Distortion correction and robust tensor estimation for MR diffusion imaging

J.-F. Mangin[a,c,*], C. Poupon[a,b,c,d], C. Clark[a,c], D. Le Bihan[a,c], I. Bloch[b,c]

[a]*Service Hospitalier Frédéric Joliot, CEA, 91401 Orsay, France*
[b]*Département Signal et Images, CNRS URA 820, ENST, Paris, France*
[c]*Institut Fédéral de Recherche 49, Paris, France*
[d]*GE Medical Systems, Buc, France*

Abstract

This paper presents a new procedure to estimate the diffusion tensor from a sequence of diffusion-weighted images. The first step of this procedure consists of the correction of the distortions usually induced by eddy-current related to the large diffusion-sensitizing gradients. This correction algorithm relies on the maximization of mutual information to estimate the three parameters of a geometric distortion model inferred from the acquisition principle. The second step of the procedure amounts to replacing the standard least squares-based approach by the Geman–McLure M-estimator, in order to reduce outlier-related artefacts. Several experiments prove that the whole procedure highly improves the quality of the final diffusion maps.
© 2002 Elsevier Science B.V. All rights reserved.

Keywords: Registration; Diffusion; Distortion; Tensor; Robust

1. Introduction

There is currently considerable interest in the use of MRI for imaging the apparent diffusion of water in biological tissues (Le Bihan et al. 2001). The physical process underlying this diffusion is the random walk motion of the molecules in a fluid: due to thermal agitation, the molecules are constantly moving and colliding with neighbors. When the fluid is embedded into the complex geometry of biological tissue, however, the collisions with cell membranes and macromolecules and the restriction to various compartments highly influence this process. Hence, by probing the microscopic motion of tissular water, diffusion-weighted imaging provides a unique in vivo tool for studying the structure of biological tissue. In particular, this imaging modality gives access to various information about the brain microstructures that could be used to improve the interpretation of functional imaging studies.

Diffusion-weighted MRI relies on a phenomenon that could have been considered as a technical problem only. When protons are placed into a static magnetic field ($B0$), they begin to precess (i.e. their magnetic vector rotates around $B0$). A key point underlying MRI is the linear relationship between the local field strength and the frequency of the proton precession motion, which is the frequency of the signal produced in the receiving antenna. Adding a gradient to $B0$ encodes the spin localization into this frequency, which leads to images after Fourier transform. Without special preparation, the protons (spins) precessing in a static magnetic field ($B0$) do not produce signal in the receiving antenna (in x–y plane) because of lack of coherence between the individual precessions (they are all out of phase and hence have no net transverse component). By applying a 90 degree radio-frequency (RF) pulse, the frequency of which matches the frequency of precession of protons, the spins can be made to be in phase and have a net transverse component, producing signal in an antenna. After the 90 degree RF pulse the spins will again go out of phase, mainly because of the effect of external field inhomogeneities. For static spins, the dephasing caused by external field inhomogeneities can be eliminated with a 180 degree pulse leading to what is called a spin echo. This is not possible for spins undergo-

ing diffusion because they are not static (their position fluctuates randomly because of the random character of the thermal spin motion). The result is diffusion-related signal attenuation.

While with standard MR sequences, the diffusion-related signal attenuation is negligible, diffusion imaging sequences increase this effect with the addition of two strong diffusion sensitizing gradient pulses (Stejskal and Tanner, 1965). These additional gradients increase the attenuation of the signal produced by the spins that move along the gradient direction. Within a simple isotropic medium like a glass of water, the attenuation is related to an exponential of the medium property called the diffusion coefficient D (the standard coefficient of Fick's law). By using, for instance, an image without diffusion weighting and one diffusion-weighted image, we can calculate a D value for each voxel. As a consequence of their spatial structure, however, many substances and biological tissues exhibit anisotropic diffusion behavior: the computed diffusion coefficient depends on the direction of the sensitizing gradient. Therefore, when anisotropy of the 3D diffusion process is of interest, for instance for fiber bundle tracking (Poupon et al., 2001), a symmetric diffusion tensor D has to be calculated for each voxel from a series of diffusion-weighted volumes (Basser et al., 1994a,b). Each such volume is acquired with a different applied diffusion-sensitizing gradient (Stejskal and Tanner, 1965). These gradients are applied in order to vary a symmetric matrix b (s/mm^2) that depends on the gradient direction, strength and timing (Mattiello et al., 1994). The diffusion-sensitizing gradient affects the signal intensity of any given voxel in a manner that can be described by the linear equation

$$\ln S(b) = \ln S(0) - D_{xx}b_{xx} - 2D_{xy}b_{xy} - 2D_{xz}b_{xz} - D_{yy}b_{yy}$$
$$- 2D_{yz}b_{yz} - D_{zz}b_{zz}, \qquad (1)$$

where S denotes the signal of the selected voxel. When a sufficient number of different b matrices is used (related to at least six different gradient directions), the diffusion tensor D can be estimated. Such calculations are simple if each voxel in the different volumes represents the same point in the anatomy of the subject, but can be impractical

if different volumes of the series are distorted relative to each other. Diffusion-weighted images, however, are often acquired using echo-planar imaging (EPI), to reduce acquisition time and artefacts related to physiological motions. Unfortunately, this fast acquisition scheme is highly sensitive to eddy currents induced by the large diffusion gradients (Haselgrove and Moore, 1996). These eddy currents can cause significant distortions in the phase-encoding direction where the image bandwidth is quite low (see Fig. 1). Since the degree and nature of this artefact typically vary both with the strength and orientation of the diffusion-sensitizing gradient, distortions can dramatically change the estimated diffusion tensor.

The methods for reducing the effects of eddy currents may be divided into three categories. The first one simply consists of modifications of the gradient sequences (Alexander et al., 1997). This approach, however, seems insufficient to get completely rid of artefacts. A second family of approaches, which rely deeply on MR physics, require additional experimental data (Jezzard et al., 1998; Horsfield, 1999). Since the eddy-current distortions do not rely on the subject's head geometry, these cumbersome additional acquisitions can be done on phantoms only during a calibration operation. Unfortunately, the obtained correction scheme has to be updated on a regular basis because of some slow variations of the magnet (Bastin and Armitage, 2000).

The last kind of approaches are purely retrospective and can be considered as registration methods. They use a distortion geometric model inferred from the acquisition principle, which leads to estimate a few parameters using a standard similarity measure like cross-correlation (Haselgrove and Moore, 1996; Calamante, 1999; Bastin, 1999). Such simple similarity measures, however, are not sufficient to perfectly take into account the complex dependencies embedded in Eq. (1) (Bastin and Armitage, 2000). In this paper we propose to estimate the few parameters of the distortion geometric model using the mutual information as similarity measure in order to achieve a more robust correction scheme (Maes, 1997; Wells III et al., 1996).

It should be understood that EPI imaging leads to

no gradient   13mT/m (-1,1,0)   22mT/m (-1,1,0)
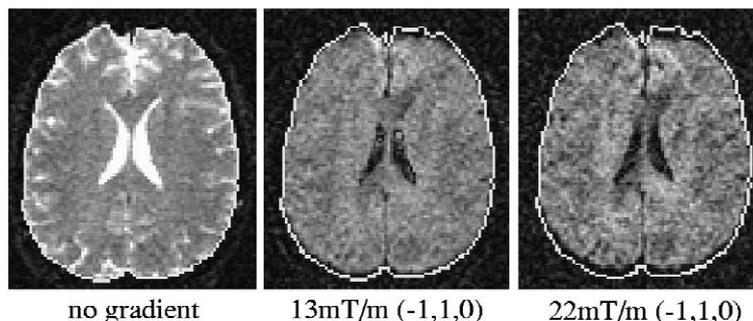
Fig. 1. Example of eddy-current-related distortions (8 mm in the worst case) for two different diffusion sensitizing gradient magnitudes.

another kind of distortion induced by susceptibility artefacts (Jezzard and Clare, 1999) that is not addressed in this paper because it is not dependent on the diffusion gradient. These distortions are non-linear and depend on the subject's head geometry. They have to be corrected to relate functional MRI experiments performed with EPI with standard high resolution anatomical images. Several registration schemes have been proposed for this purpose using either a free deformation model with a high number of parameters (Hellier and Barillot, 2000), or a more constrained model taking into account a priori knowledge about the main distortion direction and a difference of squares based similarity measure (Kybic et al., 2000).

This paper proposes a second improvement of the standard calculation of the diffusion tensor *D*. The linearity of Eq. (1) usually leads to a least squares-based regression method (Basser et al., 1994a). This approach, however, is not robust to the various kinds of noises that can be observed in diffusion-weighted data (Bastin et al., 1998; Basser and Pajevic, 2000; Skare et al., 2000; Anderson, 2001). Non-Gaussian noise can stem for instance from physiological motions (brain beat), subject motions (Atkinson et al., 2000; Clark et al., 2000) or residual distortions. While careful acquisition schemes including cardiac gating (Dietrich et al., 2000) and navigator echo (Butts et al., 1996; Clark et al., 2000) may reduce some of these problems, some weaknesses of the tensor diffusion model lead to other regression problems: each voxel includes several water compartments endowed with different diffusion processes that are mixed up in the data (Clark and LeBihan, 2000). Furthermore, the choice of the gradient directions used by the MR sequence can lead to very different estimation situations (Papadakis et al., 2000). Sophisticated restoration schemes dedicated to diffusion-weighted data are bound to be developed in the future to improve the situation using for instance anisotropic smoothing (Parker et al., 2000). In our opinion, however, these restoration methods cannot be perfect because of the poor quality of the raw data. Hence, in order to overcome the influence of outliers on the tensor estimation, we propose the use of a standard robust M-estimator (Meer et al., 1991). A comparison of the behaviour of both regression methods in the presence of various levels of corrupted data proves the interest of the robust approach. An earlier shorter version of this paper was published in the MICCAI proceedings (Mangin et al., 2001).

2. Distortion correction

In the following, echo-planar diffusion-weighted images were acquired in the axial plane. Blocks of eight contiguous slices were acquired, each 2.8 mm thick. Seven blocks were acquired covering the entire brain corresponding to 56 slice locations. For each slice location, 31 images were acquired; a T2-weighted image with no diffusion sensitization followed by 5 diffusion sensitized sets (*b* values linearly incremented to a maximum value of 1000 s/mm^2) in each of 6 non-collinear directions. In order to improve the signal-to-noise ratio this was repeated 4 times, providing 124 images per slice location. The image resolution was 128×128, field of view 24 cm $\times$ 24 cm, TE = 84.4 ms, TR = 2.5 s.

For each slice, each acquisition is aligned with the first image of the series, which is the standard T2-weighted image without diffusion sensitization. This alignment is done one acquisition at a time, like for fMRI motion correction procedures. For convenience we use the notation that the image is in the *XY* plane, and the phase-encoding direction lies along *Y*. Simple considerations about MR physics lead to the following distortion model (Haselgrove and Moore, 1996):

- A residual gradient in the slice-encoding direction *Z* produces uniform translation along *Y*.
- A residual gradient in the frequency-encoding direction *X* produces a shear parallel to *Y* (a translation linearly related to *X*).
- A residual gradient in the phase-encoding direction *Y* produces a uniform scaling in *Y* direction.

Hence, the geometric model (see Fig. 2) can be written for each column *X* as

$$Y' = SY + T_0 + T_1 X, \qquad (2)$$

which amounts to a simple slice-dependent affine transformation (*S* is the scale factor, T_0 a global translation, and T_1 a shear). An additional global multiplicative correction by $1/S$ has to be applied to the slice intensities, which is done after estimation of (S, T_0, T_1). This last correction stems from an energy conservation-based MR principle.

In order to estimate the affine transformation, which brings a diffusion-weighted slice into spatial alignment with the standard T2-weighted slice, a similarity measure taking into account the complex dependence between intensities embedded in Eq. (1) has to be chosen. The

stretch translation shear
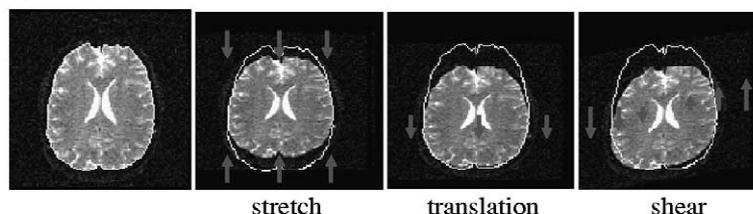
Fig. 2. The simple geometric model of eddy-current-related distortions.

purpose of this similarity measure is to return a value indicating how well the two slices match given a certain transformation (Maintz and Viergever, 1998; Hill et al., 2001). Ideally, by maximizing this similarity measure one should find the affine transformation that registers the slices. The optimal transformation, however, usually depends on the chosen similarity measure and on the implementation of its optimization. The cross-correlation measure performs poorly on this problem (Bastin, 2001) because the hypothesis of a linear relationship between the T2-weighted intensities and the diffusion-weighted ones is not verified (otherwise, diffusion imaging would be much less interesting). Furthermore, an isoset in the T2-weighted slice may correspond to a large range of values in the diffusion-weighted slice. This T2-based isoset, indeed, hides the large variability of the underlying diffusion tensors. Hence, the functional dependence assumed by the similarity measures based on weighted sums of isoset variances is not verified too (Woods et al., 1993; Roche et al., 1998).

Therefore, we have chosen to rely on the mutual information, which is a measure originating from information theory assuming the least about intensity dependence (Maes et al., 1997; Wells III et al., 1996). The underlying concept is entropy. The entropy of an image can be thought of as a measure of dispersion in the distribution of the image gray values. Given two images A and B, the definition of the mutual information MI(A,B) of these images is MI(A,B) = E(A) + E(B) − E(A,B) with E(A) and E(B) the entropies of the images A and B, respectively, and E(A,B) their joint entropy. The joint entropy E(A,B) measures the dispersion of the joint probability distribution $p(a,b)$: the probability of the occurrence of gray value a in image A and gray value b in image B (at the same position), for all a and b in the overlapping part of A and B. The joint probability distribution should have fewer and sharper peaks when the images are matched than for any case of misalignment. Therefore maximization of mutual information should correspond to the optimal affine transformation (S^{opt}, T_0^{opt}, T_1^{opt}).

Since the two images to be aligned are 128×128 slices to be compared to the usual 3D situation, a Parzen window is used to get a robust estimation of the joint intensity distribution. This Parzen window is a truncated Gaussian kernel sufficient to smooth the joint histogram. This approach turned out to be crucial to prevent the maximization algorithm to be trapped in MI local maxima. Hence, estimation of MI(S, T_0, T_1) consists of a linear resampling of the image to be aligned according to Eq. (2), followed by the application of the Parzen window to the joint histogram. Then MI can be computed from

$$\text{MI}(S, T_0, T_1) = \sum_{i_t=0}^{M-1} \sum_{i_a=0}^{M-1} p(i_t, i_a) \log \frac{p(i_t, i_a)}{p(i_t) \, p(i_a)},$$

where M is the sampling of the joint probability dis-

tribution (in practice $M = 64$, because of the few samples leading to the joint histogram computation), i_t the intensity in the target image and i_a the intensity in the image to be aligned. The marginal probabilities $p(i_t)$ and $p(i_a)$ are computed by row and column summation.

Since 123 realignments have to be performed for each slice of the volume, a fast optimization scheme is required. Fortunately, the use of a Parzen window leads to a rather smooth MI landscape around the global maximum (see Fig. 3). In some cases, several maxima have been observed near the global one. In such situations, however, we could not claim that the global maximum was a better solution than the surrounding maxima. Hence, Powell algorithm has been used in the following to maximize MI (Powell, 1964). A multiresolution approach, however, may be used in the future to get a more reliable selection of the optimal maximum (Pluim et al., 2001). For each new image, the initial position is (1,0,0), namely the no distortion situation.

Thanks to the four repetitions embedded in our acquisition process, the accuracy and the robustness of the correction process can be evaluated. A first experiment consists of comparing the results obtained using mutual information with the results obtained using another similarity measure: the correlation ratio (Roche et al., 1998). This experiment aims at evaluating whether the arguments about intensity dependence leading to choose MI could be supported by some experimental results. While similarity measures like the cross correlation assuming a linear dependence between intensities are clearly ill-adapted to our problem, it was interesting to test an intermediate measure assuming a weaker relationship but less degrees of freedom than MI. Robustness, indeed, is usually ob-
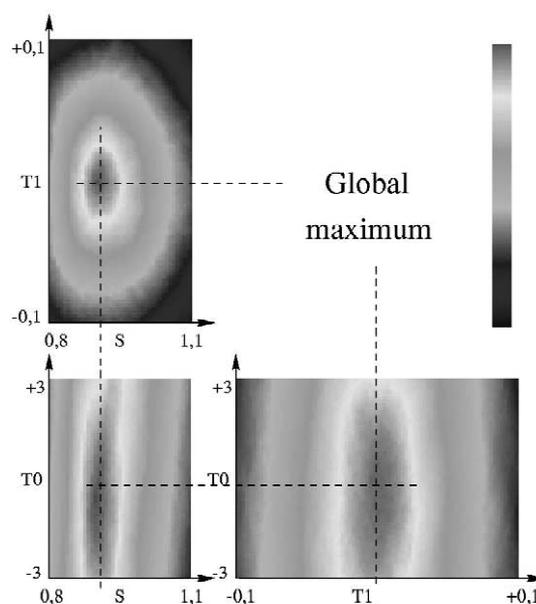
Fig. 3. Orthogonal slices of the mutual information crossing at the global maximum.

tained when strong a priori knowledge is embedded in the similarity measure.

The correlation ratio is an asymmetrical similarity measure assuming a functional dependence between both image intensities (Roche et al., 1998). This measure, which has some link with the inter-modality measure proposed in Woods et al. (1993), requires that the reference image be partitioned into a number of intensity isosets, namely broken up into areas of similar intensity. Our local implementation relies on a partition of the T2-weighted slice in 64 isosets corresponding to intensity ranges with equal length. These areas are placed over the transformed diffusion-weighted slice. Then the variance within each area is calculated and the similarity measure is defined from a weighted sum of the variances. The correlation ratio assumes that each intensity isoset in the reference image should correspond to a low dispersion intensity range in the image to be aligned.

While both methods have given similar results, the variability across the four repetitions was higher for the correlation ratio (see Fig. 4). This lower reproducibility could be predicted since the functional relationship between the two slices assumed by the correlation ratio is not verified. Therefore, this result tends to prove that mutual information is more adapted to our problem.

In general, the correction was reproducible across the four repetitions (see Fig. 4). The largest variability was observed for the translation parameter T_0, which can be understood from the shape of MI landscape (see Fig. 5). MI isophotes, indeed, are rather cylindrical with a T0 oriented axis along which some local maxima can be observed. An interesting result is the fact that the highest variability is obtained for the three repetitions of the pure T2-weighted target (no sensitizing gradient) for which $S = 1$, $T_1 = 0$ but $T_0 \neq 0$. This observation tends to prove that eddy currents have long-term trends that corrupt several consecutive acquisitions. Finally, the estimated distortions fit well with the physical interpretation mentioned above: the xy and yz gradients induce a scaling, the xz and yz gradients induce a global translation, and the xy and xz gradients induce a shearing.

3. Robust tensor estimation

The estimation of the diffusion tensor is done from linear equation (1). The B matrix in this equation depends only on the sequence, and can be computed from Bloch's equations, either formally or numerically (Mattiello et al., 1994, 1997). Each diffusion-weighted gradient choice
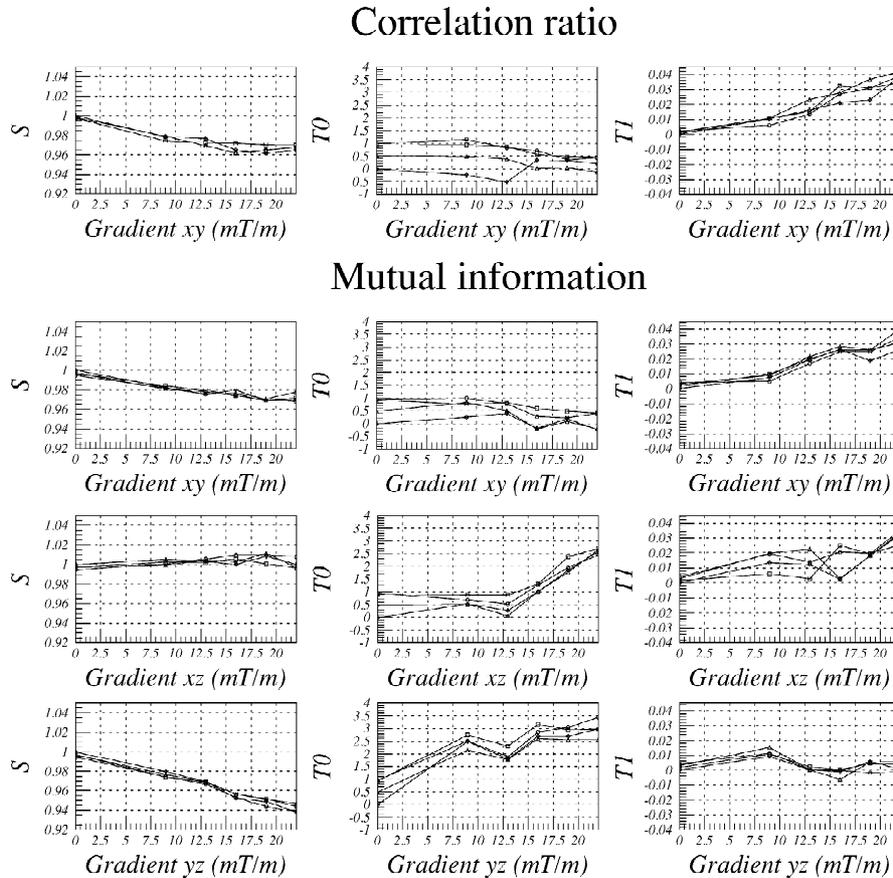
Fig. 4. Comparison between mutual information and correlation ratio for one gradient direction: (1,1,0). Reproducibility of the correction process across the four repetitions in three gradient directions, (1,1,0), (1,0,1) and (0,1,1), with six different strengths.

target   raw image   correction
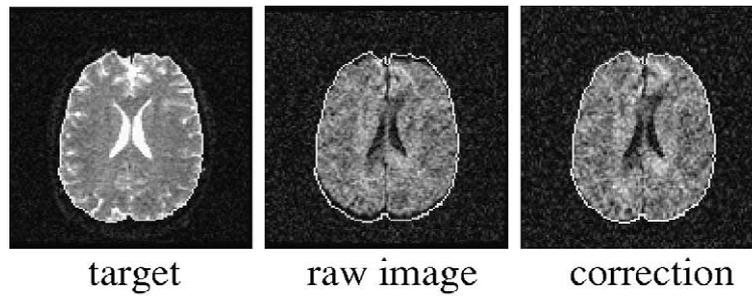
Fig. 5. Example of distortion correction.

leads to a different B matrix. To get a solvable linear system, at least six different directions of diffusion-weighted gradients have to be chosen. In the simplest case with only one gradient strength for each of six directions, the linear system is made up of six equations and six unknown values to be estimated (the symmetric tensor components). In such situations, the solution is straightforward.

In more realistic situations, however, more acquisitions are performed to improve the signal to noise ratio and a standard estimation procedure has to be designed. In the case of the data used in this paper, for instance, five different gradient strengths were used for each of six directions, leading to 30 equations. This scheme has been repeated four times leading to 120 equations. Another kind of dataset consists of using only one gradient strength but a lot of different directions (Frank, 2001). This choice allows a minimal influence of the direction choice on the tensor estimation. This influence, indeed, stems from the fact that the tensor model is usually not sufficient to perfectly describe the actual diffusion process. This remark concerns not only the problem of fiber crossing (Tuch et al., 2001; Mangin et al., 2002), but also the multi water compartment issue (Clark et al., 2000). Future work, indeed, may lead to acknowledge a complex dependence between the diffusion tensor and the gradient strength. Anyway, when the model to be estimated is a single tensor, a second model has to be chosen for noise first, then an estimator can be derived.

While tradition and ease of computation have made the least squares method the popular approach for the regression analysis yielding the tensor matrix estimation (Basser et al., 1994), this method becomes unreliable if some outliers are present in the data. The noise model underlying the least squares estimator, however, is deeply related to the Gaussian distribution. Unfortunately, MR artefacts and some weaknesses of the tensor model lead us to claim that the data do include some outliers relatively to this distribution. Robust regression methods can be used in such situations (Meer et al., 1991). The M-estimators are the more popular robust methods. These estimators minimize the sum of a symmetric, positive-definite function $\rho(\epsilon_i)$ of the residuals ϵ_i, with a unique minimum at $\epsilon_i = 0$. A residual is defined as the difference between the data point and the fitted value. For the least squares method $\rho(\epsilon_i) = \epsilon_i^2$. Several other ρ functions have been proposed which reduce the influence of large residual values on the estimated fit. We have chosen one of the most popular ones, the Geman–McLure estimator $\rho(\epsilon_i) = \epsilon_i^2 / \epsilon_i^2 + C^2$, where $C = 1.48$ median$_i\{\|\epsilon_i\|\}$. This estimator cancels out the influence of the large residuals located largely beyond C because ρ is flat outside a centered basin. In contrast, the least squares approach is mainly driven by these large residuals. The M-estimate of the diffusion tensor is obtained by converting the minimization into an iterated weighted least squares problem. The initial guess is the solution of the standard least squares.

In order to compare the behaviour of the two estimators, raw data have been corrupted with various levels of outliers (of course some actual outliers are also present in these data). For a given experiment, a percentage P of the 124 images is modified. For such images, an additional error e is added to each voxel. This error is sampled from a
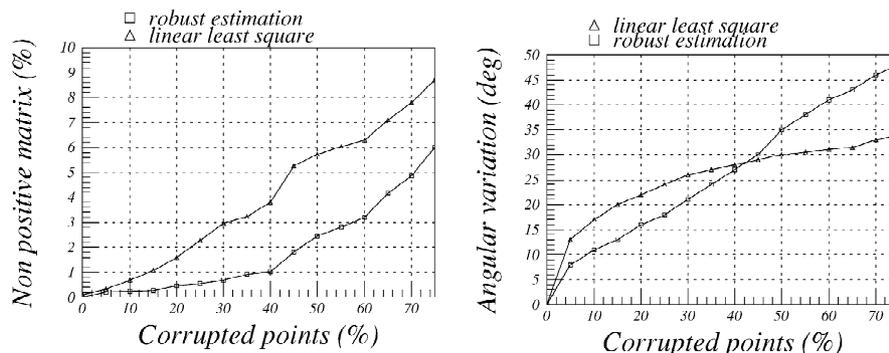
Fig. 6. Influence of outliers on the number of non-positive tensor matrices and on the direction of highest diffusion.
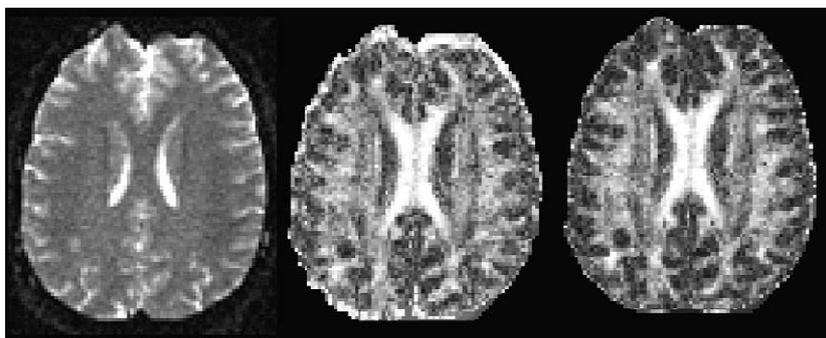
Fig. 7. Left: raw T2-weighted image. Right: tensor fractional anisotropy (Pierpaoli and Basser, 1996) without and with distortion correction and robust regressor. Fractional anisotropy is a ratio which measures the variability between the diffusion tensor eigenvalues. Without the corrections, the distortions in the Y directions create a layer of very anisotropic tensors. With the correction, this layer has almost disappeared and the boundary between grey and white matter has been largely improved.

Gaussian distribution whose mean is the mean intensity inside the brain, and whose standard deviation is a tenth of the mean. Two measures computed for the voxels located inside the brain allow us to assess the effect of these outliers. The first one is the number of non-positive estimated tensors, which have no physical interpretation. Such situations can occur because no positivity constraint is embedded in the fitting process. The second measure is the mean angular variation between the direction of highest diffusion according to the tensors estimated with and without outliers. This direction corresponds to the tensor eigenvector associated with the largest eigenvalue. The evolution of these measures relative to the percentage of outliers P is proposed in Fig. 6. The superiority of the Geman–McLure estimator is straightforward.

It should be noted that a simple way of adding a positivity constraint into the tensor estimation would consist in projecting the estimated tensor on the manifold of positive tensors. Unfortunately, this approach would amount to replacing negative eigenvalues with zero, which would lead to non-definite tensor with no more physical meaning. In our opinion, the existence of a few non-positive tensors may stem from some weaknesses of the tensor model, which is too simple to account for the complexity of the diffusion process occurring into a brain voxel. Therefore the future of diffusion imaging is in the design of more sophisticated models inferred for instance from higher sampling of the sensitizing gradient direction (Frank, 2001; Tuch et al., 1999, 2001). Whatever the model to be estimated, however, robust approaches will be required to deal with the bad quality of diffusion-weighted raw data.

4. Conclusion

This paper has presented a robust procedure to estimate the diffusion tensor from a sequence of diffusion-weighted images. This procedure is made up of a robust correction of the eddy-current-related distortions and of a robust

estimation of the tensor matrix. Further work, however, could still improve this procedure. For instance, the issue of distortion correction in the presence of subject motion remains completely open, like in the case of functional MRI. Nevertheless, our new procedure already highly improves the quality of the diffusion map which is illustrated by anisotropy images in Fig. 7. The whole scheme has been applied with success to more than 100 datasets, including numerous clinical cases.

The basic tenet underlying our procedure is the robustness to outliers that can bias either cross-correlation-based realignment procedures or least squares-based tensor estimation. While this tenet is very usual in the field of computer vision, the field of MR physics is much more used to least squares-based approaches that are bound to fail with poor data. Therefore, we think that the field of diffusion imaging may become in a near future propitious for rich collaborations between the two research fields. Foreseeable development of MR diffusion imaging will rapidly call for re-evaluation of the procedure proposed in this paper. The apparition of ultrafast parallel imaging approaches using several coils (Bammer et al., 2001) and high angular resolution imaging sequences (Frank, 2001; Tuch et al., 1999, 2001), indeed, will father new post-processing challenges.

References

Alexander, A.L., Tsuruda, J.S., Parker, D.L., 1997. Elimination of eddy current artefacts in diffusion-weighted echo-planar images: the use of bipolar gradient. Magn. Reson. Med. 38, 1016–1021.

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

Atkinson, D., Porter, D.A., Hill, D.L., Calamante, F., Connely, A., 2000. Sampling and reconstruction effects due to motion in diffusion-weighted interleaved echo planar imaging. Magn. Reson. Med. 44 (1), 101–109.

Bammer, R., Keeling, S.L., Augstin, M., Pruessmann, K.P., Wolf, R., Stollberger, R., Hartung, H.P., Fazekas, F., 2001. Improved diffusion-weighted single-shot echo-planar imaging (EPI) in stroke using sensitivity encoding (SENSE). Magn. Reson. Med. 46 (3), 548–554.

Basser, P.J., Pajevic, S., 2000. Statistical artifacts in diffusion tensor MRI (DT-MRI) caused by background noise. Magn. Reson. Med. 44 (1), 41–50.

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

Basser, P.J., Mattiello, J., LeBihan, D., 1994b. MR diffusion tensor spectroscopy and imaging. Biophys. J. 66, 259–267.

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

Bastin, M.E., 2001. On the use of the FLAIR technique to improve the correction of eddy current induced artefacts in MR diffusion tensor imaging. Magn. Reson. Imag. 19 (7), 937–950.

Bastin, M.E., Armitage, P.A., 2000. On the use of water phantom images to calibrate and correct eddy current induced artefacts in MR diffusion tensor imaging. Magn. Reson. Imag. 18 (6), 681–687.

Bastin, M.E., Armitage, P.A., Marshall, I., 1998. A theoretical study of the effect of experimental noise on the measurement of anisotropy in diffusion imaging. Magn. Reson. Imag. 16 (7), 773–785.

Butts, K., de Crespigny, A., Pauly, A.J.M., Moseley, M., 1996. Diffusion-weighted interleaved echo-planar imaging with a pair of orthogonal navigator echoes. Magn. Reson. Med. 35 (5), 763–770.

Calamante, F., Porter, D.A., Gadian, D.G., Connelly, A., 1999. Correction for eddy current induced B0 shifts in diffusion-weighted echo-planar imaging. Magn. Reson. Med. 41, 95–102.

Clark, C.A., LeBihan, D., 2000. Water diffusion compartmentation and anisotropy at high b values in the human brain. Magn. Reson. Med. 44, 852–859.

Clark, C.A., Barker, G.J., Tofts, P.S., 2000. Improved reduction of motion artifacts in diffusion imaging using navigator echoes and velocity compensation. J. Magn. Reson. 142 (2), 358–363.

Dietrich, O., Heiland, S., Benner, T., Sartor, K., 2000. Reducing motion artefacts in diffusion-weighted MRI of the brain: efficacy of navigator echo correction and pulse triggering. Neuroradiology 42 (2), 85–91.

Frank, L.R., 2001. Anisotropy in high angular resolution diffusion-weighted MRI. Magn. Reson. Med. 45 (6), 935–939.

Haselgrove, J.C., Moore, J.R., 1996. Correction for distorsion of echo-planar images used to calculate the apparent diffusion coefficient. Magn. Reson. Med 36.

Hellier, P., Barillot, C., 2000. Multimodal non-rigid warping for correction of distortions in functional MRI. In: MICCAI, LNCS-1935. Springer Verlag, pp. 512–520.

Hill, D.L.G., Batchelor, P.G., Holden, M., Hawkes, D.J., 2001. Medical image registration. Phys. Med. Biol. 46, R1–R45.

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

Jezzard, P., Clare, S., 1999. Sources of distortion in functional MRI data. Hum. Brain Mapp. 8 (2–3), 80–85.

Jezzard, P., Barnett, A., Pierpaoli, C., 1998. Characterization of and correction for eddy current artefacts in echo planar diffusion imaging. Magn. Reson. Med. 39, 801–812.

Kybic, J., Thévenaz, P., Nirkko, A., Unser, M., 2000. Unwarping of EPI images. IEEE Trans. Med. Imag. 19 (2), 80–93.

Le Bihan, D., Mangin, J.-F., Poupon, C., Clark, C.A., Pappata, S., Molko, N., Chabriat, H., 2001. Diffusion tensor imaging: concepts and applications. J. Magn. Reson. Imag. 13, 534–546.

Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P., 1997. Multimodality image registration by maximisation of mutual information. IEEE Trans. Med. Imag. 16 (2), 187–198.

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

Mangin, J.-F., Poupon, C., Clark, C.A., LeBihan, D., Bloch, I., 2001. Eddy-current distortion correction and robust tensor estimation for MR diffusion imaging. In: MICCAI '01, Utrecht, The Netherlands, LNCS-2208. Springer Verlag, pp. 186–194.

Mangin, J.-F., Poupon, C., Cointepas, Y., Rivire, D., Papadopoulos-Orfanos, D., Clark, C.A., Rgis, J., Le Bihan, D. A framework based on spin glass models for the inference of anatomical connectivity from diffusion-weighted MR data. NMR Biomed., in press.

Mattiello, J., Basser, P.J., LeBihan, D., 1994. Analytical expression for the b matrix in NMR diffusion imaging and spectroscopy. J. Magn. Reson. A 108, 131–141.

Mattiello, J., Basser, P.J., LeBihan, D., 1997. The b matrix in diffusion tensor echo-planar imaging. Magn. Reson. Med. 37, 292–300.

Meer, P., Mintz, D., Rosenfeld, A., 1991. Robust regression methods for computer vision: a review. Int. J. Comput. Vis. 6 (1), 59–70.

Papadakis, N.G., Murrills, C.D., Hall, L.D., Huang, C.L., Carpenter, A.T., 2000. Minimal gradient encoding for robust estimation of diffusion anisotropy. Magn. Reson. Imag. 18 (6), 671–679.

Parker, G.J., Schnabel, J.A., Symms, M.R., Werring, D.J., Barker, G.J., 2000. Nonlinear smoothing for reduction of systematic and random errors in diffusion tensor imaging. J. Magn. Reson. Imag. 11 (6), 702–710.

Pierpaoli, C., Basser, P.J., 1996. Toward a quantitative assessment of diffusion anisotropy. Magn. Reson. Med. 36, 893–906.

Pluim, J.P.W., Maintz, J.B.A., Viergever, M.A., 2001. Mutual information matching in multiresolution contexts. Image Vis. Comput. 19 (1–2), 45–52.

Poupon, C., Mangin, J.-F., Clark, C.A., Frouin, V., LeBihan, D., Bloch, I., 2001. Towards inference of the human brain connectivity from MR diffusion tensor data. Medical Image Analysis 5, 1–15.

Powell, M., 1964. An efficient method for finding the minimum of a function of several variables without calculating derivatives. Comput. J. 7, 155–162.

Roche, A., Malandain, G., Pennec, X., Ayache, N., 1998. The correlation ratio as a new similarity measure for multimodal image registration. In: MICCAI '98, MIT, USA, LNCS-1496. Springer Verlag, pp. 1115–1124.

Skare, S., Hedehus, M., Moseley, M.E., Li, T.Q., 2000. Noise considerations in the determination of diffusion tensor anisotropy. Magn. Reson. Imag. 18 (6), 659–669.

Stejskal, E.O., Tanner, J.E., 1965. Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. J. Chem. Phys. 42 (1), 288–292.

Tuch, D.S., Weisskoff, R.M., Belliveau, J.W., Wedeen, V.J., 1999. High angular resolution diffusion imaging of the human brain. In: VIIth ISMRM, Philadelphia, USA.

Tuch, D.S., Wiegell, M.R., Reese, T.G., Belliveau, J.W., Van Wedeen, J., 2001. Measuring cortico-cortical connectivity matrices with diffusion spectrum imaging. In: ISMRM-ESMRMB, Glasgow, p. 502.

Wells III, W., Viola, P., Atsumi, H., Nakajima, S., Kikinis, R., 1996. Multi-modal volume registration by maximization of mutual information. Medical Image Analysis 1, 35–51.

Woods, R.P., Mazziotta, J.C., Cherry, S.R., 1993. MRI-PET registration with automated algorithm. J. Comput. Assist. Tomogr. 17, 536–546.