.sciencedirect.com.)

¹ Note that there is a field strength dependence for some of the parameters such as T2* and T1, but they are not dependent on the acquisition parameters.

This invariance is accomplished by including the parameters and their effect on image contrast into the segmentation, effectively predicting the intensity distribution of each brain structure under the set of parameters used in the acquisition.

A further advantage of the integration of the imaging physics into the segmentation process is that it allows the optimization of the MR pulse parameters to be phrased directly in terms of reducing the probability of misclassification. This allows one to specify the imaging time available for a study, then to compute the MR acquisition parameters that are optimal for the segmentation. Further, if a subset of structures are of interest, the optimization can be naturally constrained to those structures, yielding sequences that are, for example, optimal for segmenting the hippocampus. It is important to note in this context that this is not equivalent to maximizing the SNR of a structure, but rather to maximizing the Mahalanobis distance between each pair of tissue classes that occurs in proximity to one another. This is in contrast to the more standard technique of maximizing overall gray/white contrast, which does not account for the spatial distribution of heterogeneous tissue such as cortical gray matter, nor the fact that the accuracy of the segmentation of the gray matter is dependent on how much the intensity distribution of the gray matter overlaps with other structures in the skull that can be adjacent to gray matter, such as cerebrospinal fluid (CSF) and dura.

The explicit estimation of the underlying tissue parameters has other potential benefits as well. For example, changes in T1 and proton density separately may not be reflected in the corresponding T1-weighted images, as they may offset each other. In addition, cell death or gliosis may entail changes in the parameters, but not in morphometric properties such as volume. Variations of this nature would thus be undetectable to a standard morphometric analysis procedure but would be clearly revealed by cross-group comparison of the tissue parameters.

Finally, it is clear that multispectral information (e.g., different flip angles or different TRs) provides additional information that is not present in a single T1-weighted image. This is due to the highly correlated nature of T1 and proton density (their correlation coefficient in brain is greater than 0.5) and the fact that they influence contrast in opposite directions. For example, cortical gray matter has a longer T1 and a higher proton density than the underlying white matter. The longer T1 tends to darken the gray matter, while the higher PD brightens it, resulting of course in the reversal of contrast seen in T1-weighted versus PD-weighted images. Thus, it is clear separating the effects of different tissue properties can significantly enhance class separation.

Methods

In previous work (Fischl et al., 2002), we developed a procedure for automatically and accurately labeling each voxel in the brain as one of 40 structures (e.g., thalamus, hippocampus, amygdala, etc.), building on the work of many researchers (Geman and Geman, 1984; Kapur et al., 1998; Pham et al., 1997; Wells et al., 1996; Zhang et al., 2001). Briefly, this procedure is based on modeling the segmentation as a nonstationary anisotropic Markov Random Field (MRF), in which the probability of a label is modulated by the probability of its neighbors, with the probabilities computed separately at each

position in an atlas, for each pair of tissue classes and for each of the six cardinal directions. This results in a maximum a posteriori (MAP) estimation problem, in which the segmentation W is given by:

$$\begin{aligned} W(\mathbf{r}) &= \arg \max_c p[W(\mathbf{r}) = c | W(\mathbf{r}_i), \mathbf{I}(\mathbf{r}'), \mathbf{r}_i] \\ &= p[\mathbf{I}(\mathbf{r}') | W(\mathbf{r}) = c] p[W(\mathbf{r}) = c] \\ &\quad \times \prod_{i=1}^K p[W(\mathbf{r}_i) | W(\mathbf{r}) = c, \mathbf{r}_i] \end{aligned} \quad (1)$$

where $\mathbf{r}' = f(\mathbf{r})$ is the image coordinate corresponding to the atlas coordinate \mathbf{r} , as given by the atlas function f .

In this paper, we present a number of enhancements to this procedure. First, the likelihood term $p(\mathbf{I}|W)$ is modified to include a forward model of image formation, yielding a segmentation procedure that can model changes in acquisition parameters and hence be insensitive to them. Second, the explicit modeling of the image formation process into the segmentation allows the derivation of MR acquisition parameters that are optimal for the purposes of segmentation. Third, we derive a nonlinear atlas function f that increases the anatomical accuracy of the atlas coordinate system. Finally, the intrinsic tissue parameters that are the source of contrast in MR imaging are themselves estimated as part of the procedure, giving access to biologically relevant parameters that are independent of the acquisition details and may vary with a variety of conditions and disorders.

Intrinsic tissue parameter estimation

MRI is an amazingly versatile technology that allows one to probe various properties of the brain through the manipulation of magnetic and radio frequency (RF) fields. Standard structural imaging is typically based on some combination of three time constants that vary by tissue type (T1, T2, and T2*), as well as the density of protons (P) that are being imaged. Here, we focus on a class of acquisition protocols known as either fast low-angle shot (FLASH) or spoiled gradient recalled echo (SPGR). These steady-state saturation recovery gradient echo sequences have a number of advantages:

1. They are easily modeled with a well-known equation for image formation.
2. They can be manipulated to generate contrast differences arising from different intrinsic tissue parameters.
3. They are commonly available on the vast majority of MR scanners.

Specifically, for these sequences, the measured signal S can be modeled as a function of the intrinsic tissue parameters $\boldsymbol{\beta} = [T1, P, T2^*]^T$ by solving the steady state Bloch equation via:

$$S(\mathbf{m}, \boldsymbol{\beta}) = P \sin \alpha \left(\frac{1 - e^{-TR/T_1}}{1 - \cos \alpha e^{-TR/T_1}} \right) e^{-TE/T_2^*} \quad (2)$$

where $\mathbf{m} = [TR, TE, \alpha]^T$ are the acquisition parameters that the user is free to modify. Estimation of the tissue parameters $\boldsymbol{\beta}$ is a well-posed problem assuming at least as many FLASH images have been collected as there are parameters to be solved for. In general,

though, making the problem overdetermined by collecting additional measurements will result in less noise in the parameter estimates. Specifically, we estimate the parameters by maximizing the likelihood of the observed signals \mathbf{I} given our estimated parameters $\boldsymbol{\beta}$:

$$\boldsymbol{\beta}(\mathbf{r}) = \arg \max p[\mathbf{I}(\mathbf{r})|\boldsymbol{\beta}(\mathbf{r})] \quad (3)$$

Assuming Gaussian noise in the image intensities, we have:

$$p[\mathbf{I}(\mathbf{r})|\boldsymbol{\beta}(\mathbf{r})] \propto e^{-0.5[\mathbf{I}(\mathbf{r}) - \mathbf{S}(\boldsymbol{\beta}(\mathbf{r}))]^T C^{-1}[\mathbf{I}(\mathbf{r}) - \mathbf{S}(\boldsymbol{\beta}(\mathbf{r}))]} \quad (4)$$

Further assuming independent identically distributed unit-variance noise and taking the log of Eq. (4), results in a standard least mean squared (LMS) estimate of $\boldsymbol{\beta}$:

$$\boldsymbol{\beta}(\mathbf{r}) = \arg \min D(\boldsymbol{\beta}, \mathbf{r}, \mathbf{m}_i) = \frac{1}{2} \sum_i [I_i(\mathbf{r}) - \mathbf{S}(\mathbf{m}_i, \boldsymbol{\beta}(\mathbf{r}))]^2 \quad (5)$$

Subject motion between FLASH scans acquired with different pulse parameters is a potential source of error in the parameter estimation procedure. Standard motion-correction algorithms (Woods et al., 1992) have difficulty accurately coregistering the images in this domain as the images have variable contrast properties. While other techniques can be employed that are less sensitive to direction of contrast (Jenkinson et al., 2002; Viola and Wells, 1995), some accuracy may be sacrificed. Here, we directly solve for the rigid alignment parameters that minimize the parameter estimation error. The parameters are then reestimated using the new alignment, and this procedure is iterated until it converges. Again, denoting the observed images by \mathbf{I}_i and the forward model given by Eq. (2), relating the tissue parameters $\boldsymbol{\beta}$ and the pulse parameters \mathbf{m}_i to a predicted image by $\mathbf{S}(\mathbf{m}_i, \boldsymbol{\beta})$, we minimize the following energy functional to solve for the rigid alignment parameters \mathbf{R}_i :

$$D(\boldsymbol{\beta}, \mathbf{m}) = \frac{1}{2} \sum_i \iiint [I_i(\mathbf{R}_i \mathbf{r}) - \mathbf{S}(\mathbf{m}_i, \boldsymbol{\beta}(\mathbf{r}))]^2 d\mathbf{r} \quad (6)$$

We minimize Eqs. (5) and (6) using a multiscale global search of the six-dimensional rigid alignment parameter space, starting with angles in the interval $[-13^\circ, 13^\circ]$ and translations in $[-10, 10]$ mm after aligning centroids. The estimation of the intrinsic tissue parameters in Eq. (5) is simplified by noting that the proton density is a scale factor that can be removed by normalizing the length of $\boldsymbol{\beta}$. The T1 can then be estimated using a precomputed table that maps T1 values discretized in 1-ms intervals from 10 ms to 10 s into normalized signal intensities in $\mathcal{O}[\log(n)]$ time with a binary search. The proton density can then be computed as the ratio of the norm of the image intensities at the voxel to the norm of the signal table at the appropriate T1 for unit proton density.

Cross-sequence prediction of means and covariance matrices

The estimation of the intrinsic tissue parameters $\boldsymbol{\beta}$ solves part of the problem of attaining the desired invariance to modifications in pulse parameters. However, in order to segment a novel image with an arbitrary set of acquisition parameters, we must be able to predict the intensity distribution of each tissue class c under the new set of MR parameters m [that is, $p(I|m, c)$]. With the assumption that the noise is Gaussian distributed, this amounts to being able to predict the mean vector and covariance matrix of each class at each atlas location in the new signal space. Given the tissue parameters $\boldsymbol{\beta}$, the means are easily predicted by using

them as input to the forward model of Eq. (2), where the MR parameters m are typically contained in most current image headers:

$$\hat{\mu}_c(\mathbf{m}, \mathbf{r}) = \mathbf{S}[\mathbf{m}_{\text{predicted}}, \mathbf{b}(\mathbf{m}_{\text{training}}, \mathbf{l}_c(\mathbf{r}))] \quad (7)$$

where $\mathbf{m}_{\text{training}}$ is the MR parameters used in the construction of the atlas, and $\mathbf{m}_{\text{predicted}}$ is the MR parameters of the image to be segmented.² Global gain factors are accounted for by normalizing the mean overall proton density to a predefined constant. Bias fields caused by the sensitivity profile of the receive coil also corrupt the estimates of the proton density P . These are accounted for by automatically selecting a set of points whose posterior probability of being white matter is high, then using these as fixed points in a soap bubble interpolation algorithm to fix the white matter intensity to a desired value (Dale et al., 1999).

For predicting the covariance structure, we decompose the noise into two parts. The first is anatomical variability in the intrinsic tissue properties of the various brain structures. The second is white noise inherent in the imaging process. This results in the predicted covariance structure given by:

$$\hat{\Sigma}_c(\mathbf{m}) = J_{\text{predicted}} \left(J_{\text{training}c}^+ \Sigma_c J_{\text{training}}^+ \right) J_{\text{predicted}}^T + \lambda \text{Id} \quad (8)$$

where J is the Jacobian matrix of \mathbf{S} , J^+ denotes the pseudo-inverse of matrix J , Id is the identity matrix of the same dimensionality as predicted covariance matrix, and λ is a constant that reflects the component of the noise that is scan dependent. The scan-dependent part of the noise encapsulated in λ reflects factors such as averaging multiple acquisitions and the band width of the scan.

Sequence optimization

Much effort has been devoted in the MRI community towards finding pulse parameters that are optimal for various tasks (Baker, 1991; Constable and Henkelman, 1991; Constable et al., 1995; Edelstein et al., 1986; Epstein et al., 1994; Grief et al., 1985; Venkatesan and Haacke, 1997), including optimization for segmentation (Prince et al., 1995). An advantage of phrasing the segmentation procedure in terms of models of image acquisition is that it provides a natural framework for formulating an energy functional for sequence optimization specifically for segmentation. In particular, we can compute the probability of mistakenly labeling a voxel as class c_1 when the true class was c_2 as a function of the MR parameters \mathbf{m} :

$$\begin{aligned} p[c_1(\mathbf{r})c_2(\mathbf{r}), \mathbf{m}] &= \int p[c_i(\mathbf{r})|\mathbf{I}(\mathbf{m})]p[\mathbf{I}(\mathbf{m})c_2(\mathbf{r})]d\mathbf{I} \\ &= p[c_1(\mathbf{r})] \int p[\mathbf{I}(\mathbf{m})c_i(\mathbf{r})]p[\mathbf{I}(\mathbf{m})c_2(\mathbf{r})]d\mathbf{I} \end{aligned} \quad (9)$$

² Note that transmit inhomogeneities can also affect P via the principle of reciprocity. If the transmit and receive coils are the same, then this completely defines the bias field. Transmit inhomogeneities also introduce a space-varying gain factor on the flip angle, which is much more difficult to deal with than standard bias fields as it changes image contrast. Fortunately, at 1.5 T with a body transmit coil, this effect is minimal.

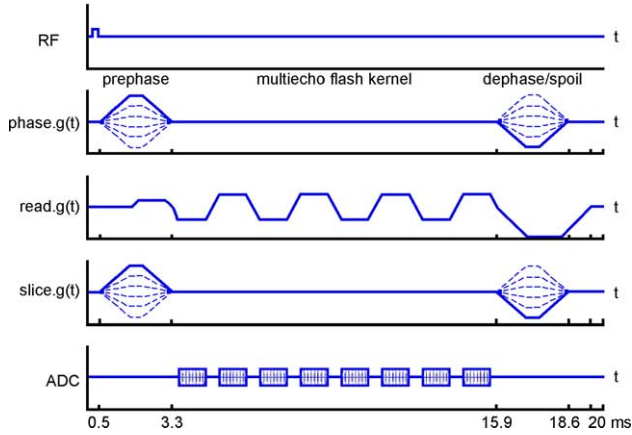


Fig. 1. Multiecho FLASH sequence timing diagram with eight echoes, non-selective RF excitation and 3-D encoding.

The total probability of misclassification at a voxel \mathbf{r} is then given by:

$$p[\mathbf{r}, \mathbf{m}] = \sum_{c_1 \neq c_2} \sum_{c_2} p[c_1(\mathbf{r})c_2(\mathbf{r}), \mathbf{m}] \quad (10)$$

By writing the synthesized images $I_i = S(\mathbf{m}_i)$ as functions of the MR parameters $\mathbf{m}_i = [TR_i, TE_i, \alpha_i]^T$ using Eq. (2), we can integrate Eq. (10) over space to yield an overall measure of the ambiguity associated with a certain combination of MR parameters:

$$A(\mathbf{m}) = \int \int \int p[\mathbf{r}, \mathbf{m}] d\mathbf{r} \quad (11)$$

We then seek the combination of MR parameters \mathbf{m} that minimize this ambiguity over all pairs of tissue classes that occur together anywhere in the brain:

$$\hat{\mathbf{m}} = \arg \min_{\mathbf{m}} A(\mathbf{m}) = \int \int \int \sum_{c_2 \neq c_1} \sum_{c_1} p[c_1 c_2, \mathbf{m}, \mathbf{r}] d\mathbf{r} \quad (12)$$

The likelihood terms $p[\mathbf{I}c_i(\mathbf{r})]$ are assumed to be normally distributed with means and covariances given by Eqs. (7) and (8).

$$p[\mathbf{I}c(\mathbf{r})]N(\hat{\boldsymbol{\mu}}_c, \hat{\boldsymbol{\Sigma}}_c) \quad (13)$$

Eq. (9) is integrated over the region of intensity that is nonzero for both classes c_1 and c_2 , typically ± 5 SD from the mean. Minimizing Eq. (12) thus amounts to reducing the amount of overlap of the distributions for tissue classes that are likely to occur at the same location.

It is important to note that the ambiguity measure captures the difficulty of segmentation in several crucial ways. First, it allows the intrinsic properties of the tissue classes to vary over space, as the ambiguity is computed separately for each atlas location. This is important, as the tissue characteristics do show considerable spatial variability, as will be shown in the Intrinsic tissue parameter distributions section. Second, only tissue classes that co-occur in a given atlas voxel contribute to the ambiguity. Thus, the difficulty of segmenting for example cortical gray matter from dura would effect the sequence optimization, but not cortical gray matter from the caudate, as these structures never occur in the same region of atlas space.

Developing multiecho FLASH acquisitions

The FLASH protocol used above is available on the vast majority of clinical scanners. As such, this sequence is limited in that it does not take advantage of recent advances in imaging hardware and acquisition techniques. In particular, it is typically a relatively low band width sequence, implying that distortions due to magnetic field inhomogeneities and susceptibility artifacts can be substantial. The low band width is of course used to increase SNR. In general, there is a trade-off between high band width, low distortion, low SNR images, and low band width, high distortion, and high SNR images. That is, SNR and distortion both go down with bandwidth.

In order to avoid this trade-off, we have developed a high band width multiecho FLASH (MEF) sequence that minimizes distortions while maximizing SNR. In a single 8-min scan, the same amount of time required for a $1.3 \times 1 \times 1$ -mm single-echo FLASH scan, this sequence provides eight high band width images at different echo times. While the individual scans can be quite noisy, the information in the ensemble is significantly greater than the low band width FLASH scans. In addition, the higher band width of the multiecho FLASH sequence, coupled with the fact that alternating echoes are collected with opposite read-out directions, results in less distortion in the images due to B0 effects (chemical shift and susceptibility distortion) (Haacke and Lenz, 1987). This is particularly important for longitudinal studies in which different shim settings can result in substantial differential distortions between scan sessions for low band width sequences. Physiologic and bulk motion during the readout also result in fewer artifacts due both to the shorter readouts of the multiecho sequence and to the averaging of the readouts with alternating directions. Finally, image reconstruction techniques can exploit the alternating readout direction to recover parts of the image previously lost due to susceptibility artifacts (Chen and Wyrwicz, 1999; Kadah and Hu, 1998; Schmithorst et al., 2001). While these reconstruction techniques are usually applied to EPI, they can also be adapted for use with multiecho FLASH images.

The MEF sequence was generated by modifying the standard Siemens gradient echo sequence to allow up to 1024 echoes following a single excitation pulse (although all results in this paper use eight echoes). The sequence can be run in 2-D or 3-D mode and maintains the standard features of the Siemens sequence, such as the option to include preparation pulses commonly used to generate magnetization transfer contrast. We increased the amount of gradient spoiling to eliminate artifacts in the images and developed a modified image reconstruction program to conveniently stream the raw k -space data to additional storage on the image reconstruction computer. A nonselective 100- μ s rectangular pulse provides the excitation. Prephasing pulses on all three gradients select the starting point in k -space for the particular set of lines as usual for a 3-D sequence. The same line is read back and forth during the train of pulses on the readout gradient and the integral of the gradients determines the position in k -space. Since the readout gradients for adjacent readouts have opposite polarities,

μ_c	$0.5*\mu_c$	$2*\mu_c$
---------	-------------	-----------

Fig. 2. One-dimensional sample image.




Fig. 3. Example of errors induced by interpolating I instead of $p(I/f)$.

the direction of the k -space traversal alternates direction, and errors in the trajectory due to B_0 errors also alternate in effect. A gradient echo is formed at the center of each readout. Dephasing gradients on the two-phase encoding gradients refocus the gradients so that the integral is zero. Spoiling is done on the readout gradient where an additional pulse after the last readout is added so that the total integral in the readout direction is four times the integral under each readout gradient (Fig. 1).

Finally, we have preliminary data suggesting that this sequence has significantly better test–retest replicability both within and across sites, a critical issue for longitudinal analysis of morphometric changes. Specifically, the variance in intensity values is significantly reduced relative to low band width FLASH and MP-RAGE sequences due presumably to the high band width of the MEF sequence, as will be shown in the Multiecho FLASH section.

Nonlinear transform

While the linear registration procedure outlined in Fischl et al. (2002) accounts for much of the anatomical variability across subjects, there is clearly much room for improvement. Nonlinear extensions to this type of intersubject registration procedure have

been extensively investigated (Ashburner et al., 1997; Collins et al., 1994; Davatzikos, 1997; Fox et al., 1984, 1985; Woods et al., 1992). Most of these techniques have error functionals with two types of terms—data-driven terms (e.g., aligning like intensities) and smoothness terms (e.g., generating morphs with minimal metric distortion). Here we extend this type of cross-subject registration to account for the fact that while smoothness of the deformation is reasonable within tissue class, it is not desirable to enforce such constraints across the border between structures.

For example, many neurodegenerative disorders result in a contraction of either cerebral white matter or cortical gray matter and a concomitant expansion of the ventricular system. Thus, a shock exists in the metric properties of the mapping at the border between the ventricles and the white matter in which adjacent locations in the brain are experiencing deformations with very different characteristics. In order to model this type of deformation, we disable the smoothness constraints on the evolving vector field across the border between structures. Specifically, we use an energy functional with five terms, similar to Fischl et al. (1999a), one to encourage smooth deformation fields, one maximizing the log likelihood of the image given the class at each atlas location, one ensuring invertibility, and one minimizing within class metric




Fig. 4. Tissue parameter estimation. Left: map of T_1 ; right: map of proton density (regions of low PD have been masked as the T_1 estimates are arbitrary there). Note the color scale at the left is only for the T_1 image, the PD values are arbitrary. Note also that regions outside the skull with extremely low PD values in which the T_1 estimates are arbitrary have been masked.

distortion, and one that discretizes the probability field directly, as will be described below:

$$G = G_I + \lambda_T G_T + \lambda_M G_M + \lambda_S G_S + \lambda_L G_L \quad (14)$$

where G_I is an intensity term, G_T is a topology constraining one, G_M is a metric preservation term, G_S is a smoothness term, and G_L is a label term, each of which is described in detail below.

Considering the atlas an array of V nodes at locations r_i , we can write the image likelihood in discrete form as:

$$G_I = \sum_{i=1}^V [\boldsymbol{\mu}_c(\mathbf{r}_i) - \mathbf{I}(\mathbf{r}_i + \mathbf{v}_i)]^T \boldsymbol{\Sigma}_c(\mathbf{r}_i)^{-1} [\boldsymbol{\mu}_c(\mathbf{r}_i) - \mathbf{I}(\mathbf{r}_i + \mathbf{v}_i)], \quad (15)$$

where $\boldsymbol{\mu}_c$ and $\boldsymbol{\Sigma}_c$ are the mean vector and covariance matrix for class c , respectively (predicted for the acquisition parameters of the

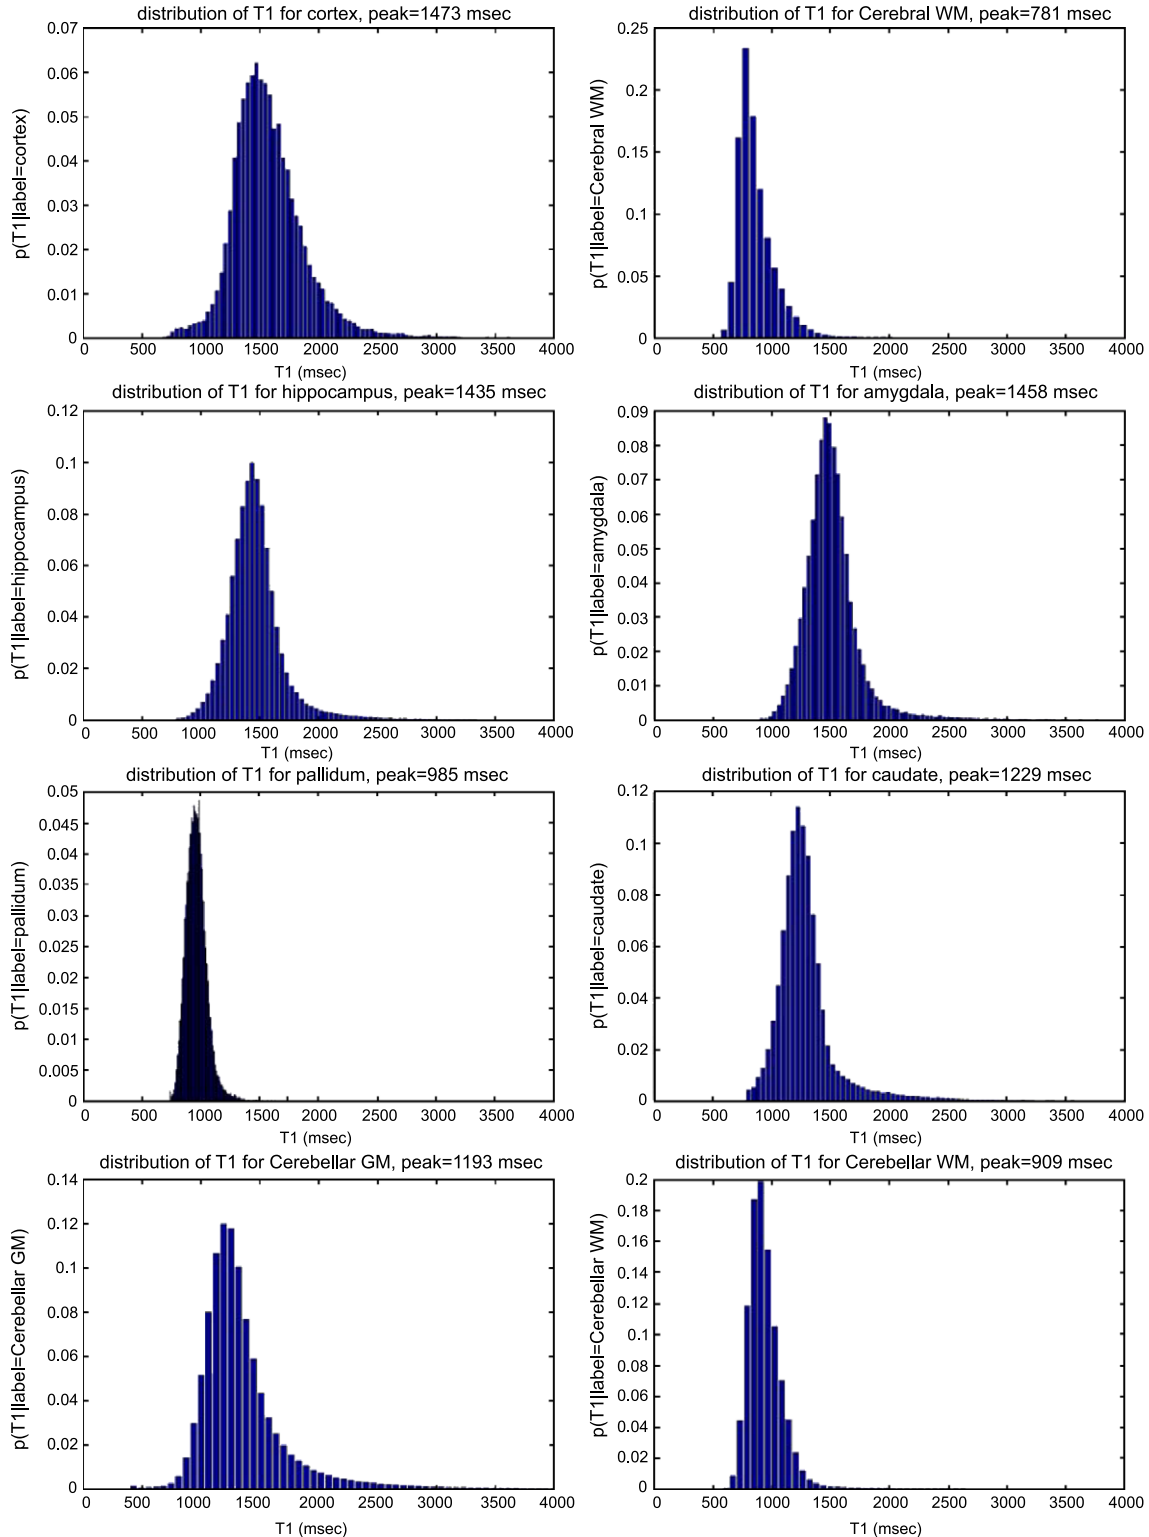


Fig. 5. T1 distributions by structure.

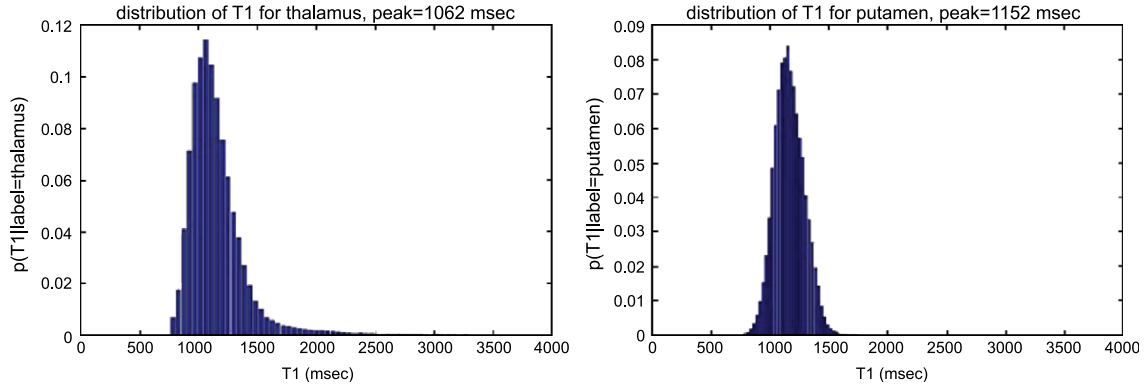


Fig. 5 (continued).

input scans, as detailed in the Cross-sequence prediction of means and covariance matrices section, and v_i is the vector field that we are seeking. The topology term G_T is given by:

$$G_T = \sum_{i=1}^V \left(\frac{\log(1 + e^{kR_i})}{k} - R_i \right), R_i = \frac{A_i^t}{A_i^0}, \quad (16)$$

where A_i^0 is the determinant of the Jacobian of the initial linear mapping, and A_i^t is the determinant of the current mapping. The nonlinearity in Eq. (16) only penalizes regions that are highly compressed or negative definite, and thus the coefficient on this term can be set to an arbitrarily large value as it only affects regions that are either noninvertible or almost noninvertible.

The metric distortion of the morph can be computed in a straightforward manner as the mean squared difference between the distance of the nodes at time t denoted d^t , and their original distance d^0 . In a similar manner, the smoothness of the deformation can be quantified as the difference between the vector at location i and the mean of its neighbors. We then modulate the metric term with the Kronecker delta δ_{ij} , in order to limit its application to within a neuroanatomical label, where $\delta_{ij} = 1$ if the highest probability class labels at locations r_i and

r_j are the same, and 0 otherwise. Denoting the set of Eq. (6) neighbors of each node by $N(i)$, these terms are given by:

$$G_M = \frac{1}{4V} \sum_{i=1}^V \sum_{n \in N(i)} \delta_{in} (d_{in}^t - d_{in}^0)^2, d_{in}^t = \|r_i + v_i^t - r_n + v_n^t\| \quad (17)$$

$$G_S = \frac{1}{2V} \sum_{i=1}^V \left\| v_i - \left(\frac{1}{N(i)} \sum_{n \in N(i)} v_n \right) \right\|^2. \quad (18)$$

Density-based morphing

Most volumetric intersubject morphing procedures assume an additive Gaussian noise model. For the scalar case, this gives rise to a log likelihood term that is the scalar analog of Eq. (15):

$$\log(p(I(\mathbf{r})|f, c) \propto -\left(\frac{I(f(\mathbf{r})) - \mu_c(\mathbf{r})}{\sigma_c} \right)^2 \quad (19)$$




Fig. 6. Average maps of T1 in cortical gray matter for the left hemisphere (left) and right hemisphere (right) across 10 subjects, shown in lateral (top), ventral (middle), and medial (bottom) views.

Differentiating Eq. (19) with respect to the atlas coordinate $f(\mathbf{r})$ we are seeking results in:

$$\frac{d \log(p(I(\mathbf{r})|f, c))}{df} \propto - \left(\frac{I(f(\mathbf{r})) - \mu_c(\mathbf{r}))}{\sigma_c} \right) \nabla I(f(\mathbf{r})) \frac{df}{d\mathbf{r}} \quad (20)$$

where the gradient is taken with respect to the coordinates $f(\mathbf{r})$. Eq. (20) simply states that if the image intensity at $f(\mathbf{r})$ is smaller than the atlas at \mathbf{r} , then f should be adjusted such that $f(\mathbf{r})$ maps to a new point \mathbf{r}' , where \mathbf{r}' is in the direction of $\nabla I[f(\mathbf{r})]$ with respect to \mathbf{r} . At first glance, one would think that discretizing Eq. (19) would

result in a similar minimization to discretizing Eq. (20). However, this is not the case, as implementing minimization techniques based on Eq. (20) can result in incorrect solutions, a problem not suffered by the direct discretization of Eq. (19). To see why this is the case, consider a one-dimensional registration problem in which the image consists of three voxels and the atlas consists of a single one (Fig. 2).

Further, assume the initial conditions are such that the atlas point with mean μ_c and unit variance is initially aligned with the central voxel. Given this configuration, let us examine the gradient

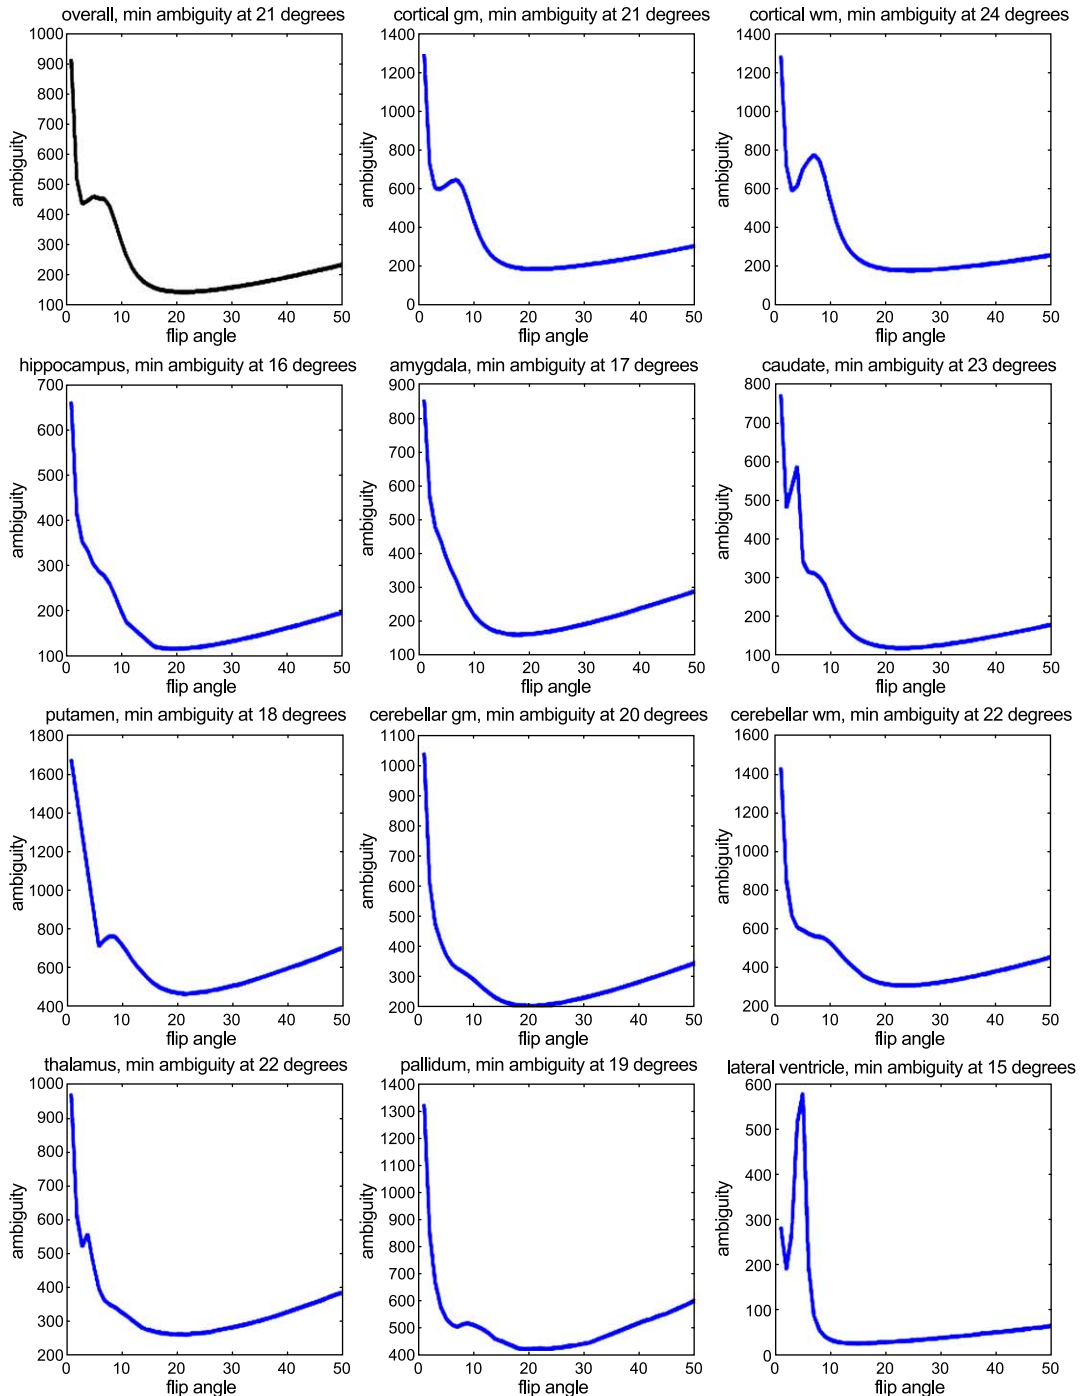


Fig. 7. Ambiguity for various brain structures versus flip angle.




Fig. 7 (continued).

of Eq. (19) given by direct discretization and interpolation of $-\log[p(I(\mathbf{r})|f,c)]$ versus discretizing Eq. (20) and computing $-\log[p(I(\mathbf{r})|f,c)]$ by interpolating image values, as is more commonly done. This situation is depicted in Fig. 3, with the interpolated image values shown in green, the direct interpolation of $-\log[p(I(\mathbf{r})|f,c)]$ in blue, and $-\log[p(I(\mathbf{r})|f,c)]$ computed from the interpolated image in red. Note that at the central point (voxel 1 on the x -axis), the gradient of the red and blue curves points in opposite directions. Specifically, interpolating the image first results in the movement of the atlas point towards the right, as the error functional decreases more sharply in that direction. This is a direct result of the assumption of continuity in the image—since voxel 1 has a value of $\mu_c/2$ and voxel 2 has a value of $2\mu_c$, continuity implies that somewhere between the two is a point with exactly the value of μ_c . Unfortunately, in Magn. Reson. Imaging, this will not in general be the case. In fact, partial volume effects frequently cause the truth to lie towards the opposite extreme. That is, if we assume that the true underlying image is made up of piecewise constant regions corresponding to different objects being imaged, with noise added and blurring induced by the imaging process (such as partial volume effects), then the intensity distribution induced by imaging the underlying objects is less continuous than the actual ones. In this case, interpolating image values can lead to incorrect results. Directly interpolating the probability field (or its log), however, avoids these problems as the assumption of continuity is not violated. Another way to look at this issue is that interpolation should be done after nonlinearities, resulting in the more reasonable interpolation of the blue line in Fig. 3, the gradient of which points in the correct direction.

Unfortunately resolving this problem requires a significant increase in computational complexity. If a scale-space approach is to be taken, the probability field (or its log) must be smoothed as opposed to the image. This implies that a smoothed probability field must be generated for every atlas point (as each one has different means and covariance matrices) at each step in the numerical minimization. Instead of replacing the likelihood term G_I in Eq. (15), we implement a preliminary version of the density-based morphing only for white matter regions inferior to the hippocampus and append it as a separate term G_L with coefficient λ_L in Eq. (15), as it is these regions that most frequently misalign due to very bright blood vessels inferior to cortex in this part of the brain. Note that this term is essentially identical to Eq. (5), but its numerical implementation and smoothing must be done quite differently, as a smoothed probability field must be computed for each atlas location separately, then its derivative with respect to a

change in the associated vector results in the delta that must be applied to the current vector field.

Eq. (15) is minimized using a multiscale line minimization, as outlined in (Fischl et al., 1999a). The tissue classes c are initially assumed to be the class with the highest prior probability at each location. After the minimization of Eq. (15) converges, the classes are reestimated to be those with the highest posterior density. The minimization is then repeated and the classes reestimated until the entire procedure converges, with the topology checked at each step of the integration, to insure invertibility.

Results

Intrinsic tissue parameter distributions

Preliminary results of the intrinsic tissue parameter estimation procedure are given in Fig. 4, which was generated by minimizing Eqs. (5) and (6), assuming that the noise is independent and identically distributed to estimate T1 and PD from five FLASH scans (fov = 240 mm × 240 mm; matrix = 256 × 256; slice thickness = 1.5 mm; TR = 18, 18, 20, 20, and 20 ms; $\alpha = 3^\circ, 5^\circ, 20^\circ, 20^\circ, \text{ and } 30^\circ$).

Distribution of T1 by tissue class

In this study, four FLASH images (1.5 T; TR = 20; TE = min; fov = 240 × 240; matrix = 256 × 256; $\alpha = 3^\circ, 5^\circ, 20^\circ, \text{ and } 30^\circ$) were acquired for each of eight healthy subjects (five males, three females, ages 21–37, mean 25.5 ± 6.1) and the parameter estimation procedure detailed in the intrinsic tissue parameter estimation section was applied to yield maps of proton density and T1. Each of the subjects was manually labeled using the procedures given in (Kennedy et al., 1993), and the T1 maps were combined with the segmentations in order to estimate the distribution of T1 on a structure by structure basis. The results of this study are given in Fig. 5. Note the variability in some structures, such as cortical gray matter, while others have very tight T1 distributions, such as the pallidum. The maximum likelihood estimate of the T1 for each structure is given in the title bar of each figure. There are a number of interesting points to note in this data set. First, the rightwards tail of many of the distributions is caused by the partial volume effects with CSF, which has an extremely long T1. This and the coloring effect of the nonlinearity involved in inverting Eq. (2) result in the non-Gaussian nature of the distributions, indicating the desirability of

classifying in the signal space where the Gaussian assumption is more valid.

Distribution of T1 across the cortex

The parameter estimation procedure can be combined with cortical surface models (Dale and Sereno, 1993; Fischl et al., 1999b; 2001) and high-resolution intersubject averaging procedures (Fischl et al., 1999a) in order to compute maps that reveal the spatial distribution of T1 across the cerebral cortex. In this study, 10 healthy subjects (four males, six females, ages 22–34, mean 25.7 ± 5.5) were scanned using the same protocol detailed above, and parameter maps were generated in the same way. Eq. (2) was then used to synthesize a strongly T1-weighted image with a flip angle of 30° , which was used as input to automated cortical surface reconstruction procedures (Dale and Sereno, 1993; Fischl et al., 1999b; 2001) to generate explicit models of both the gray/white junction and the pial surface. The T1 and proton density of the cortical gray matter were estimated at each position on the surface of each subject by sampling at the midpoint of the cortical ribbon defined by these surfaces. These T1 maps were then averaged using a spherical morphing procedure that aligns cortical folding patterns (Fischl et al., 1999a).

The results of this analysis are given in Fig. 6, for the left and right hemisphere (left and right), in lateral (top), inferior (middle), and medial (bottom) views of an average inflated surface. Note the spatial heterogeneity of the tissue characteristics across the surface, with limbic and association cortices having significantly longer T1 (and hence darker gray matter) than other cortical regions, agreeing with other published studies (Cho et al., 1997; Ogg and Steen, 1998; Steen et al., 2000). This spatial variability underscores the difficulty of the segmentation problem as the underlying characteristics of the tissue are changing across space in a way that cannot be accounted for by a multiplicative bias field correction.

There are several possible confounds to this analysis. The principle one is that the results reflect inhomogeneities in the RF transmit field, which would then masquerade as T1 variability due to the erroneous assumption that the flip angle is constant across space in Eq. (2). These inhomogeneities would derive from two sources—dielectric resonance, and the sensitivity profile of the transmit coil. While dielectric effects cannot be ruled out entirely, we think they are unlikely to be a major source of error in this study as it was carried out exclusively at 1.5 T where dielectric effects are minimal. In addition, the pattern of T1 variability is not what one would expect from dielectric resonance, in which a radially symmetric pattern is typical. Similarly, transmit inhomogeneities due to the coil are unlikely to be a major source of error as a body transmit coil with an extremely uniform spatial profile over the head was used in all data acquisition. An additional point to note is that the T1 estimates are completely insensitive to the more standard bias fields due to the sensitivity profile of the receive coil, as these effects, being multiplicative, are accounted for in the PD estimates.

Partial volume effects are harder to avoid, particularly in thin regions such as primary sensory ones. We minimize these effects by sampling at the midpoint of the cortical ribbon. One indication of the effectiveness of this technique is that the primary sensory areas typically show shorter T1 than the average, whereas partial volume effects should bias the estimates of the T1 in gray matter to be longer than the true value due to the extremely long T1 of CSF.

Sequence optimization

The sequence optimization procedure detailed in the Sequence optimization section is a quite general one. In this section, we present results for optimal acquisitions for one, two, and three standard single-echo FLASH images. In all cases, we hold TR and TE constant in (20 and 6 ms, respectively) order to fix the image noise and the flip angles induced by the vary RF excitation. Note that the noise parameter λ in Eq. (8) was varied across three orders of magnitude without substantially changing the results and was ultimately fixed at $\lambda = 100$ for all studies presented here.

Optimal single angle acquisition

One important advantage of phrasing the sequence optimization as in Eq. (12) is that one has the freedom to restrict the classes c_1 and c_2 that are considered. In this section, we carry out an optimization for a single FLASH acquisition in order to determine the optimal flip angle for segmentation purposes. This is done across all classes, as well as on a class-by-class basis, in order to determine, for example, the optimal flip angle for segmenting hippocampus from all other structures that it abuts. It is important to note that the minimum ambiguity for each structure is not simply determined by the flip angle that maximizes the signal-to-noise ratio (SNR) of that tissue class, which would be given by the Ernst angle. For example, for cortical gray matter, which has a T1 of approximately 1300 ms at 1.5 T, the Ernst angle for a TR of 20 ms would be 10.0° . However, the optimal flip angle for segmenting cortical gray matter from other structures that it is adjacent to is 21° , reflecting the T1/PD characteristics of cerebral white matter, CSF, and dura that the gray matter must be distinguished from.

Fig. 7 gives the ambiguity measure for a set of structures for flip angles from 0° to 50° , with the optimum given in the title of each figure. Most structures share a characteristic U-shaped curve, with very low flip angles resulting in very low signal and hence high ambiguity. The ambiguity then typically reduces until the global minimum is reached between 15° and 25° . The upward spike in

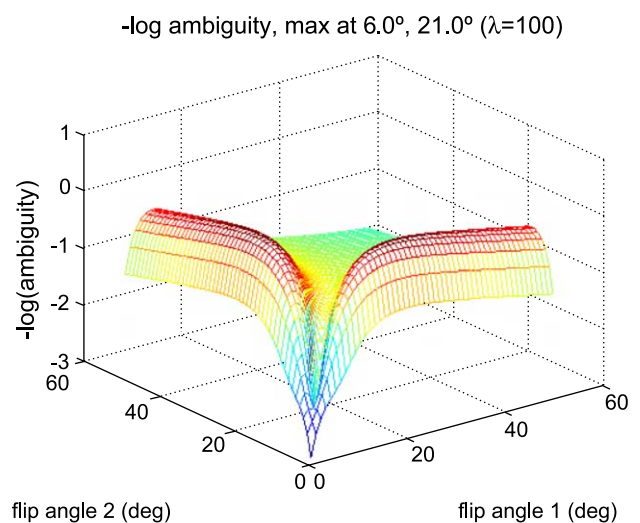


Fig. 8. $-\log(A)$ measure for two FLASH scans with different flip angles. The minimum ambiguity (a maximum in the $-\log(A)$ plot) occurs at the location $6^\circ, 21^\circ$ indicating that having different contrast properties in the two images helps the segmentation procedure (the 6° would be strongly proton density weighted, and the 21° strongly T1 weighted).




Fig. 9. MEF acquisitions. Top: proton density-weighted scan on a 1.5-T Siemens Sonata (TR = 20 ms; flip angle = 5°; band width = 615 Hz/voxel; TE = 1.85, 3.67, 5.49, 7.31, 9.13, 10.95, 12.77, 14.59 ms). Bottom: T1-weighted scan (flip angle = 30°, other parameters same as above).

many of the plots that occurs at approximately 7° or 8° no doubt reflects the isointensity point of gray and white matter. In contrast, the ventricular plots, particularly the lateral ventricles and the third ventricle, show increased ambiguity at lower flip angles due to the gray or CSF isointensity point, which occurs at about 5°. This effect is also apparent in the ambiguity curve for the caudate due to its proximity to the lateral ventricles for much of its extent, again highlighting the fact that the optimal acquisition parameters reflect the properties of a structure, and the properties of the structures that occur near it.

Optimal dual-angle acquisition

Most morphometry studies have sufficient acquisition time for at least two 3-D structural images to be acquired. These images most commonly have the same pulse parameters, allowing them to be easily motion corrected and averaged, resulting in a $\sqrt{2}$ improvement in the noise. In order to investigate the efficiency of this approach, we again held TR and TE constant, while allowing two flip angles to vary from 0° to 60°. We fixed TR to be 20 ms, resulting in an acquisition time of approximately 16.5 min. The results of this analysis are given in Fig. 8, which plots $-\log(A)$ for pairs of flip angles. Note the shoulder in the graph that runs parallel to the x - and y -axes, indicating the optimum pair of flip angles is composed of a low (6°) and a high (21°) angle. Intuitively, this makes sense, as the low and high flip angle images contain qualitatively different information, being strongly proton density and T1-weighted, respectively. The ambiguity is relatively insensitive to the exact value of the larger flip angle, as evidenced by the gentle slope of the plot in the region around 20°.

Optimal triple-angle acquisition

In this study, we searched the full 11-dimensional space to find the optimal triplet of flip angles for segmentation purposes. This is a difficult global optimization problem due to its high-dimensional nature. That is, three flip angles must be searched over, each of which contains an eight-dimensional optimization (two classes, three spatial dimensions, and three intensity dimensions). Carrying this calculation out for the >8000 flip angle triples required over 80,000 h of compute time on our 200 node Beowulf cluster, with the optimum occurring at 5°, 16°, and 31°. As might be expected, the addition of a third flip angle occurs in the middle of the previous two and pushes them outwards so that the low flip angle is lowered and the higher one is increased.

Multiecho FLASH

An example of images acquired using the multiecho FLASH sequence is given in Fig. 9, which shows the echoes (left to right) of a 5° (top) and a 30° (bottom scan), with increasing echo numbers going from left to right. The sequence was run in 3-D mode with timing approximately as shown in Fig. 1 (note that the duration of the readout ramps has been exaggerated for visualization purposes). The sequence has eight echoes following each excitation pulse, a band width of 651 Hz/pixel for each readout, a voxel size of $1.3 \times 1.0 \times 1.3$ mm with a field of view of 256 mm, and 128 partitions (192×256 imaging matrix). Acquisition time is 8 min and 12 s.

An additional advantage of this sequence is that the acquisition of data at varying echo times allows the explicit estimation of T2* from a single 8 min scan set. Examining the two image sequences in Fig. 9, the T2* decay is apparent, as the images get uniformly darker from left to right. Using a log-linear fitting procedure, we used manually drawn ROIs in the cortical gray matter and the underlying white matter to estimate T2* for each tissue class for both 1.5 and 3 T scanners. The results of this analysis are given in Table 1 for 1.5 T (top row) and 3 T (bottom row).

Optimizing multiecho FLASH for segmentation

In order to gain some intuition as to the contrast properties of the different echoes in the MEF acquisition, we manually labeled an ROI in the cortical gray matter and another in the cerebral white matter. The ROIs were chosen in perirolandic regions where the contrast-to-noise ratio (CNR) between gray and white is reduced relative to other cortical regions due to intrinsic variations in the tissue parameters. The CNR between gray and white matter in this region is quantified in Fig. 10 for the 30° (left) and 5° scans for each of the eight echoes.

As can be seen from these plots, the gray/white CNR generally decreases with increasing echo number for the 30° scan and increases for the 5° scan. We then used the T2* values tabulated in

Table 1
T2* values for gray matter and white matter at 1.5 and 3 T

Field strength (T)	Gray matter T2* (ms)	White matter T2* (ms)
1.5	75	55
3	68	53

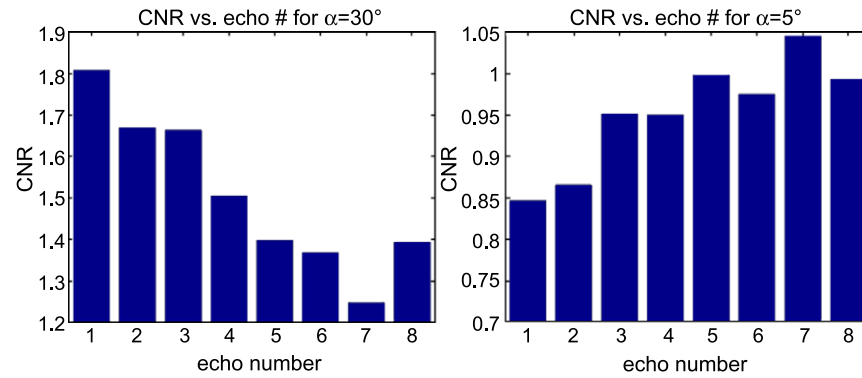


Fig. 10. CNR between gray and white matter as a function of echo no. for 30° (left) and 5° images (right).

Table 1 and Eq. (2) to simulate the evolution of the contrast properties of the different flip angles with increasing echo time. The results of this analysis are given in Fig. 11. Examining this figure, it is apparent that the longer $T2^*$ of GM results in it darkening more gradually than WM, implying that for scans in which it is initially darker (such as the 30°), $T2^*$ decay will reduce CNR at longer echo times. Conversely for scans in which GM is initially brighter than WM (such as the 5°), the converse holds and the $T2^*$ effect will increase the CNR at longer echo times. An additional observation of note is that much of the contrast in a “proton-density-weighted image” seems to be due to $T2^*$ decay, implying longer TEs will increase CNR for low flip angle FLASH scans. An additional implication of these results is that the optimal method of combining the different echoes is not simply to average them but rather to weight them unequally based on their contrast. In future work, we will implement an optimization procedure to minimize the ambiguity A given by Eq. (11) as a function of echo weighting.

Multiecho FLASH test–retest stability

In order to assess the test–retest repeatability of the multiecho FLASH acquisition protocol, we acquired four images for each of three different acquisition types [single echo FLASH, MP-RAGE (Mugler, 1999), and multiecho FLASH] within a single scan session. The images were each rigidly aligned to the first acquisition of each type (Jenkinson and Smith, 2001). Nonbrain tissue in these images was removed using a hybrid watershed or deformable surface algorithm (Segonne et al., in press), resulting in four coregistered brain images for each structural image type. The intensity values in each scan were then scaled to be a fraction

of the mean brain intensity value, and the cross-scan intensity standard deviations were computed at each voxel. The mean of these values across the brain represents the stability of the intensity of each acquisition type, from scan to scan. The results of this analysis are given in Fig. 12 for standard MP-RAGE (left), FLASH (center), and multiecho FLASH (right). Note that the standard errors of the mean are too small to see on this plot, indicating that all the results are significant (due mainly to the large number of degrees of freedom in the standard error calculation—one for each brain voxel). As can be seen, the test–retest intensity variation is significantly reduced in the MEF acquisitions, a critical issue for longitudinal analysis of morphometric changes. Presumably, the reduction in the variance in intensity values is due to the high band width of the MEF sequence, whereas the higher variability of the MP-RAGE is due to its relatively low SNR with respect to FLASH acquisitions (due among other factors to its high band width).

Nonlinear density-based morphing

As pointed out in the Nonlinear transform section, discretizing the gradient of the standard log likelihood term commonly used in nonlinear image registration can result in incorrect results. This is due to the assumption of spatial continuity implicit in this derivation. In this formulation, extremely bright structures in the image will attract atlas regions that are initially aligned with low-intensity portions of the image. Computing the gradient of the discretized log likelihood instead can resolve this problem, as illustrated in Fig. 13, which shows a representation of the atlas (top left), the initial image after optimal linear alignment (top right), and

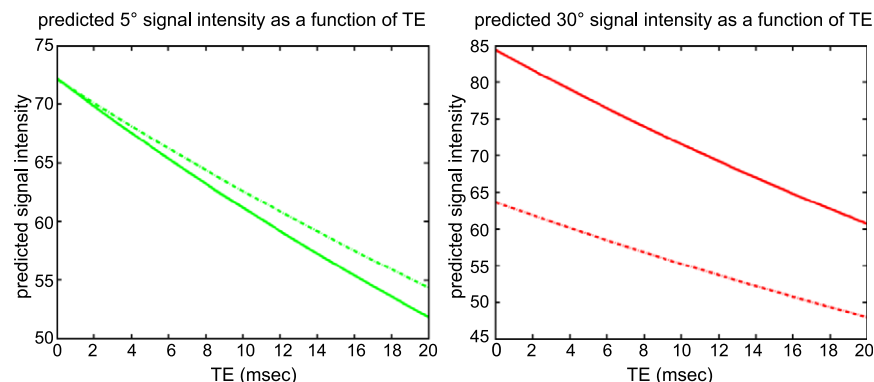


Fig. 11. $T2^*$ decay curves for gray matter (solid) and white matter (dotted) at 30° (left) and 5° (right).

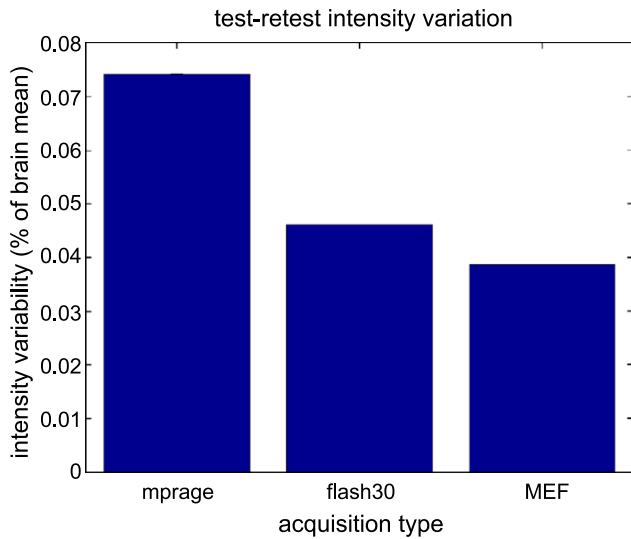


Fig. 12. Test–retest intensity variation for three different structural acquisition protocols.

the final aligned image without (bottom left) and with (bottom right) a term discretizing the likelihood instead of its derivative. Note the distortions induced by very bright blood vessels in inferior temporal regions. Computing the gradient after discretization (bottom right) instead of before (bottom left) can resolve this issue, as atlas voxels are attracted to regions of high probability instead of along the intensity gradient.

It is important to note that this problem is exacerbated if a standard scale-space approach is taken and ∇I is convolved with a

large Gaussian in the initial stages of the numerical minimization. In this case, the interpolation of gradient values is carried out over long distances, resulting in effects such as those seen in Fig. 13: All atlas near the edge of the brain or the ventricles that are initially aligned with image locations brighter than $\mu_c(\mathbf{r})$ are drawn towards the outside of the brain or towards the ventricles as the image intensity gradient points in these directions. Note that the presence of the label term eliminates this effect and properly aligns the temporal white matter, even in the presence of dramatically enlarged ventricles (bottom right).

Cross-sequence segmentation

In order to test the sequence independence of the segmentation procedure, nine FLASH scans were acquired on a single subject (TR = 20 ms; TE = 6 ms; flip angles $\alpha = 30^\circ, 2^\circ, 15^\circ, 3^\circ, 10^\circ, 20^\circ, 4^\circ, 7^\circ,$ and 25°). Each set of three consecutive scans (consisting of a low, middle, and high flip angle) was used to separately estimate β using Eqs. (5) and (6). The forward Bloch Eq. (2) was then used to synthesize an image for each set with $\alpha = 30^\circ$. These images were labeled using the technique described in (Fischl et al., 2002), and the volumes of 14 major brain structures were computed for each of the three data sets. The results of this study are given in Fig. 14. As can be seen, the mean difference in the structure volumes between the scans was found to be 1.8%, indicating the insensitivity to the image contrast found in the original data.

In a final study, four FLASH images (1.5 T; TR = 20 ms; TE = min; fov = 240×240 ; matrix = 256×256 ; $\alpha = 3^\circ, 5^\circ, 20^\circ,$ and 30°) were acquired for each of eight healthy subjects (five males, three females, ages 21–37, mean 25.5 ± 6.1). A jackknifing




Fig. 13. Distortions induced by intensity-based morphing that can be eliminated with discretization of the probability field. Top left: target image. Top right: initial image after optimal linear alignment. Bottom left: morph results with $\lambda_L = 0$. Bottom right: morph results with $\lambda_L = 1$. Note that the red cross-hair represents corresponding points across all images, and the green arrow indicates the location of a blood vessel in the atlas. $\lambda_S = 1, \lambda_T = 1, \lambda_L = 0$ (bottom left) or $\lambda_L = 1$ (bottom right), $\lambda_M = 0, k = 5$ in Eq. (14).

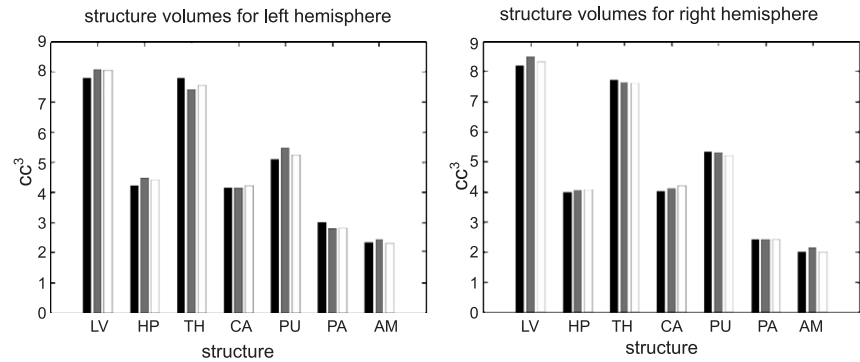


Fig. 14. Test–retest structure volumes measured from three separate data sets (dark, medium, and light bars) acquired on the same subject. Note that each data set had different acquisition parameters (LV = lateral ventricle, HP = hippocampus, TH = thalamus, CA = caudate, PU = putamen, PA = pallidum, AM = amygdala).

procedure was used to generate an atlas for each subject by using the remaining seven subjects' data. In order to simulate the variability in acquisition parameters, each subject's four FLASH scans were grouped into four sets of two, consisting of a high and a low flip angle ($3^\circ, 20^\circ$), ($3^\circ, 30^\circ$), ($5^\circ, 20^\circ$), ($5^\circ, 30^\circ$). Each of these scan pairs was then registered and labeled using Eqs. (7) and (8) to map the atlas into the appropriate FLASH space. The mean and standard errors across subjects for four brain structures for each of the scan pairs are given in Fig. 15, together with the mean and standard error of the manual labeling (in blue). The agreement of the structure volumes across flip angle is in general quite good, with a mean variation of less than 3% (standard deviation/mean = 2.9%, mean across structures). The residual variability may be due to poorly regularized estimates of the covariance structure. This leads to ill-conditioned inverse covariance matrices being used in the estimation of the image likelihood and hence errors in the MAP estimates of the segmentation. This type of problem may be improved by regularization techniques such as those suggested in Efron and Morris (1976), Haff (1979), Rayens and Greene (1991), and Tadjudin and Landgrebe (1999). In addition, there are many

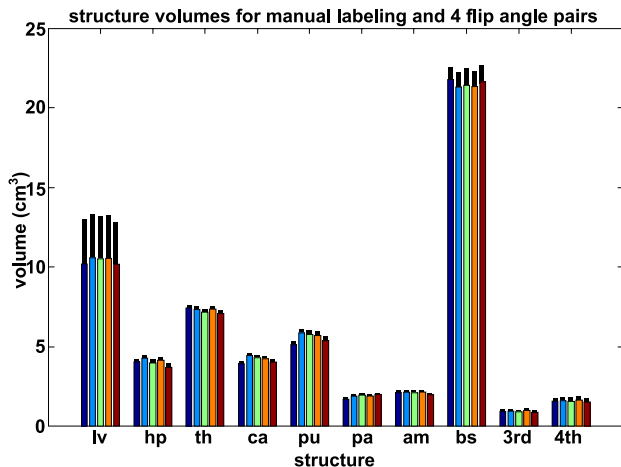


Fig. 15. Structure volumes for different pairs of flip angles across eight subjects. Structures: lv = lateral ventricles, hp = hippocampus, th = thalamus, ca = caudate, pu = putamen, pa = pallidum, am = amygdala, bs = brainstem, 3rd = third ventricle, 4th = fourth ventricle (blue = manual; cyan = $3^\circ, 20^\circ$; green = $3^\circ, 30^\circ$; orange = $5^\circ, 20^\circ$; red = $5^\circ, 30^\circ$).

MRI effects not modeled by Eq. (2) such as magnetization transfer, inflow effects, etc., which may also contribute to the residual variability.

Conclusion

Current anatomical studies are frequently limited by the requirement that the initial MR protocol used be held fixed throughout the length of the study. This is because automated and semiautomated techniques are typically only optimized for a given MR pulse sequence. Thus, measurements of structural properties may vary as a function of the particular sequence used in acquiring the data. This is of course undesirable, particularly for longitudinal studies that are then unable to take advantage of advances in imaging technology.

The ability to combine data across sites in neuroimaging studies would also be of great potential utility, as it would facilitate the analysis of large subject populations. Unfortunately, this type of cross-site study is frequently hampered by the lack of standard parameter settings for structural MRI acquisition. This problem is again compounded by the dependence of postprocessing software on a particular set of acquisition parameters.

High-resolution structural MR images are usually generated from so-called "T1-weighted" pulse sequences. Nevertheless, all structural scans also image the density of protons in the underlying tissue, as well as the transverse relaxation time $T2^*$ (or $T2$ for spin-echo sequences). Thus, it is impossible to infer from changes in a single type of structural scan which of these tissue parameters (i.e., $T1$, $T2$, $T2^*$, and proton density) is responsible for the variation. T1-weighted morphometry must therefore frequently be solely phenomenological, or descriptive, as opposed to providing a more comprehensive understanding of disease processes. As part of the current work, we have extended a standard clinical sequence (FLASH) in order to allow the estimation of transverse relaxation time $T2^*$. An additional benefit of this sequence is that it has a high band width, essentially eliminating geometric distortions and signal loss due to magnetic susceptibility interfaces. The multiecho nature of the sequence allows the high band width to be achieved without sacrificing SNR and has the added benefit of reducing test–retest intensity variability.

In this paper, we have presented a set of techniques to address the problems associated with the sequence dependence of morphometry tools. First, we derived and implemented a parameter

estimation procedure that yields measures of intrinsic tissue properties rather than the somewhat arbitrary intensity of an MR image. These parameters can be derived from a set of standard structural scans (SPGR or FLASH) and are properties of the tissue and hence insensitive to changes in the acquisition parameters. We then embedded the solution to the steady state Bloch equations into the registration and segmentation procedures, allowing the prediction of the intensity distributions of the underlying structures for any set of MR acquisition parameters. The segmentation and registration thus use the physics of image formation in order to map an atlas into the signal space of a novel image. This results in procedures that are largely invariant to variations in the acquisition details, as these variations are explicitly modeled.

The explicit estimation of the underlying tissue parameters may also be beneficial for detecting pathology-related changes in T2*, T1, and proton density. In particular, these changes may be correlated but have opposite effects on image contrast, making them difficult or impossible to detect from a single T1-weighted image. This type of alteration in the imaging properties of the tissue, possibly due to processes such as cell death or gliosis, may entail changes in the parameters but not in morphometric properties such as volume. Variations of this nature would thus be undetectable to a standard morphometric analysis procedure but would be clearly revealed by cross-group comparison of the tissue parameters.

Another advantage of embedding models of the physics of image formation into the segmentation procedure is that it allows the explicit optimization of the imaging parameters for the segmentation to be used on them. This can be carried out for all tissue classes or can be naturally restricted to a subset of the structures that can be labeled. This optimization reveals that ambiguity in the segmentation is minimized by the conjunction of proton density and T1-weighted images, in contrast to the more standard practice of acquiring two strongly T1-weighted images and averaging them.

Acknowledgments

Support for this research was provided in part by the National Center for Research Resources (P41-RR14075, R01 RR16594-01A1, and the NCCR BIRN Morphometric Project BIRN002), as well as the Mental Illness and Neuroscience Discovery (MIND) Institute. We thank Evelina Busa for extensive testing of the proposed algorithms.

References

- Ashburner, J., Neelin, P., Collins, D.L., Evans, A.C., Friston, K.J., 1997. Incorporating prior knowledge into image registration. *NeuroImage* 6, 344–352.
- Baker, J.R., 1991. *Spatially Variant Tomographic Imaging: Estimation, Identification, and Optimization*. University of California, Berkeley.
- Chen, N.-k., Wyrwicz, A.M., 1999. Correction for EPI distortions using multi-echo gradient-echo imaging. *Magn. Reson. Med.* 41, 1206–1213.
- Cho, S., Jones, D., Reddick, W.E., Ogg, R.J., Steen, R.G., 1997. Establishing norms for age-related changes in proton T1 of human brain tissue in vivo. *Magn. Reson. Imaging* 15 (10), 1133–1143.
- Collins, D.L., Neelin, P., Peters, T.M., Evans, A.C., 1994. Data in standardized Talairach space. *J. Comput. Assist. Tomogr.* 18 (2), 205–292.
- Constable, R.T., Henkelman, R.M., 1991. Contrast, resolution, and detectability in MR imaging. *J. Comput. Assist. Tomogr.* 15, 297.
- Constable, R.T., Skudlarski, P., Gore, J.C., 1995. An ROC approach for evaluating functional brain MR imaging and postprocessing protocols. *Magn. Reson. Med.* 34, 57–64.
- Dale, A.M., Sereno, M.I., 1993. Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: a linear approach. *J. Cogn. Neurosci.* 5 (2), 162–176.
- Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis. I: segmentation and surface reconstruction. *NeuroImage* 9, 179–194.
- Davatzikos, C., 1997. Spatial transformation and registration of brain images using elastically deformable models. *Comput. Vis. Image Underst.* 66 (2), 207–222.
- Edelstein, W.A., Glover, G.H., Hardy, C.J., Redington, R.W., 1986. The intrinsic signal-to-noise ratio in NMR imaging. *Magn. Reson. Med.* 3, 604.
- Efron, B., Morris, C., 1976. Multivariate empirical Bayes and estimation of covariance matrices. *Ann. Stat.* 4, 22–32.
- Epstein, F., Mugler, J.I., Brookeman, J., 1994. Optimization of parameter values for complex pulse sequences by simulated annealing: application to 3D MP-RAGE imaging of the brain. *Magn. Reson. Med.* 31, 164–177.
- Fischl, B., Sereno, M.I., Dale, A.M., 1999a. Cortical surface-based analysis. II: inflation, flattening, and a surface-based coordinate system. *NeuroImage* 9 (2), 195–207.
- Fischl, B., Sereno, M.I., Tootell, R.B.H., Dale, A.M., 1999b. High-resolution inter-subject averaging and a coordinate system for the cortical surface. *Hum. Brain Mapp.* 8 (4), 272–284.
- Fischl, B., Liu, A., Dale, A.M., 2001. Automated manifold surgery: constructing geometrically accurate and topologically correct models of the human cerebral cortex. *IEEE Trans. Med. Imaging* 20 (1), 70–80.
- Fischl, B., Salat, D.H., Albert, M., Dieterich, M., Haselgrove, C., Kouwe, A.v.d., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., Dale, A.M., 2002. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* 33 (3), 341–355.
- Fox, P.T., Perlmutter, J.S., Raichle, M.E., 1984. Stereotactic method for determining anatomical localization in physiological brain images. *J. Cereb. Blood Flow Metab.* 4 (4), 634.
- Fox, P.T., Perlmutter, J.S., Raichle, M.E., 1985. A stereotactic method of anatomical localization for positron emission tomography. *J. Comput. Assist. Tomogr.* 9 (1), 141–153.
- Geman, S., Geman, D., 1984. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* 6, 721–741.
- Grief, W.L., Buxton, R.B., Lauffer, R.B., Saini, S., Stark, D.D., Wedeen, V.J., Rosen, B.R., Brady, T.J., 1985. Pulse sequence optimization for MR imaging using a paramagnetic hepatobiliary contrast agent. *Radiology* 157, 461–466.
- Haacke, E.M., Lenz, G.W., 1987. Improving MR image quality in the presence of motion by using rephasing gradients. *Am. J. Roentgenol.* 148 (6), 1251–1258.
- Haff, L.R., 1979. Estimation of the inverse covariance matrix: random mixtures of the inverse Wishart matrix and the identity. *Ann. Stat.* 7, 1264–1276.
- Jenkinson, M., Smith, S., 2001. A global optimisation method for robust affine registration of brain images. *Med. Image Anal.* 5 (2), 143–156.
- Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage* 17 (2), 825–841.
- Kadah, Y.M., Hu, X., 1998. Algebraic reconstruction for Magn. Reson. Imaging under B0 inhomogeneity. *IEEE Trans. Med. Imaging* 17 (3).
- Kapur, T., Grimson, W.E.L., Wells, W.M., Kikinis, R., 1998. Enhanced spatial priors for segmentation of magnetic resonance imagery. First International Conference on Medical Image Computing and Computer-Assisted Intervention, Cambridge, MA.

- Kennedy, D.N., Meyer, J.W., Filipek, P.A., Caviness, V.S., 1993. MRI-based topographic segmentation. In: Thatcher, R.W., et al. (eds.), *Functional Neuroimaging: Technical Foundations*. Orlando, Academic Press Inc.
- Killiany, R.J., Gomez-Isla, T., Moss, M., Kikinis, R., Sandor, T., Jolesz, F., Tanzi, R., Jones, K., Hyman, B.T., Albert, M.S., 2000. The use of structural MRI to predict who will get Alzheimer's disease. *Ann. Neurol.* 47, 430–439.
- Mugler, J.P.I., 1999. Improved three-dimensional GRASE imaging with the SORT phase-encoding strategy. *J. Magn. Reson. Imaging* 9, 604–612.
- Ogg, R.J., Steen, R.G., 1998. Age-related changes in brain T1 are correlated with iron concentration. *Magn. Reson. Med.* 40 (5), 749–753.
- Pham, D.L., Prince, J.L., Dagher, A.P., Xu, C., 1997. An automated technique for statistical characterization of brain tissues in Magnetic Resonance Imaging. *Int. J. Pattern Recognit. Artif. Intell.* 11 (8), 1189–1211.
- Prince, J.L., Tan, Q., Pham, D., 1995. Optimization of MR pulse sequences for Bayesian image segmentation. *Med. Phys.* 22 (10), 1651–1656.
- Puri, B.K., Richardson, A.J., Oatridge, A., Hajnal, J.V., Saeed, N., 1999. Cerebral ventricular asymmetry in schizophrenia: a high resolution 3D Magn. Reson. Imaging study. *Int. J. Psychophysiol.* 34, 207–211.
- Rayens, W., Greene, T., 1991. Covariance pooling and stabilization for classification. *Comput. Stat. Data Anal.* 11, 17–42.
- Schmithorst, V.J., Dardzinski, B.J., Holland, S.K., 2001. Simultaneous correction of host and geometric distortion artifacts in EPI using a multiecho reference scan. *IEEE Trans. Med. Imaging* 20 (6).
- Schmitt, F., 1985. Correction of geometrical distortions in MR-images. *Computer Assisted Radiology CAR*. Heidelberg, Springer Verlag.
- Segonne, F., Dale, A.M., Busa, E., Glessner, M., Salvolini, U., Hahn, H.K., Fischl, B., 2004. A hybrid approach to the skull-stripping problem in MRI. *NeuroImage* (in press).
- Steen, R.G., Reddick, W.E., Ogg, R.J., 2000. More than meets the eye: significant regional heterogeneity in human cortical T1. *Magn. Reson. Imaging* 18 (4), 361–368.
- Tadjudin, S., Landgrebe, D.A., 1999. Covariance estimation with limited training samples. *IEEE Trans. Geosci. Remote Sens.* 37 (4), 1–13.
- Venkatesan, R., Haacke, E.M., 1997. Role of high resolution in magnetic resonance (MR) imaging: applications to MR angiography, intracranial T1-weighted imaging, and image interpolation. *Int. J. Imaging Syst. Technol.* 8, 529.
- Viola, P., Wells, W., 1995. Alignment by maximization of mutual information. 5th International Conference on Computer Vision.
- Wells, W., Grimson, W., Kikinis, R., Jolesz, F., 1996. Adaptive segmentation of MRI data. *IEEE Trans. Med. Imaging* 15 (4), 429–442.
- Wolf, H., Grunwald, M., Kruggel, F., Riedel-Heller, S.G., Angerhofer, S., Hojjatoleslami, A., Hensel, A., Arendt, T., Gertz, H.-J., 2001. Hippocampal volume discriminates between normal cognition; questionable and mild dementia in the elderly. *Neurobiol. Aging* 22, 177–186.
- Woods, R.P., Cherry, S.R., Mazziotta, J.C., 1992. Rapid automated algorithm for aligning and reslicing PET images. *J. Comput. Assist. Tomogr.* 16 (4), 620–633.
- Zhang, Y., Brady, M., Smith, S., 2001. Segmentation of Brain MR Images through hidden Markov Random Field model and the expectation maximization algorithm. *IEEE Trans. Med. Imaging* 20 (1), 45–57.