Estimating intensity variance due to noise in registered images: Applications to diffusion tensor MRI

Gustavo K. Rohde,[a,b,]* Alan S. Barnett,[c] Peter J. Basser,[a] and Carlo Pierpaoli[a]

[a]STBB/LIMB/NICHD, National Institutes of Health, Building 13, Room 3w16, 13 South Drive, Bethesda, MD 20892, USA
[b]Applied Mathematics and Scientific Computation Program, University of Maryland, College Park, MD 20742, USA
[c]NIMH, National Institutes of Health, Bethesda, MD 20892, USA

Image registration techniques which require image interpolation are widely used in neuroimaging research. We show that signal variance in interpolated images differs significantly from the signal variance of the original images in native space. We describe a simple approach to compute the signal variance in registered images based on the signal variance and covariance of the original images, the spatial transformations computed by the registration procedure, and the interpolation or approximation kernel chosen. The method is general and could handle various sources of signal variability, such as thermal noise and physiological noise, provided that their effects can be assessed in the original images. Our approach is applied to diffusion tensor (DT) MRI data, assuming only thermal noise as the source of variability in the data. We show that incorrect noise variance estimates in registered diffusion-weighted images can affect DT parameters, as well as indices of goodness of fit such as chi-square maps. In addition to DT-MRI, we believe that this methodology would be useful any time parameter extraction methods are applied to registered or interpolated data, such as in relaxometry and functional MRI studies.
© 2005 Elsevier Inc. All rights reserved.

Keywords: Image registration; Interpolation; Noise; Variance; Parameter estimation

Introduction

Post-acquisition image alignment (registration) is routinely performed in biomedical research and clinical practice (Maintz and Viergever, 1998; Pluim et al., 2003). Applications using image registration techniques include motion and distortion correction in functional MRI (fMRI), diffusion tensor MRI (DT-MRI), and MR relaxometry experiments. In addition, image registration procedures are increasingly being used in computational-based studies of neuroanatomy. This involves understanding the variability of tissue properties, including shape, across specific populations. An example is voxel-based morphometry, described in Ashburner and Friston (2000).

In general, many of the current post-processing methodologies can be summarized as follows. A set of medical images is acquired and reconstructed using standard methodologies. This step may include filtering to avoid "ringing" artifacts, denoising, intensity corrections, etc. Next, using one of many available algorithms, images are registered to ensure, as much as possible, that a fixed image coordinate corresponds to the same structure, or anatomical coordinate, in all images acquired. This step is necessary because the subject being imaged may move during data acquisition. In addition, images may contain geometric distortions with respect to each other. In echo planar (EPI) MRI, these distortions can be caused by magnetic field susceptibility related artifacts. In EPI-based diffusion-weighted imaging, significant geometric distortions may also occur due to eddy-currents induced by the rapidly switched diffusion weighting magnetic field gradients applied during imaging. Corrections to account for such misregistration artifacts are absolutely necessary to ensure the data analysis is reliable. In addition to correcting for motion and geometric distortions, the entire image sequence may also be aligned to a standard template image, using stereotaxic normalization techniques, for example, so that the data analysis results can be more conveniently interpreted. Data analysis consists of extracting or estimating some physically meaningful parameters from the sequence of medical images. In DT-MRI, a 3×3 symmetric diffusion tensor is estimated, based on which several other quantities such as measures of diffusion anisotropy and depictions of fiber tracts can be generated. In fMRI, these may be statistical parametric maps (Friston et al., 1995), for example.

In many of these applications, the analysis of the registered images involves fitting or estimating model parameters from the intensity values of the images. For such tasks, it is crucial to know the correct signal variance of the registered images so that least-squares procedures, for example, can be properly implemented.

* Corresponding author. Fax: +1 301 435 5035.
 E-mail address: rohdeg@math.umd.edu (G.K. Rohde).
 Available online on ScienceDirect (www.sciencedirect.com).

Though significant research has been devoted to estimating signal variance in medical images–some examples in MRI include Gudbjartsson and Patz (1995), Henkelman (1985), and Sijbers et al. (1998) among others–it is important to recognize that the signal variances in the registered and the original unregistered images differ. This is because the image interpolation or approximation step generally required in image registration can, as will be shown later, significantly change the noise properties of the image. We will show how a simple formula can be used to compute the appropriate signal variance in registered images. The analysis of diffusion-weighted MRI data using the diffusion tensor model will be used as a case study. That is, given a set of diffusion-weighted MR images (DWI), we use an existing software to register the DWIs to remove rigid body motion and eddy-current-related distortions prior to tensor computation. We then show that noise variance in the registered images differs from the noise variance in the original images. However, even though DT-MRI is the only application discussed in detail in this paper, we believe that the general approach described in this paper should be considered whenever registered images are being analyzed using procedures that require knowledge of the variance in the image intensity values.

At the time of writing not much related work can be found in the biomedical imaging literature. Friston et al. (1996) address the problem of removing movement-related artifacts, such as those caused by intensity fluctuations due to the change in position of the imaged object with respect to the reference frame of the scanner. In Thacker et al. (1999) and Grootoonk et al. (2000), the authors investigate the error in the intensity values produced by interpolation procedures applied on the registered images. Maas and Renshaw (1999) discuss artifacts related to high frequency losses on registered (interpolated) data. Pluim et al. (2000) report that interpolation methods may cause undesirable artifacts when estimating the mutual information similarity measure. Nickerson et al. (2003) describe a method through which the local intensity variance in positron emission tomography (PET) can be estimated from the operations performed during image reconstruction. None of these works, however, detail the importance of, and methods for, obtaining correct estimates of the signal variance at each coordinate of each registered image.

In the field of diffusion-weighted imaging and diffusion tensor MRI, several researchers have investigated methods for performing post-acquisition motion and distortion correction of data (Andersson and Skare, 2002; Bastin, 1999; Haselgrove and Moore, 1996; Horsfield, 1999; Mangin et al., 2002; Rohde et al., 2004). Though the registration methods differ, most of these works use linear interpolation to produce the series of DWIs. This series is then used to estimate one diffusion tensor for each voxel via least-squares fitting procedures similar to the χ^2 minimization procedure described in Basser et al. (1994). We show in this paper that least-squares fitting procedures that extract diffusion tensor estimates from registered data can be affected by the changes in image noise properties due to interpolation. We also provide a simple method for obtaining correct variance estimates for the registered images.

Theory

In practice, the process of registering two images is usually approached within an optimization framework in which the goal is to find a spatial transformation $f(x)$, where $f : \Re^2 \to \Re^2$, or $f : \Re^3 \to$ $\Re^3$ for volumetric images, that maximizes some similarity measure I between the digitized target $T(x)$ and source $S(x)$ images:

$$\max_f I(S(f(x)), T(x)). \tag{1}$$

The function $f(x)$ may be a rigid body, affine, or higher-order transformation, depending on the application. The function I usually measures the similarity between the images being registered by computing some form of statistical dependency between the intensity values of the images. In the processing pipeline described above, the problem defined by Eq. (1) is usually solved for K images in the image sequence $\{S_1(x),\ldots, S_K(x)\}$, so it is clear that the sequence of images $\{S_1(f_1(x)), S_k(f_k(x)),\ldots, S_K(f_K(x))\}$ is properly aligned. Note that in cases where $f_k(x)$ is used to correct for geometric distortions caused by imperfect magnetic field gradients in MRI, for example, the intensity value of the corrected images may also have to be multiplied by a correction factor (Rohde et al., 2004; Studholme et al., 2000):

$$\tilde{S}_k(f_k(x)) = S_k(f_k(x))\det|\mathrm{Jac}(f_k(x))|, \tag{2}$$

where $\det|\mathrm{Jac}(f_k(x))|$ stands for the determinant of the Jacobian matrix of the transformation $f_k(x)$.

Independently of how the solution to Eq. (1) is actually computed for each image in the sequence, many imaging applications require knowing the value of the registered images $\{S_1(f_1(x_i)),\ldots, S_K(f_K(x_i))\}$ for some arbitrary coordinate x_i. Since in general the point $f_k(x)$ will not coincide with a sampling coordinate of image S_k, an interpolation or approximation strategy must be used to produce the image value $S_k(f_k(x))$. Many approximation and interpolation methods can be chosen to perform such tasks (Meijering et al., 2001). Most estimate the value of $S_k(f_k(x))$ based on a linear combination of the intensity values of image S_k around the point $f_k(x)$. Fig. 1 illustrates this process. Note that w refers to grid coordinates of the image S_k. Mathematically, this interpolation or approximation procedure can be expressed as:

$$S_K(f_K(x)) = \sum_{w_i \in \Theta} \alpha_i S_K(w_i), \tag{3}$$

where Θ defines a set of sampling coordinates that surround $f(x)$ (see Fig. 1). The coefficients α_i of the linear combination Eq. (3) as well as the size of Θ are determined solely by the choice of interpolation or approximation kernel. For the linear interpolation method, one of the most popular image interpolation methods, the value of the image S at coordinate $f(x)$ is given by:

$$S(f(x)) = \sum_{i=1}^{2} \sum_{j=1}^{2} \sum_{k=1}^{2} (1 - V_i)(1 - P_j)(1 - Q_k)S(x_i, y_j, z_k), \tag{4}$$

where $V_i = |f(x)_x - x_i|$, $P_j = |f(x)_y - y_j|$, $Q_k = |f(x)_z - z_k|$, and $\{x_i, y_j, z_k\}$ are image grid coordinates for which $|f(x)_x - x_i| < 1$, $|f(x)_y - y_i| < 1$, $|f(x)_z - z_i| < 1$. Thus the coefficients of the linear combination (Eq. (3)) are given by:

$$\alpha_{i,j,k} = (1 - V_i)(1 - P_j)(1 - Q_k). \tag{5}$$

The set Θ, in this case, are the coordinates w_i for which $|f(x) - w_i| \leq 1$ holds. Note that Eqs. (4) and (5) represent the three-dimensional case, while Fig. 1 depicts a 2-dimensional situation. Naturally, when different interpolation or approximation methods are used, different formulas are needed for estimating the variance
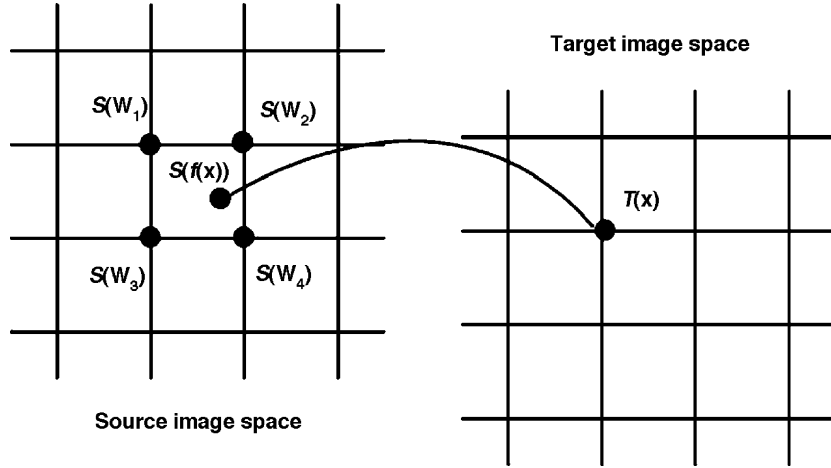
Fig. 1. Illustration of an interpolation or approximation procedure for image registration. First, a coordinate x in the target image space is transferred to a coordinate in the "source" image space via $f(x)$. The value of the source image at $f(x)$ is computed using neighboring values of the source image at that coordinate, $s(w_1)$, $s(w_2)$, etc.

of any given interpolated image value. Refer to Appendix A for the general formula for the variance of an interpolated image value given a general (separable) basis function.

Because of random variability introduced at several steps during image acquisition, the measurement $S_k(w_i)$ should be considered a random variable with a variance $\mathrm{Var}(S_k(w_i))$. For MR images, it is customary to assume that noise variance, denoted by λ^2, is uniform throughout the imaging volume. Note that, though it can be assumed that $S_k(w_i)$ and $S_k(w_j)$, where $i \neq j$, have equal variances, in general they are not independent measurements because several image reconstruction steps effectively correlate measurements from different image coordinates. Correlation in the data due to the reconstruction procedure can arise from filtering during analog to digital conversion, filtering to remove ringing artifacts, filtering to remove noise, correcting for ghosting artifacts (particularly salient in EPI reconstructions), and others. Correlation between values in different image coordinates occurs not only in MRI, but also X-ray-based computed tomography and positron emission tomography (PET). This is because most reconstruction algorithms use filtering operations that correlate intensity values of

Fig. 2. Ordering of voxels used to compute the correlation matrix (E9). Corr(1,3), for example, corresponds to the correlation coefficient between the image value at location with index 1, and the image value at location with index 3.

different image coordinates. A simple method for estimating this correlation in MRI will be described in the next section.

In short, because of the noise variability introduced during image acquisition and processing, the measurements $S_k(w_i)$ and $S_k(w_j)$ are random variables with variance $\mathrm{Var}(S_k(w_i))$ and $\mathrm{Var}(S_k(w_j))$, respectively, and covariance $\mathrm{Cov}(S_k(w_i), S_k(w_j))$. Thus, $S_k(f_k(x))$, as defined by Eq. (3), is also a random variable with variance (Hogg and Craig, 1995):

$$\mathrm{Var}(S_k(f_k(x))) = \left(\sum_{w_i \in \Theta} \alpha_i^2 \mathrm{Var}(S_k(w_i)) \right) + 2 \left(\sum_{\{w_i, w_j\} \in \Theta, i < j} \alpha_i \alpha_j \mathrm{Cov}\big(S_k(w_i), S_k(w_j)\big) \right).$$
(6)

If it can be assumed that $\mathrm{Var}(S_k(w_i)) = \lambda^2$ is approximately constant for all values of the image, Eq. (6) simplifies to,

$$\mathrm{Var}(S_k(f_k(x))) = \lambda^2 \left(\sum_{w_i \in \Theta} \alpha_i^2 \right) + 2 \left(\sum_{\{w_i, w_j\} \in \Theta, i < j} \alpha_i \alpha_j \mathrm{Cov}\big(S_k(w_i), S_k(w_j)\big) \right).$$
(7)

In cases when the intensity correction function defined in Eq. (2) needs to be applied to the registered image $S_k(f_k(x))$ to obtain intensity corrected value $\tilde{S}_k(f_k(x))$, it is easy to show that the correct formula for the variance becomes:

$$\mathrm{Var}\big(\tilde{S}_k(f_k(x))\big) = (\det|\mathrm{Jac}(f_k(x))|)^2 \left(\lambda^2 \left(\sum_{i \in \Theta} \alpha_i^2 \right) + 2 \left(\sum_{\{i,j\} \in \Theta, i < j} \alpha_i \alpha_j \mathrm{Cov}\big(S_k(w_i), S_k(w_j)\big) \right) \right).$$
(8)

Note that if nearest neighbor interpolation is used, the variance of each value in the interpolated image would be equal to the
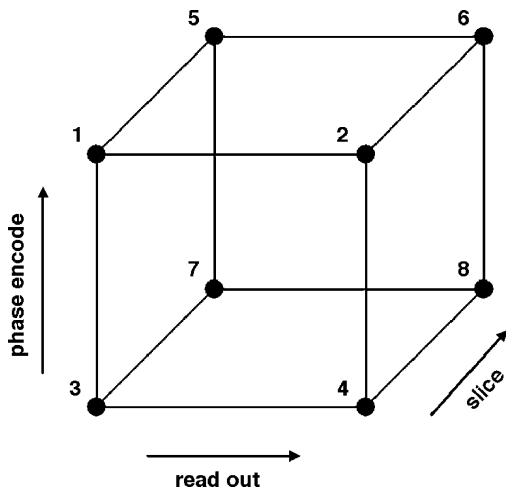
variance of the nearest neighbor voxel, multiplied by the Jacobian correction term when appropriate. We have implemented Eq. (8), using the linear interpolation method, in our DTI processing pipeline and we show next that, because the fitting procedure includes estimates of the noise variance in each image, Eq. (8) should be used to re-estimate the variance of the signal in each voxel in each image that has been registered.

Methods

MRI data acquisition

The data sets in the demonstrations used throughout this paper were acquired with a standard single-shot multi-slice spin-echo EPI sequence (i.e., fat suppression pulse, $90°$ pulse, first diffusion gradient, $180°$ pulse, second diffusion gradient, EPI readout). Scans were performed on a 1.5-T GE Signa system equipped with a whole-body gradient coil able to produce gradient pulses up to 50 mT/m (GE Medical Systems, Milwaukee, WI). The imaged volume was composed of 80 contiguous slices with 2-mm slice thickness and 2-mm in-plane resolution. The echo-time was 82.7 ms, the read-out time 50 ms, and the repetition time was greater than 10 s with cardiac gating (4 acquisitions per heart beat starting with a 150-ms delay after the rise of the sphygmic wave as measured with a peripheral pulse oxymeter). The gradient strength was 49 mT/m, yielding a b value (i.e., trace of the b-matrix) of 1120 s/mm^2. A total of 56 3D images were acquired by repeating 8 times a diffusion sampling scheme described previously (Pierpaoli et al., 1996), which includes one volume with no diffusion weighting followed by the same volume six times, acquired with diffusion gradients applied in different directions. The total imaging time was approximately 20 min. Replicate volumes were acquired for signal to noise considerations in order to improve the quality of the estimated diffusion tensor parameters. The signal to noise ratio, as measured by the mean signal in the region of the thalamus divided by the estimated standard deviation of the signal (see section below), was about 13 for the T2-weighted images and about 7 for the diffusion-weighted images.

MRI noise estimation

The sources that introduce uncertainty in each voxel intensity are many and are generally put into one of two categories: thermal noise and physiological noise. Other sources may also exist in the electronics of the acquisition system, such as digitization, etc., but these can be minimized in an ideal experiment. Thermal noise is usually considered as "white noise" because it is expected that its power should be equal for all frequencies within the readout bandwidth. Because the images are reconstructed using the Fourier transform, the variance that characterizes the uncertainty due to thermal noise is constant throughout the imaging volume (Haacke et al., 1999). Naturally, the same cannot be said about physio-logical noise.

In our experiments, we are only able to estimate the variance that characterizes the uncertainty of the MR measurement due to thermal noise. We do so by computing the variance of magnitude reconstructed intensity values in an artifact-free background region and propagating it to regions with strong signal from the brain through the method described in Gudbjartsson and Patz (1995) and Henkelman (1985). The correction factor described in

Gudbjartsson and Patz (1995) and Henkelman (1985) uses the assumption that Gaussian distributed noise is added to the real and imaginary channels of the receiver system. If possible, we would also like to estimate the variance component due to physiological noise such as flow, MR spin history errors, etc. To do so, however, would require many repeated acquisitions. It would also be difficult to isolate the variance due to patient motion in such repeated measurements (which is something the registration step is actually trying to diminish). Because of these difficulties, we are not able to estimate the variance introduced by physiological effects. Thus the variance estimate we are able to compute for each voxel intensity value is a biased lower bound estimate of the variance when all sources of uncertainty are included. We would like to note, however, that if the total variance, and covariance, in the signal (from all sources) does somehow become available in the future, the same formulas described in the Theory section can be used to propagate the known variance beyond the interpolation step.

The correlation matrix used in our experiments was estimated empirically. Though theoretically possible, it could be very cumbersome to account for all of the filtering steps applied to the data before it becomes a magnitude image. In addition, some steps taken during analog to digital conversion of the free induction decay signals may be proprietary and thus inaccessible. Instead, we acquired and reconstructed several 3D images of pure noise. Using this pure noise image data, we computed the correlation coefficient between the original volumes and the same volumes shifted by one pixel in the x, y, and z directions. Note that because we are using linear interpolation, it is only necessary to include 1 voxel shift in the computation Eq. (8). When bases functions of wider support are used in the interpolation or approximation procedure, the correlations of larger shifts may be required. Using this method, we computed the following 8×8 correlation matrix:

$$
\mathrm{Corr}(i,j) =
\begin{bmatrix}
1 & 0.35 & 0.40 & 0.25 & 0 & 0 & 0 & 0 \\
0.35 & 1 & 0.25 & 0.40 & 0 & 0 & 0 & 0 \\
0.40 & 0.25 & 1 & 0.35 & 0 & 0 & 0 & 0 \\
0.25 & 0.40 & 0.35 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0.35 & 0.40 & 0.25 \\
0 & 0 & 0 & 0 & 0.35 & 1 & 0.25 & 0.40 \\
0 & 0 & 0 & 0 & 0.40 & 0.25 & 1 & 0.35 \\
0 & 0 & 0 & 0 & 0.25 & 0.40 & 0.35 & 1
\end{bmatrix}.
\tag{9}
$$

Fig. 2–which defines the ordering of the voxel coordinates– helps explain the correlation matrix expressed in Eq. (9). Because we are assuming that most of the correlation is caused by linear filtering operations applied on the image data, the noise correlation matrix (Eq. (9)) should be approximately constant throughout the domain of the original magnitude reconstructed images. Note that since our acquisition is based on a 2D EPI pulse sequence, measurements between one slice and the next show no significant correlation. Also note that the correlations in the x and y directions are not equal, since additional operations are performed in the phase encode (y in this case) direction to minimize ghosting artifacts. Lastly, since we are also assuming that the noise variance in the original magnitude reconstructed image is constant, the covariance matrix used in Eq. (8) is given by:

$$
\mathrm{Cov}(i,j) = \mathrm{Corr}(i,j) \times \lambda^2.
\tag{10}
$$

Diffusion tensor estimation

The diffusion tensor model was estimated in each voxel x from the diffusion-weighted data by minimizing the following equation:

$$\chi^2(D(x),A(x)) = \frac{1}{K-7}$$
$$\times \sum_{k=1}^{K} \frac{\left(A(x)e^{(-D(x):b_k)} - \tilde{S}_k(f_k(x))\right)^2}{\mathrm{Var}\left(\tilde{S}_k(f_k(x))\right)}, \qquad (11)$$

where $D(x)$ is a 3×3 symmetric matrix, $A(x)$ is the amplitude term, and b_k is the b-matrix for image k, and $D{:}b$ stands for the matrix dot product (Basser et al., 1994). The minimization was performed using the Levenberg–Marquardt least-squares method.

Simulated data experiments

As an initial test of our variance estimation software, we performed simulation experiments using artificially constructed data. In this experiment, one thousand 2D images of Gaussian distributed random noise with mean zero and variance one were rotated about their centers by $5°$ using bilinear interpolation. In this simulation, the correlation matrix used was approximately:

$$\mathrm{Corr}(i,j) = \begin{bmatrix} 1 & 0.25 & 0.25 & 0.23 \\ 0.25 & 1 & 0.23 & 0.25 \\ 0.25 & 0.23 & 1 & 0.25 \\ 0.23 & 0.25 & 0.25 & 1 \end{bmatrix}.$$

For a fixed pixel coordinate x, the variance across all of the rotated images was computed and displayed. The purpose of this experiment is to show that the variance in the images acquires a particular striped structure. The origin of the striped structure shown stems from the fact that each intensity value in the rotated image was computed by interpolating the original image on a particular non-grid-point coordinate. This estimate comes from a linear combination of the intensity values from around the transformed sampling coordinate (see Fig. 1). The coefficients of the linear combination are computed from the distance of the transformed coordinates to its nearest neighbors. For a specific degree of rotation, this distance will repeat itself every so often throughout the rotated image domain. Since the variance of the rotated image is determined by the coefficients of the linear combination, the variance value of the rotated image will also repeat itself every so often throughout the rotated image domain. We show that by using Eq. (8) the variance in the interpolated images can be predicted exactly.

Experimental data

The diffusion-weighted data used in the examples in this paper were registered to account for patient motion and eddy-current-induced geometric distortions using the methodology described in Rohde et al. (2004). When using this method, first a non-diffusion-weighted image is chosen from the same DWI dataset to be the reference image to which all remaining images are aligned. The registration of each image is done in series and independently from the registration of the other images in the same set. This approach uses a mutual information-based registration technique and a

spatial transformation model containing parameters that correct for eddy-current-induced image distortion and rigid body motion in three dimensions. Each registration consists of estimating 14 parameters in total: 6 for rigid body motion and 8 for the model of eddy-current-induced distortions which consist of a spherical harmonics series expansion in Cartesian coordinates, up to quadratic terms. All 14 parameters for each image in the set are estimated simultaneously. Optimization is performed using a gradient-ascent-type technique within a multi-resolution framework. Initial estimates of the registration parameters are obtained using low-resolution approximations of the images. These estimates are then used to initialize the optimization using higher-resolution representations of the data. The images can also be registered to an arbitrary template with a single interpolation step without additional significant computational cost, though this feature was turned off in all of the experiments shown here. The registered images are created using trilinear interpolation. Following registration, the signal amplitude of each DWI volume is corrected to account for size variations of the object produced by the distortion correction, and the b-matrices are properly recalculated to account for any rotation applied during registration.

The diffusion tensor at each voxel was computed using the registered images by solving Eq. (11) as described above. For comparison purposes, we also estimate the diffusion tensor from the registered images using Eq. (11), but using a constant term for the noise variance $\mathrm{Var}(\tilde{S}_k(f_k(x))) = \lambda^2$. The χ^2 measure at each voxel is compared for both methods. In addition to χ^2, we also compare the estimated tensor parameters to investigate whether or not they are significantly affected when the incorrect noise variance is used.

Results

The results of the simulation experiments are shown in Fig. 3. Part a shows a sample noisy image computed as described above. Part b shows the same image rotated by $5°$ about its center. Values outside the original image were assumed to be zero. Part c shows an image of the variance of the one thousand rotated images computed at each pixel. Clearly the variance became non-uniform and acquired a striped pattern throughout the domain of the image. This variance image was computed analytically using Eq. (7), and the result is shown in part d.

A similar effect can be seen in real data experiments using diffusion-weighted images. Though these striped artifacts are practically invisible in the interpolated DWI volumes, they become evident in the χ^2 maps computed using Eq. (11). Some results are shown in Fig. 4. In this experiment, a set of DWI volumes was rotated about its horizontal axis by about $7.5°$, thus causing interpolation to be performed between values of different slices, as well as between values of different lines in the logical y direction. For this experiment, the same rotation transformation was applied to each DWI volume, that is: $f_1(x) = f_k(x) \ldots = f_K(x)$. Part a of Fig. 4 shows the χ^2 map computed using a single value, λ^2, for the variance of each voxel in each image. Horizontal stripes are visible along the vertical axis of the image, reflecting the different amounts of interpolation performed at each voxel location. Part b shows the variance predicted using Eq. (8). Part c shows the χ^2 map computed using the variance given by Eq. (8). The "striping" patterns become negligible when compared to those shown in part a of the same figure. Note that the dynamic ranges of

Fig. 3. Simulation showing how the interpolation necessary to relate measurements in two images can significantly affect the noise properties of the interpolated image. Part a, an image of simulated noise. Part b is the image in part a rotated by 5°. Part c is the variance of image (b) computed by repeating the rotation experiment 1000 times. Part d shows the variance of image (b) predicted by Eq. (8).

both χ^2 maps in this example were auto-scaled to obtain maximum contrast.

Fig. 5 displays an additional comparison of χ^2 maps computed on unregistered and registered data, with and without the estimation of intensity variance method we propose above. Unlike the example above, the dynamic range of all χ^2 maps was set to [0,5]. Note that the original image data used in this experiment was significantly misregistered due to relatively large subject motion. Part a shows the χ^2 computed from the original, unregistered images using a single variance value estimated from the background of the images. Fig. 5b shows the χ^2 map computed from the registered DW images with the same variance value used in part a. Note that the χ^2 values of the registered images are generally lower than the chi-squared values of unregistered images. Finally, Fig. 5c shows the χ^2 maps computed from the registered DW images using the variance values produced by Eq. (8). Note also that the chi-squared values for part c are generally higher than those of part b.

We also compared some of the most well-known parameters derived from the diffusion tensor computed from the fitting of Eq. (11). For reference, the amplitude, trace, and fractional anisotropy index (Basser, 1995; Basser and Pierpaoli, 1996) are shown in Fig. 6, parts a through c, respectively. Fig. 7a shows the relative error between the trace parameter computed with and without the variance correction scheme proposed above. The relative error was computed using the following formula: $|v_{\text{corrected}} - v_{\text{uncorrected}}|/v_{\text{corrected}}$, where '$v$' stands for the voxel's specific value for the trace of the diffusion tensor. The absolute value of the difference between the fractional anisotropy values computed with and without the variance correction described above is shown in part b of Fig. 7.

Discussion

The rotation experiments performed with the simulated noisy images demonstrate qualitatively and quantitatively the effect that image interpolation can have on the noise variance in registered or interpolated images—the variance becomes non-uniform. The experiment also shows that Eq. (8) can be used to estimate the variance in the interpolated images.

Experiments using real DWIs showed that the change of image noise properties caused by the registration (interpolation) procedure can significantly affect parameter estimation procedure in DT-MRI. First, the alignment of the entire DWI dataset to a standard template can cause χ^2 maps to acquire a striped pattern if a single value for the image intensity variance is used during tensor estimation. The pattern can be explained by the non-uniform intensity variance introduced by the image interpolation step. The patterns disappear when the correct noise variance in each voxel of each image, given by Eq. (8), is used to compute the diffusion tensor. The striped pattern in the χ^2 values is negligible if the DWI dataset was not aligned to a standard template, in addition to being corrected for motion and distortion, even if a single value for the intensity variance is used in estimating the tensor model. Nonetheless, Eq. (8) should be used in this case–because the images have suffered interpolation–to ensure an estimation of the correct variance values. Our results showed that in general the χ^2 computed from registered images is lower than the χ^2 computed from unregistered images when significant misregistration due to motion was present. However, the χ^2 values computed using a single variance value estimated from the original (unregistered) images were lower than the χ^2 values computed


Fig. 4. Demonstration of bias in χ^2 between the DT model and registered DWI data. Part a shows the χ^2 map computed using a single value for the variance in the data. Part b shows the non-uniform variance estimated using Eq. (8). Part c shows the same χ^2 map, however, this time computed using the variance values displayed in part b.

Fig. 5. Part a, χ^2 images computed from the raw (unregistered) data. χ^2 images computed from registered images with (part c) and without (part b) the noise variance formula given in Eq. (8). The χ^2 values computed using the "correct" noise variance values are generally higher than the values computed using a single noise variance estimated on the original (unregistered) images. The actual noise variance in registered images is generally lower than the original noise variance because of the linear combinations performed during image interpolation.

using Eq. (8) to estimate the correct intensity variance. This is to be expected since the variance of registered images at any given voxel location is less than or equal to the variance of the original (unregistered) images because of the interpolations necessary for registration. Thus, if a single variance value estimated from original (unregistered) images is used for the tensor computation, the overall effect will be an artificial decrease in the χ^2 maps derived from tensor fitting.

We have also shown that the estimation of the trace and fractional anisotropy parameters of the diffusion tensor can be affected by incorrect noise variance estimates. In the experiment shown, the error between the parameters estimated with and without the variance correction to account for image registration was small: a few relative percentage points for the trace of the diffusion tensor and a few absolute percentage points for the fractional anisotropy index. We expect that the error caused by inappropriate weights in computing the actual parameters of the diffusion tensor model will be largest when the data being fit differ substantially from the model being used. To understand this, one only has to think of the extreme case in which the model fits the data without error. In this case, the weights being used become irrelevant since the numerator of the chi-squared equation becomes zero. The error between the data and the model arises from normally distributed thermal noise, physiological noise, as well as regions where it is known that the DT model poorly describes the underlying diffusion process, e.g., regions of crossing fibers. When considering only thermal, normally distributed additive noise, as we do throughout this paper, errors caused by incorrect variance estimates are not expected to be large and may diminish as the number of diffusion-weighted images increases. As shown in the Results section, however, these errors are expected to be in the order of a few percent.

The precise effect that changed image noise properties due to interpolation or approximation will have on DT estimation procedures cannot be determined a priori and will depend on several aspects of the registration and data processing procedures. Some of these are the spatial transformations used to register the images, the interpolation or approximation kernel used, the noise variance and covariance in the original images, and the anatomical content of the images. However, it is worth noting that a translation of 0.5 pixels in all three dimensions can cause the variance of the signal to be reduced to 0.125 of the original variance of the signal when the linear interpolation method is used and if the data are spatially uncorrelated. If the correct noise variance value is not used, the resultant χ^2 measure will be underestimated by 8 times. Using the correlation matrix stated in Eq. (9), a translation of 0.5 pixels in all three dimensions would cause the variance in the interpolated image to be 0.25 of the variance in the original data. This would cause the χ^2 measure to be underestimated by 4 times if all images in the dataset suffered similar interpolation.


Fig. 6. Tensor-derived quantities computed after registration with 'correct' variance estimates. Part a, amplitude image. Part b, trace of the diffusion tensor. Part c, fractional anisotropy image.

Fig. 7. Part a, relative error (absolute value of the difference divided by the 'correct' value) between the trace of the diffusion tensor computed with and without the variance estimate given by Eq. (8). Part b, absolute value of the difference between fractional anisotropy values computed with and without the variance estimated by Eq. (8).

Note that, although noise variance, and thus covariance, may vary even between datasets acquired using the same magnet and reconstructed using the same procedure, because of receive coil temperature or amplification settings, for example, the noise correlation should not vary greatly. This is because the most significant correlations are introduced almost entirely by data-independent post-processing operations performed during magnitude image reconstruction. Thus, we expect that the method we propose to estimate the noise correlations in the images to be well suited when the MR images are reconstructed using the same procedure.

By inspecting images a and c displayed in Fig. 4 closely, the reader may notice a slight vertical dark band running through the center of the images. We believe that this is due to non-uniformity in noise variance through the field of view caused by noise aliasing in the frequency encode direction (logical x direction) during Fourier transform-based image reconstruction. The magnet receiver chain includes an analog filter, A/D converter, and a digital decimation filter. The filters reduce the response to higher frequencies. If the filters are not properly chosen, high frequency noise will be aliased into the Nyquist band. The observed pattern reflects the shape of the filter. There is no modulation of the brightness of an object in the field because the object fits into the FOV, so no aliasing takes place. There is no modulation of the noise in the phase encode direction due to its low bandwidth, ~100× lower than in the frequency-encode direction. We are in the process of determining the exact causes and remedies for the problems outlined above, but we do not believe that the slight (though noticeable) pattern in the noise materially affects our results.

Implications for analysis of variance of DT parameters

Knowledge of the uncertainty in the estimated diffusion tensor model parameters is important for assessing the significance of results of inter-subject or inter-acquisition comparisons. It is also worth noting that thermal noise variance not only plays a role in estimating the parameters of the model but also their uncertainty. From Basser et al. (1994), it is known that when multivariate log-

linear regression is used to compute the diffusion tensor parameters, the error variances of the estimated diffusion parameters are given by the diagonal elements of the matrix $(\mathbf{B}^T \Sigma_e^{-1} \mathbf{B})^{-1}$ (see Appendix B), where $\mathbf{B}$ is the "design" matrix for the experiment, computed from the vectors that define the diffusion weighting gradients being used, and the diagonal values of Σ_e^{-1} given by $\tilde{S}_k^2/\mathrm{Var}(\tilde{S}_k)$, where $\tilde{S}_k$ represents the intensity value of the kth image (for a fixed spatial coordinate) in the experiment. As shown in Appendix B, if incorrect values of $\mathrm{Var}(\tilde{S}_i)$ are used, the variance of the estimated parameters is no longer $(\mathbf{B}^T \Sigma_e^{-1} \mathbf{B})^{-1}$ and it is given by Eq. (B3). Methods for estimating the uncertainty in parameters computed through nonlinear models usually rely on Monte Carlo-type simulations for which it is necessary to know the variance that characterizes the uncertainty of each image intensity value (Behrens et al., 2003).

Implications for DT-MRI-based tractography

One application which may be particularly affected by incorrect estimate intensity variance due to random thermal noise is DT-MRI-based tractography. Intensity variations due to thermal noise cause uncertainty in the orientation of greatest diffusivity measured in a DTI experiment. This uncertainty is normally computed using bootstrap (Jones, 2003) or Monte Carlo methods (Behrens et al., 2003). Such approaches are general in the sense that they can be used with both linear and nonlinear regression methods. On the down side, they are computationally intensive. In addition, bootstrap methods such as the one discussed in Jones (2003) require the acquisition of an additional amount of data. In both cases, testing the effect of different experimental setups (diffusion-weighted directions, diffusion weighting strength, number of image replicates, etc.) can be cumbersome.

Alternatively, given a specific set of b-matrices and a diffusion tensor, a root mean square estimate of the uncertainty in orientation as a function of thermal noise variance can be derived using the theory of linear regression (see Appendix B). This result can be used to calculate the approximate effect that incorrect intensity variance estimates can have on the variability of the principal diffusivity direction. We use a set of 22 b-matrices derived using the scheme described in Jones et al. (1999) and an anisotropic diffusion tensor specified by the eigenvectors $g_1 = \{1,0,0\}$, $g_2 = \{0,1,0\}$, $g_3 = \{0,0,1\}$, and eigenvalues $d_1 = 1,685\mathrm{e}^{-6}$, $d_2 = 287\mathrm{e}^{-6}$, $d_3 = 109\mathrm{e}^{-6}$ mm^2/s, and SNR = 15 to demonstrate the following example. When correct variance values are used in the estimation process, the covariance matrix of the estimated DT parameters is given by Eq. (B4) and the root mean square estimate in angle deviation when using the correct variance values is about 2.5°. If the entire set of images is translated by 0.5 pixels in all three dimensions, using linear interpolation and using the covariance matrix stated in Eq. (9), for example, and the variance of each intensity value is not recomputed using the method described above, the covariance matrix of the estimated parameters is given by Eq. (B3). The root mean square estimate of angle deviation in this case increases to about 10°. This result seems counter intuitive since data interpolation should reduce the intensity variance of image values. This, in turn, should reduce the variability of the measurement of principal direction. We point out, however, that this is only caused by neglecting to account for the variance reduction due to image interpolation. If the variances of the image intensity values are appropriately recalculated, the root

mean square of the angle variation is reduced to about $1.3°$. We point out that these results are only approximations since they were obtained using first order expansion methods. Moreover, only variability due to thermal noise was included. However, it seems clear that tractography methods that rely on information about the variability of diffusivity orientation should be directly and adversely affected by neglecting to recompute intensity variance estimates after image registration.

Moreover, we point out that probabilistic tractography is not the only approach that could be affected by intensity variance modifications due to interpolation. The deterministic methods presented in Basser et al. (2000) and Conturo et al. (1999) rely on estimating a continuous version of the diffusion tensor field for numerically computing the continuous path of presumed fiber tracts. In the approach described by Basser et al. (2000), the continuous tensor field is estimated using an approximate fit to the discretely sampled diffusion tensor data using cubic B-splines. Conturo et al. (1999) obtain a continuous version of the diffusion tensor field by interpolating the diffusion-weighted images where needed and fitting the DT model using the interpolated values. In both cases, the continuous approximation of diffusion tensor data produces diffusion tensors with different variance properties at different locations in the domain of the images. The variance of the interpolated diffusion tensors can be computed using the formulas given in Appendix A. Thus, such deterministic tract following approaches effectively integrate tracts by using estimated principal diffusivity directions that have different orientation uncertainty across different parts of the images, whether or not the raw DW images used for computing each diffusion tensor have been registered. At this point, it is unclear what are the effects of non-uniform variance for such deterministic tract following methods. However, we believe that further investigation in the area is merited.

Implications for functional MRI and voxel based morphometry

Note that though we used diffusion tensor imaging as a case study, we believe that the same methodology could be used whenever data analysis requiring noise variance estimates is performed on registered or interpolated data. Some application examples in biomedical imaging include fMRI data analysis, studies of tissue shape and composition using statistical analysis of image data, MR relaxometry experiments, etc. In all such applications, the goal is to detect image intensity changes that are the result of some biologically relevant phenomena. In fMRI, this may be BOLD activation correlated with some type of brain activity, while in voxel-based morphometry, for example, this may be information related to diseased tissue. Both fMRI data analysis and voxel-based morphometry methods often rely on a generalized linear model for identifying the presence, absence, and quantification of biologically relevant phenomena. In this framework, the measured image data (at a fixed voxel coordinate), defined by an N dimensional vector y, is modeled as a linear combination of explanatory coefficients arranged in an $N \times M$ matrix $\mathbf{M}$ and unknown parameters defined by an M dimensional vector a: $y = \mathbf{M}a + \mathbf{e}$, where $\mathbf{e}$ represents an N dimensional error vector whose entries are usually assumed to be independent, equally, and normally distributed. If the error values are indeed normally distributed, the maximum likelihood estimate for the model parameters is

given by $a = (\mathbf{M}^T\mathbf{M})^{-1}\mathbf{M}^Ty$, while the covariance matrix of the estimates is given by $\mathbf{S}_a = \mathbf{L}\ \mathbf{S}_y\mathbf{L}^T$, with $\mathbf{L} = (\mathbf{M}^T\mathbf{M})^{-1}\mathbf{M}^T$ and $\mathbf{S}_a$, $\mathbf{S}_y$ representing the covariance matrix of the estimated parameters and original data, respectively. Since the measurements y are usually assumed to be independently and identically distributed, the covariance matrix of the estimated parameters reduces to $\mathbf{S}_a = \lambda^2\ (\mathbf{M}^T\mathbf{M})^{-1}$, with λ^2 being the assumed noise variance. Note that this analysis is usually performed on registered images in order to account for patient motion and geometric distortions. As shown in this paper, since different images will have different spatial transformations (and thus different interpolation) applied on them, the constant noise variance assumption is no longer appropriate. That is, the variance due to noise of an image value that has suffered interpolation is expected to be different from the variance of an image value that has suffered no interpolation at all. At this point it is unclear what effect this will have on image analysis results obtained using the general linear model, though it is an issue that should be investigated further.

Summary and conclusions

As fitting and estimating procedures from registered image data become increasingly more elaborate and quantitative, knowledge of the intensity variance due to noise will become more important for increasing the accuracy and scientific value of the results obtained from them. A method for estimating the variance in registered images is presented. The general approach can be summarized as follows. The output of the registration procedure is computed using an image interpolation or approximation procedure. The interpolation or approximation procedure can be written as a linear combination of the values of the image being registered. The coefficients of the linear combination are determined by the choice of interpolation or approximation kernel. Since the values of the image being registered are typically corrupted by noise, this operation can be viewed as a linear combination of random variables. The variance of the linear combination is given by well-known statistical formulas.

The image interpolation or approximation generally required by image registration procedures will inevitably affect the noise variance properties of the images. We have shown that incorrect variance estimates can have a significant effect on diffusion tensor estimation procedures. The method we proposed for estimating the noise variance in registered images was shown to be successful in both simulated and real data experiments. Since χ^2 measures and noise variance estimates are used more and more frequently in diffusion data analysis–examples include image registration (Andersson and Skare, 2002), diffusion model selection (Alexander et al., 2002; Shrager et al., 2002), robust tensor estimation (Chang et al., 2004), and brain tumor pathology detection (Maier et al., 2003)–correct variance estimates from registered image data will become increasingly important.

The methods described here could also be useful in other biomedical imaging applications such as MR relaxometry, fMRI data analysis, voxel-based morphometry, etc. However, the effects of the technique in each of these applications are not discussed in detail here and could be the subject of future study. The techniques described here could also find applications in other image processing and data analysis fields such as automatic target

recognition and segmentation of registered data obtained from satellite or other remote sensing machinery. Statistical approaches are often used to fuse information gathered from several sensors and extract possible target matches.

Appendix A

We expect that different interpolation or approximation kernels will modify the variance in the registered images differently. The precise manner in which the choice of interpolator will affect the variance of an image is currently being investigated (Rohde, 2004). Here we give a general formula for the variance of the image intensity value produced using any kernel-based interpolation method due to a spatial transformation being applied during registration. Let $s(k)$, where $k \in Z^d$, with d being the dimension of the images, be the discretely sampled image produced by the acquisition system. A continuous approximation to $s(k)$ is given by:

$$\tilde{s}(x) = \sum_{k \in Z^d} s(k)h(x - k) \tag{A1}$$

Note that the summations are carried from $-\infty$ to ∞ by making the images periodic. Naturally, if we would like the values $\tilde{s}(x)$ to be equal to the values of $s(k)$ at coordinates $x = k$, then $h(x)$ must have the following properties:

$$h(k) = 0 \forall k \neq 0, \tag{A2}$$

and

$$h(0) = 1. \tag{A3}$$

Examples of such kernels are the linear "hat" function (also known as B-spline of order 1), and the popular sinc kernel given by:

$$h(x) = \begin{cases} 1 - |x|, |x| \leq 1 \\ 0, |x| > 1 \end{cases}, \tag{A4}$$

and

$$h(x) = \frac{\sin(\pi x)}{\pi x}, \tag{A5}$$

respectively. Note that "true" sinc interpolation is almost never used in the field of medical imaging because of the enormous computational cost associated with it. Since the support of sinc is infinite, in theory, the sum in Eq. (A1) should be evaluated from $-\infty$ to ∞. Because of such computational costs and other reasons (i.e., "ringing" artifacts), researchers in the field prefer to use truncated and apodized versions of Eq. (A5) (Hajnal et al., 1995; Lehmann et al., 1999, 2001; Meijering et al., 2001; Thevenaz et al., 2000). Note also that in the cases where $d > 1$, the interpolation kernel is replaced by

$$\hat{h}(x) = \prod_{i=1}^{d} h(x_i). \tag{A6}$$

If the basis function being used does not satisfy the properties stated in Eqs. (A2) and (A3), examples include the

popular B-splines of order 2 or greater, Eq. (A1) needs to be adjusted. Let $b(x)$ be a basis function such that properties of Eqs. (A2) and (A3) do not hold. The interpolation equation then becomes:

$$\tilde{s}(x) = \sum_{k \in Z^d} c(k)b(x - k). \tag{A7}$$

The coefficients $c(k)$ are given by:

$$c(k) = \left(b^{-1} * s\right)(k), \tag{A8}$$

where b^{-1} is the uniquely defined convolution-inverse (Unser et al., 1993a,b). As shown in Unser et al. (1993a,b), we can substitute Eq. (A8) into Eq. (A7) to see that

$$\tilde{s}(x) = \sum_{k \in Z^d} \left(b^{-1} * s\right)(k)b(x - k)$$

$$= \sum_{k_1 \in Z^d} \sum_{k_2 \in Z^d} b^{-1}(k_2)s(k_1 - k_2)b(x - k_1)$$

$$= \sum_{k \in Z^d} s(k)h(x - k) \tag{A9}$$

where the new interpolation kernel is given by:

$$h(x) = \sum_{k \in Z^d} b^{-1}(k)b(x - k). \tag{A10}$$

Thus, the variance of the interpolated image intensity value due to spatial transformation $f(x)$ is given by:

$$\text{Var}(\tilde{s}(f(x))) = \sum_{k \in Z^d} \text{Var}(s(k))(h(f(x) - k))^2 + \sum_{i \in Z^d}$$
$$\times \sum_{j \in Z^d, j \neq 1} h(f(x) - i)h(f(x) - j)\text{Cov}(s(i), s(j)). \tag{A11}$$

Applications such as geometrical distortion correction due to imperfect magnetic fields in MRI require the formula above to be multiplied by the square of the determinant of the Jacobian matrix of f, as in Eq. (8):

$$\text{Var}(\tilde{s}(f(x))) = (\det|\text{Jac}(f(x))|)^2$$
$$\times \left(\sum_{k \in Z^d} \text{Var}(s(k))(h(f(x) - k))^2 + \sum_{i \in Z^d} \right.$$
$$\left. \times \sum_{j \in Z^d, j \neq 1} h(f(x) - i)h(f(x) - j)\text{Cov}(s(i), s(j)) \right). \tag{A12}$$

Appendix B

Using the log-linear diffusion tensor model for the diffusion-weighted image data, we analyze the error distribution of the estimated diffusion tensor parameters. Let $y = \{\ln(S_1), \ldots, \ln(S_N)\}^T$, where S_i represents the ith measurement in a typical DTI acquisition, and $a = \{D_{xx}, D_{yy}, D_{zz}, D_{xy}, D_{xz}, D_{yz}, \ln(A_0)\}^T$

represents the diffusion tensor model parameters. To first order, the log linear model can be written as

$$y = \mathbf{B}a + e \tag{B1}$$

where the jth row of $\mathbf{B}$ is composed of the b-matrix entries of the jth diffusion-weighted acquisition parameters $\{-b_{xxj}, -b_{yyj}, -b_{zzj}, -2b_{xyj}, -2b_{xzj}, -2b_{yzj}, 1\}$, and e represents the error vector. The covariance matrix of e is denoted $(\Sigma_e)_{ij} = \sigma_i^2/\langle S_i \rangle^2$, where $\langle u \rangle$ denotes the expectation of random variable u. Since each measured data point in y was taken independently at different times $(\Sigma_e)_{ij} = 0$ $\forall i \neq j$. All terms in Eq. (B1) are considered deterministic, with exception of e which represents error due to noise in the imaging acquisition system. Therefore $\Sigma_y = \Sigma_e$. In practice however, one can only estimate Σ_e. This is usually done based on measurements from background intensity values. As shown in this paper, the estimates of the variance in each image intensity value need to be recalculated after registration or interpolation is performed on the images. We will differentiate the true covariance matrix of the data Σ_e from the estimated one $\tilde{\Sigma}_e$. The weighted least squares solution to Eq. (B1) is given by:

$$a = \left(\mathbf{B}^T \tilde{\Sigma}_e^{-1} \mathbf{B} \right)^{-1} \left(\mathbf{B}^T \tilde{\Sigma}_e^{-1} \right) y, \tag{B2}$$

while the covariance matrix of the estimated parameters is given by:

$$\Sigma_a = \langle aa^T \rangle = (\mathbf{B}^T \tilde{\Sigma}_e^{-1} \mathbf{B})^{-1} \mathbf{B}^T \tilde{\Sigma}_e^{-1} \Sigma_e \tilde{\Sigma}_e^{-1} \mathbf{B} (\mathbf{B}^T \tilde{\Sigma}_e^{-1} \mathbf{B})^{-1}. \tag{B3}$$

If our estimate of the covariance matrix of the measured data is precise and accurate, $\tilde{\Sigma}_e \approx \Sigma_e$, then Eq. (B3) reduces to:

$$\Sigma_a = (\mathbf{B}^T \tilde{\Sigma}_e^{-1} \mathbf{B})^{-1}. \tag{B4}$$

If, on the other hand, errors are made in calculating $\tilde{\Sigma}_e$, such as neglecting to account for the interpolation applied to the data during image registration, the covariance of the estimated parameters is given by Eq. (B3).

The uncertainty in the principal direction orientation in a diffusion tensor $\tilde{D}$ calculated using Eq. (B2) can be estimated by studying the effects of random perturbations ΔD on a deterministic tensor D_0 (Basser, 1997):

$$\tilde{D} = D_0 + \Delta D. \tag{B5}$$

Let d_1, d_2, d_3 and g_1, g_2, g_3 represent the eigenvalues (arranged in decreasing order) and eigenvectors, respectively, of the three-dimensional positive definite symmetric tensor D_0. We are interested in computing the perturbation $\tilde{g}_1 = g_1 + \Delta g_1$. We will assume that D_0 comes from biological tissue with high diffusion anisotropy so that sorting bias in the computed eigenvalues can be safely neglected. It can be shown (Fukunaga, 1972) that, to first order, the perturbation of the eigenvector associated with greatest diffusivity is:

$$\Delta g_1 = \sum_{i=2}^{3} \left(\frac{g_1^T \Delta D g_i}{d_1 - d_i} \right) g_i. \tag{B6}$$

The perturbation angle θ between $\tilde{g}_1$ and g_1 is thus (Basser, 1997) $\theta = \tan^{-1}(||\Delta g_1||)$. Noting that the eigenvectors g form an orthonormal basis for the 3D Euclidean space and using the small angle approximation for $\tan \theta$:

$$\theta \approx ||\Delta g_1|| = \sqrt{\sum_{i=2}^{3} \left(\frac{g_1^T \Delta D g_i}{d_1 - d_i} \right)^2}. \tag{B7}$$

Now

$$\langle \theta^2 \rangle \approx \langle ||\Delta g_1||^2 \rangle = \sum_{i=2}^{3} \left\langle \left(\frac{g_1^T \Delta D g_i}{d_1 - d_i} \right)^2 \right\rangle = \sum_{i=2}^{3} \frac{\langle (g_1^T \Delta D g_i)^2 \rangle}{(d_1 - d_i)^2} \tag{B8}$$

while, as shown in Anderson (2001),

$$\langle (\Delta D_{ii})^2 \rangle = \langle (\tilde{D}_{ii} - (D_0)_{ii})^2 \rangle = (\Sigma_a)_{i,i} \tag{B9}$$

and similarly

$$\langle (\Delta D_{ij})^2 \rangle = \langle (\tilde{D}_{ij} - (D_0)_{ij})^2 \rangle = (\Sigma_a)_{i+j+1,i+j+1}. \tag{B10}$$

Looking at the principal axis case, $g_i^T \Delta D g_j = \Delta D_{ij}$. In general, the covariance matrix can be rotated so that $\langle (g_i^T \Delta D g_j)^2 \rangle = (\mathbf{R}\Sigma_a \mathbf{R}^T)_{i+j+1, i+j+1} = \Xi_{i+j+1, i+j+1}$ (Anderson, 2001), where $\mathbf{R}$ is a rotation matrix. Therefore, Eq. (8) can be written as

$$\langle ||\Delta g_1||^2 \rangle = \sum_{i=2}^{3} \frac{\langle (g_1^T \Delta D g_i)^2 \rangle}{(d_1 - d_i)^2} = \sum_{i=2}^{3} \frac{(\Xi)_{2+i,2+i}}{(d_1 - d_i)^2}. \tag{B11}$$

Thus the root mean square angle estimate of the deviation from the principal direction is:

$$\theta_{RMS} = \sqrt{\sum_{i=2}^{3} \frac{(\Xi)_{2+i,2+i}}{(d_1 - d_i)^2}}, \tag{B12}$$

where the covariance matrix of the estimated parameters Σ_a in $\mathbf{R}\Sigma_a \mathbf{R}^T = \Xi$ is given in Eq. (B3).

References

Alexander, D.C., Barker, G.J., Arridge, S.R., 2002. Detection and modeling of non-Gaussian apparent diffusion coefficient profiles in human brain data. Magn. Reson. Med. 48 (2), 331–340.

Anderson, A.W., 2001. Theoretical analysis of the effects of noise on diffusion tensor imaging. Magn. Reson. Med. 46 (6), 1174–1188.

Andersson, J.L., Skare, S., 2002. A model-based method for retrospective correction of geometric distortions in diffusion-weighted EPI. Neuro-Image 16 (1), 177–199.

Ashburner, J., Friston, K.J., 2000. Voxel-based morphometry—The methods. NeuroImage 11 (6 Pt. 1), 805–821.

Basser, P.J., 1995. Inferring microstructural features and the physiological state of tissues from diffusion-weighted images. NMR Biomed. 8 (7–8), 333–344.

Basser, P.J., 1997. Quantifying Errors in Fiber-Tract Direction and Diffusion Tensor Field Maps Resulting from MR Noise. International Society for Magnetic Resonance in Medicine (ISMRM), Vancouver, BC, Canada.

Basser, P.J., Pierpaoli, C., 1996. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. J. Magn. Reson., B 111 (3), 209–219.

Basser, P.J., Mattiello, J., LeBihan, D., 1994. Estimation of the effective self-diffusion tensor from the NMR spin echo. J. Magn. Reson. B. 103 (3), 247–254.

Basser, P.J., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A., 2000. In vivo fiber tractography using DT-MRI data. Magn. Reson. Med. 44 (4), 625–632.

Bastin, M.E., 1999. Correction of eddy current-induced artefacts in diffusion tensor imaging using iterative cross-correlation. Magn. Reson. Imaging 17 (7), 1011–1024.

Behrens, T.E., Woolrich, M.W., Jenkinson, M., Johansen-Berg, H., Nunes, R.G., Clare, S., Matthews, P.M., Brady, J.M., Smith, S.M., 2003. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. Magn. Reson. Med. 50 (5), 1077–1088.

Chang, L.-C., Rohde, G.K., Jones, D.K., Basser, P.J., Pierpaoli, C., 2004. RESTORE: Robust Estimation of Tensors by Outlier Rejection. 12th Annual Meeting of ISMRM, Kyoto, Japan. International Society for Magnetic Resonance in Medicine (ISMRM), Wiley.

Conturo, T.E., Lori, N.F., Cull, T.S., Akbudak, E., Snyder, A.Z., Shimony, J.S., McKinstry, R.C., Burton, H., Raichle, M.E., 1999. Tracking neuronal fiber pathways in the living human brain. Proc. Natl. Acad. Sci. U. S. A. 96 (18), 10422–10427.

Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.P., Frith, C.D., Frackowiak, R.S., 1995. Statistical parametric maps in functional imaging: a general linear approach. Hum. Brain Mapp. 2, 189–210.

Friston, K.J., Williams, S., Howard, R., Frackowiak, R.S., Turner, R., 1996. Movement-related effects in fMRI time-series. Magn. Reson. Med. 35 (3), 346–355.

Fukunaga, K., 1972. Introduction to Statistical Pattern Recognition. Academic Press Inc., New York.

Grootoonk, S., Hutton, C., Ashburner, J., Howseman, A.M., Josephs, O., Rees, G., Friston, K.J., Turner, R., 2000. Characterization and correction of interpolation effects in the realignment of fMRI time series. NeuroImage 11 (1), 49–57.

Gudbjartsson, H., Patz, S., 1995. The Rician distribution of noisy MRI data. Magn. Reson. Med. 34 (6), 910–914.

Haacke, E.M., Brown, R.W., Thompson, M.R., Venkatesan, R., 1999. Magnetic Resonance Imaging: Physical Principles and Sequence Design. Wiley-Liss, New York.

Hajnal, J.V., Saeed, N., Soar, E.J., Oatridge, A., Young, I.R., Bydder, G.M., 1995. A registration and interpolation procedure for subvoxel matching of serially acquired MR images. J. Comput. Assist. Tomogr. 19 (2), 289–296.

Haselgrove, J.C., Moore, J.R., 1996. Correction for distortion of echo-planar images used to calculate the apparent diffusion coefficient. Magn. Reson. Med. 36 (6), 960–964.

Henkelman, R.M., 1985. Measurement of signal intensities in the presence of noise in MR images. Med. Phys. 12 (2), 232–233.

Hogg, R.V., Craig, A.T., 1995. Expectations of functions of random variables. In: Pirtle, R.W., ed., Introduction to Mathematical Statistics. Prentice-Hall, Inc., Upper Saddle River, NJ.

Horsfield, M.A., 1999. Mapping eddy current induced fields for the correction of diffusion-weighted echo planar images. Magn. Reson. Imaging 17 (9), 1335–1345.

Jones, D.K., 2003. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. Magn. Reson. Med. 49 (1), 7–12.

Jones, D.K., Horsfield, M.A., Simmons, A., 1999. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. Magn. Reson. Med. 42 (3), 515–525.

Lehmann, T.M., Gonner, C., Spitzer, K., 1999. Survey: interpolation methods in medical image processing. IEEE Trans. Med. Imaging 18 (11), 1049–1075.

Lehmann, T.M., Gonner, C., Spitzer, K., 2001. Addendum: B-spline interpolation in medical image processing. IEEE Trans. Med. Imaging 20 (7), 660–665.

Maas, L.C., Renshaw, P.F., 1999. Post-registration spatial filtering to reduce noise in functional MRI data sets. Magn. Reson. Imaging 17 (9), 1371–1382.

Maier, S.E., Mamata, H., Mulkern, R.V., 2003. Characterization of normal brain and brain tumor pathology by chisquares parameter maps of diffusion-weighted image data. Eur. J. Radiol. 45 (3), 199–207.

Maintz, J.B., Viergever, M.A., 1998. A survey of medical image registration. Med. Image Anal. 2 (1), 1–36.

Mangin, J.F., Poupon, C., Clark, C., Le Bihan, D., Bloch, I., 2002. Distortion correction and robust tensor estimation for MR diffusion imaging. Med. Image Anal. 6 (3), 191–198.

Meijering, E.H., Niessen, W.J., Viergever, M.A., 2001. Quantitative evaluation of convolution-based methods for medical image interpolation. Med. Image Anal. 5 (2), 111–126.

Nickerson, L.D., Narayana, S., Lancaster, J.L., Fox, P.T., Gao, J.H., 2003. Estimation of the local statistical noise in positron emission tomography revisited: practical implementation. NeuroImage 19 (2 Pt. 1), 442–456.

Pierpaoli, C., Jezzard, P., Basser, P.J., Barnett, A., Di Chiro, G., 1996. Diffusion tensor MR imaging of the human brain. Radiology 201 (3), 637–648.

Pluim, J.P., Maintz, J.B., Viergever, M.A., 2000. Interpolation artefacts in mutual information-based image registration. Comput. Vis. Image Underst. 77, 211–232.

Pluim, J.P., Maintz, J.B., Viergever, M.A., 2003. Mutual-information-based registration of medical images: a survey. IEEE Trans. Med. Imaging 22 (8), 986–1004.

Rohde, G.K., 2004. Ph.D. dissertation research. Applied Mathematics and Scientific Computation Program. College Park, University or Maryland.

Rohde, G.K., Barnett, A.S., Basser, P.J., Marenco, S., Pierpaoli, C., 2004. Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI. Magn. Reson. Med. 51 (1), 103–114.

Shrager, R., Jones, D.K., Pajevic, S., Munson, P., Basser, P.J., 2002. When is a Gaussian Displacement Distribution Adequate to Describe Water Diffusion in Tissues? Workshop on Diffusion MRI: Biophysical Issues. ISMRM, Saint-Malo, France.

Sijbers, J., den Dekker, A.J., Van Audekerke, J., Verhoye, M., Van Dyck, D., 1998. Estimation of the noise in magnitude MR images. Magn. Reson. Imaging 16 (1), 87–90.

Studholme, C., Constable, R.T., Duncan, J.S., 2000. Accurate alignment of functional EPI data to anatomical MRI using a physics-based distortion model. IEEE Trans. Med. Imaging 19 (11), 1115–1127.

Thacker, N.A., Jackson, A., Moriarty, D., Vokurka, E., 1999. Improved quality of re-sliced MR images using re-normalized sinc interpolation. J. Magn. Reson. Imaging 10 (4), 582–588.

Thevenaz, P., Blu, T., Unser, M., 2000. Interpolation revisited. IEEE Trans. Med. Imaging 19 (7), 739–758.

Unser, M., Aldroubi, A., Eden, M., 1993. B-Spline signal processing: Part I—Theory. IEEE Trans. Signal Process. 41 (2), 821–833.

Unser, M., Aldroubi, A., Eden, M., 1993. B-Spline signal processing: Part II—Efficient design and applications. IEEE Trans. Signal Process. 41 (2), 834–848.