NeuroImage **18,** 127–142 (2003)
doi:10.1006/nimg.2002.1281

An Evaluation of the Use of Magnetic Field Maps to Undistort Echo-Planar Images

Rhodri Cusack, Matthew Brett, and Katja Osswald

MRC Cognition and Brain Sciences Unit, Cambridge, United Kingdom

Received April 2, 2002

When a head is placed in an MRI scanner, differences between the magnetic susceptibility of tissue, bone, and air distort the magnetic field. While some of the resulting inhomogeneity can be corrected by the shimming process, much of it cannot, and this causes distortion (sometimes referred to as geometric distortion) of echo-planar images (EPIs). One strategy for the correction of distortion is to acquire a map of the magnetic field achieved in each subject and then to use this to undistort their EPIs after reconstruction (Jezzard and Balaban, 1995). Here, we present five experiments to evaluate the application of such a strategy on data from a 3-T scanner. We show that after undistortion, the shape of EPIs is more similar to the true shape of the brain, and we investigate the effect of head movement on the efficacy of undistortion. If undistortion was applied first, it was found that less nonlinear warping was required to transform EPIs into a standard space, particularly in the phase-encode direction. We show that if SPM 99 normalization is used to perform a nonlinear warp to standard space, the prior application of undistortion increases the statistical power of group studies with motor and auditory tasks. We show that this increase in power is due to an increase in the overlap of activation of different subjects. Finally, we evaluate where in the brain undistorting EPIs might be expected to have the greatest effect, in terms both of mislocalization of activation and of a reduction in power. Overall, undistorting EPIs using field maps has proved extremely successful, improving the anatomical localization of activation and increasing statistical power.

INTRODUCTION

Ideally, in the absence of applied gradients, the magnetic field in an MRI scanner would be homogenous throughout the bore. Unfortunately, different materials affect magnetic fields in different ways—a property summarized by their *susceptibility*. Biological tissue, which is mostly water, develops small magnetization that acts to counteract the applied field (it is diamagnetic—weak negative susceptibility). Bone and air have little effect on the magnetic field (near zero susceptibility). When a structure comprising materials of different susceptibilities, such as a head, is put into a uniform field, the field becomes distorted. The resulting inhomogeneities can be partially tackled by shimming, in which small correction gradients are applied. Scanners often have first order shim coils—which remove linear gradients across the head—and second order coils, which take out other gradual field changes. However, even after shimming there will often remain substantial variation that cannot be removed.

Hence, we cannot create a completely homogenous field, and this can effect acquisitions. There are two problems: dropout and distortion. Other authors have characterized dropout (Ojemann *et al.*, 1997; Devlin *et al.*, 2000) and suggested ways in which it might be alleviated (Deichmann *et al.*, 2002; Chen and Wyrwicz, 1999b), and we do not discuss this here. We focus instead on the problem of distortion. To build three-dimensional images, the spatial location of the signal from individual voxels is encoded by the application of gradients in the magnetic field strength across the bore. These gradients are much weaker than the main field (perhaps 10 mT as opposed to several Tesla) and even small inhomogeneities in the main field can lead to substantial distortions of the spatial encoding gradients and hence substantial distortions in the final MR images.

Several factors determine the extent to which we will suffer distortions in an acquisition. First, they are made worse as the field strength increases. Inhomogeneities result from the effect of the susceptibility differences on the static field and so get larger as the static field becomes stronger. However, there are several different constraints on how strong the spatial encoding gradients can be—the coils must produce a linear gradient, change the field very quickly, and be acoustically quiet (Jezzard and Clare, 1999)—and these can make it difficult to increase their strength. Hence, as the gradients are of similar strength, but the inhomogeneity in the static field is larger, distortions are generally worse at higher field strengths. Second, different sequences produce very different artefacts. As shown by Jezzard and Balaban (1995), echo-planar imaging (EPI) sequences show very small distortions in the read direction—less than one hundredth of a voxel—but much more significant deviations in the phase-encode direction—for some parts of the image of the order of a few voxels. Non-EPI sequences that acquire a single line of

1053-8119/02 $35.00

k-space after each pulse show very small distortions in the read direction (around 0.1 voxels) and no distortions in the phase-encode direction. Here, we consider only distortions in the phase-encode direction of EPI acquisitions.

In a typical fMRI experiment, functional data are acquired using an EPI sequence. Then, for each subject, a high-resolution structural scan is acquired using a non-EPI sequence. The distortion of EPIs in the phase-encode direction leads to two related problems. First, as it makes the EPIs a different shape to the structural scans (which are not distorted) it becomes impossible to accurately coregister the functional and structural data of a subject using standard rigid body transforms. This makes it difficult to accurately specify which part of the brain anatomy is activated. Second, as the distortions are different in each subject, it can be difficult to normalize subjects into the same standard space. Even in the absence of distortion, different brains are of slightly different sizes, and slightly different in shape, and some warping is required to put images from different subjects into a common space. This common space is usually defined by a standard template image, and several automatic or semiautomatic algorithms exist for calculating the required transformation. In SPM 99 (http://www.fil.ion.ucl.ac.uk/spm) this is referred to as "spatial normalization." The structural scan is often used to determine the warping required to transform a particular subject's brain into standard space. However, if distortion makes it difficult to accurately register each subject's EPIs to their structural, then the correct warping transform will not be applied to the EPIs. This can make it difficult to perform group studies, as different parts of the brains of different subjects will be compared.

To avoid the problematic coregistration of EPIs and structural scans, a different approach that may be used is to determine the warping required to transform EPIs into standard space directly from the EPIs. We might then hope that the distortions will be corrected by this warping process. However, due to the much lower resolution than structural scans, and poor contrast between tissues, normalization using EPIs is less robust. An improvement on the basic normalization process is also to specify a mask identifying regions of signal dropout and remove these from the matching process, in a similar way as proposed for the normalization of damaged brains (Brett *et al.*, 2001). Another approach is to use warping procedures specifically designed to handle the distortions such as those proposed by Kybic and colleagues (2000).

Rather than relying on such fitting procedures, a more elegant alternative strategy is to directly measure the inhomogeneities in the field in each subject (Weisskoff and Kiihne, 1992). These field map measurements can then used to undistort the EPI images, after they have been reconstructed (Jezzard and Balaban, 1995; Reber *et al.*, 1998; Chen and Wyrwicz, 1999a; Munger *et al.*, 2000; Jenkinson, 2001; Hutton *et al.*, 2001). The aim of this study was to describe a procedure in which this has been implemented and to evaluate the results. We tested whether undistorting leads to an improvement in the similarity of functional and structural images for a single subject (cf. Hutton *et al.*, 2001) and also whether movement across the time series reduces the efficacy of undistortion. We also evaluated whether undistortion makes data from each subject more similar to the standard template. We then tested whether it improved the overlap of data from different subjects, leading to greater statistical power, and finally evaluated for which areas in the brain undistortion is likely to be most important.

GENERAL PROCEDURE

Overview

In this section we describe the procedure for undistorting. Figure 1 shows an overview of the procedure, and Fig. 2 the data at various stages of processing. The aim is to use the information about the magnetic field inhomogeneities to shift the voxels of the EPIs in the phase-encode direction to their true locations. The magnetic field measurements are contained in the phase of the field map acquisition. However, before it can be used, it must be unwrapped, undergo some additional processing (dilation, smoothing, and rotation), and then be coregistered with the EPIs. Each of the stages is described below.

Acquisitions

Three sorts of data were used. We acquired T2*-weighted EPI data of the sort which is used to obtain an estimate of BOLD in functional studies. However, apart from in experiments 3 and 4, we will not examine the effects of the functional manipulations present in the experiments, as we primarily studied the shape of the EPI images. The second acquisition was a gradient-echo non-EPI magnetic field map image. We took two acquisitions that were identical in all ways apart from their echo time. We then subtracted the phases of the two images, which can be used to derive at each voxel the deviations of the magnetic field from the desired homogenous field. These field maps took 2 min 6 s each to acquire. The third acquisition was a T1-weighted high-resolution structural scan, which took 12 min to acquire.

All data were acquired on a Bruker MEDSPEC 3-T scanner at the Wolfson Brain Imaging Centre in Cambridge, using a head coil gradient set. The EPI data were collected with one of two sets of acquisition parameters: (1) 21 planes with a matrix size of 64 × 64; TE, 27 ms; 100-kHz voxel bandwidth (equivalent to a 1562.5-Hz bandwidth in the phase-encode direction), FA, 90°; TA, 2.2 s; and TR, 2.34 s (hereafter referred to as 64 × 64 EPIs); (2) 21 planes with a matrix size of 128 × 128; TE, 27 ms; FA, 90°; TA, 2.2 s; 200-kHz voxel bandwidth (again equivalent to a 1562.5-Hz bandwidth in the phase-encode direction); and TR, 10 s (hereafter referred to as 128 × 128 EPIs). The short TA and long TR were used because we presented auditory stimuli and didn't want the acoustic noise from the acquisition to mask the stimuli or the scanner noise to saturate the BOLD response (sparse imaging, see Hall *et al.*, 1999).

For both, phase encoding was in the anterior–posterior direction, the FOV was 250 × 250 mm, the slice thickness 4 mm, and interslice gap 1 mm. In the 128 × 128 EPIs, the acquisitions were angled to avoid the eyeballs. The field map data were acquired using a 3D gradient echo sequence

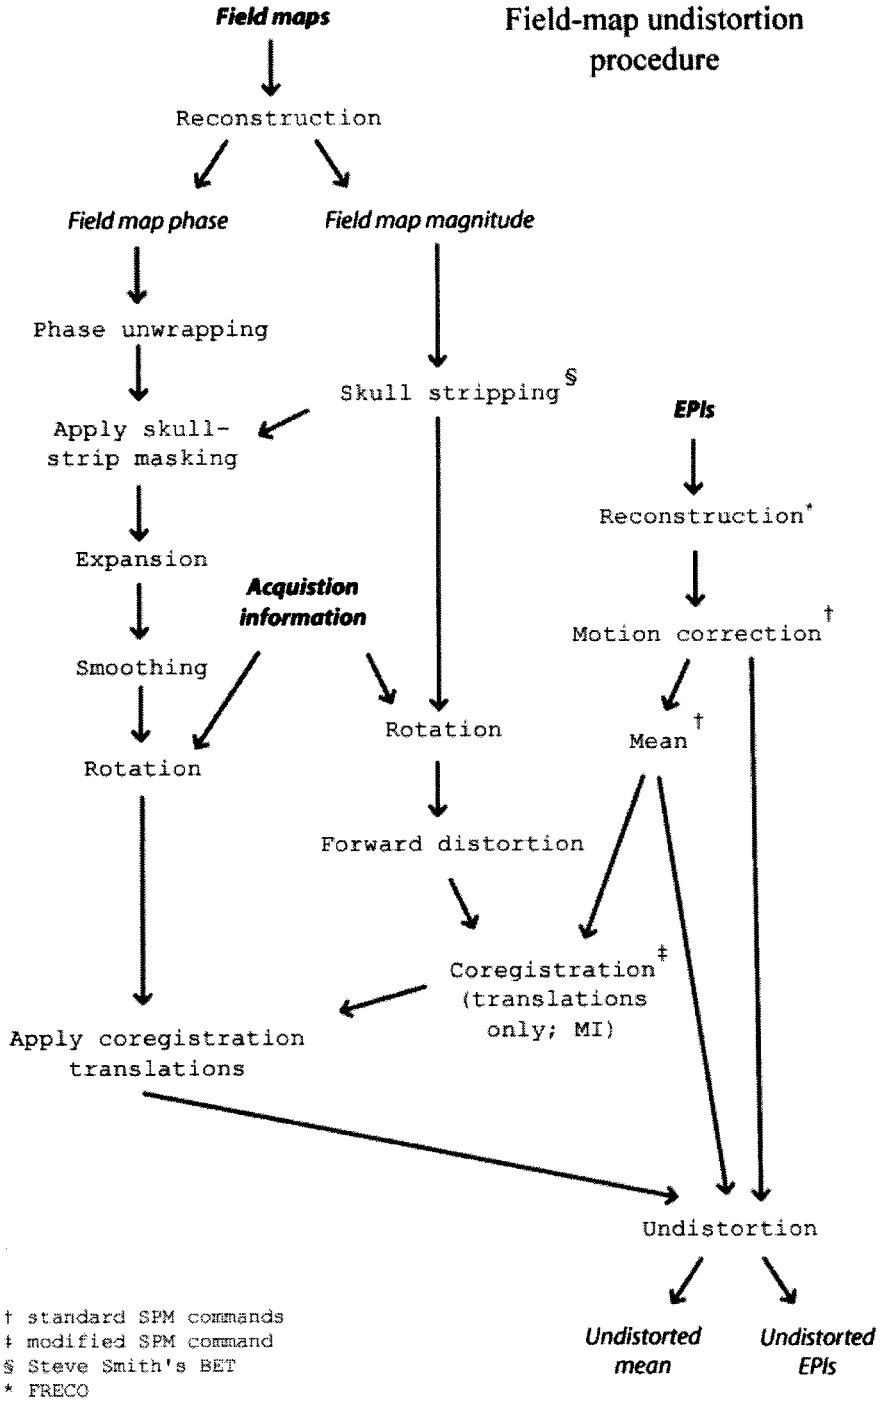

FIG. 1. An overview of the undistortion procedure.

(Bruker "GEFI TOMO") with a matrix size of 64 × 256 × 64 and a field of view 256 × 256 × 256 mm. Two volumes were acquired, differing only in their TE, which was either 1 or 10.104 ms. These values were chosen to differ by 9.104 ms, which is a multiple of the precession time for the fat signal at 3 T, calculated to be 2.27 ms. The phase of the fat signal in the short and long TE acquisitions should be similar, and it will cancel when their phases are subtracted. The structural scan was collected using a 3D T1-weighted gradient-echo sequence (Bruker GEFI TOMO; TE, 5 ms), with a matrix size of 256 × 256 × 256 and a FOV of 180 × 220 × 256 mm.

Reconstruction and Initial Processing

The EPI data were reconstructed using FRECO (Hennel, 1999). Figure 2g shows an example mean EPI image. Both

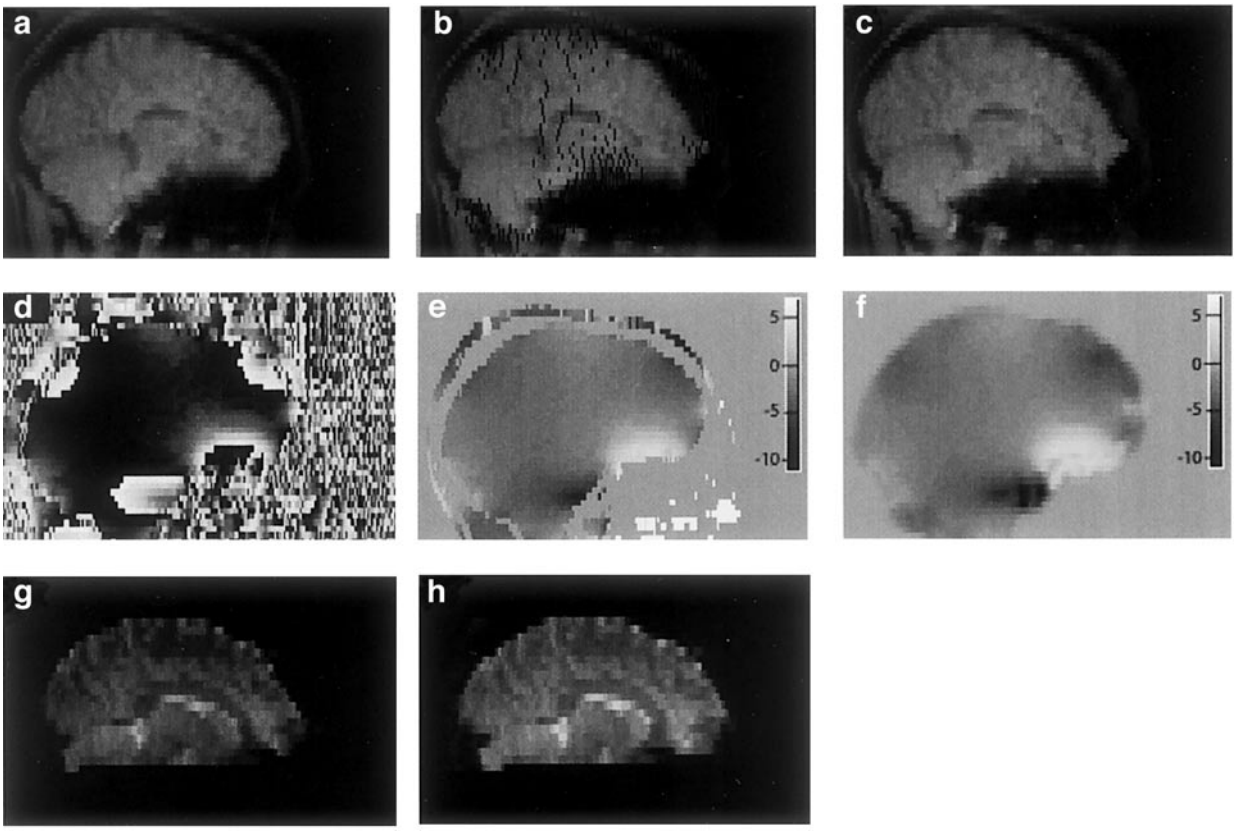

FIG. 2. The data at various stages of processing: (a) raw magnitude of field map; (b) magntiude distorted using phase information; (c) magnitude distorted and interpolated; (d) raw phase of field map; (e) unwrapped phase; (f) unwrapped and expanded phase; (g) raw EPI; (h) undistorted EPI.

field map images were reconstructed to form complex images with the phase information preserved. Each voxel in the TE = 1 ms volume was normalized to a magnitude of unity and then divided into the TE = 10.104 ms volume in complex mode. The phase and magnitude of this resultant image were then passed on to later stages for further processing. Figures 2a and 2d show the magnitude and phase of the field map, respectively.

Phase Unwrapping

The magnetic field strength is encoded in the phase of the MR signal. Unfortunately, this phase is wrapped into the range $[-\pi, +\pi]$, and to obtain a measure of the underlying magnetic field the missing multiples of 2π must be replaced. We used a robust three-dimensional phase unwrapping algorithm described elsewhere (Cusack and Papadakis, in press). An example of a wrapped field map and the resultant unwrapped image are shown in Figs. 2d and 2e. As with all spatial unwrapping algorithms, only relative phase unwrapping can be performed and there may be an absolute offset, which is an integer multiple of 2π, affecting the whole image. Our procedure for dealing with this is described in the section after the next one.

Masking, Dilation, and Conversion to Distortion Field

The field map acquisition only gives useable data within the brain. Outside the brain, the signal strength falls off rapidly and the signal to noise ratio of the phase data drops unacceptably. To avoid this noisy data interfering, we masked the field maps. This was done using the Brain Extraction Tool (http://www.fmrib.ox.ac.uk/analysis/techrep/) on the magnitude image and then applying the same estimate of the brain outline to the phase information. After masking, there will be abrupt changes in the phase around the edge of the mask. Given that there might be areas of strong signal in the EPIs where we have no phase information, due to small registration errors, or EPI signal from the scalp, this abrupt change in phase (and hence distortion) might lead to artifacts in the final undistorted image. To overcome this problem, we dilated the phase image for 10 mm, replacing each voxel with the value of the closest one for which we did have information. To enhance the signal-to-noise ratio, this map was then smoothed using a 3D Gaussian kernel with a FWHM of 5 mm. This was chosen by visual inspection of the spatial variation of the field maps.

Finally, we converted the masked, dilated, and smoothed phase map to a distortion field using the following transformation at each voxel (Jezzard and Balaban, 1995):

$$\Delta y(r) = l_y \frac{\Phi(r)}{2\pi b_y \Delta t},$$

where r indexes a particular voxel, $\Phi(r)$ is the phase of the unwrapped map, Δy is shift in voxels in the phase-encode direction, Δt is the difference in echo times between the two field map acquisitions, b_y is the bandwidth in the phase-encode direction, and l_y is number of voxels in the phase-encode direction.

Coregistration of Distorted Magnitude Image

Before the field map information can be used to undistort our EPI volumes, it is important that the two are in registration. From our acquisitions, we had information about the relative angle of the field map and EPI volumes, but not the exact translation required to put them in registration. So, we initially rotated the field map to the same orientation as the EPI volumes and then coregistered them. However, as the field maps are acquired with a non-EPI sequence, then they are not distorted, and hence, the EPIs (see Fig. 2g) and the field maps (see Fig. 2a) are different shapes, and we cannot coregister them directly. We cannot undistort the EPIs until we have coregistered them with the field maps. However, we can forward distort the magnitude of the field maps (Fig. 2c). After distortion, this should have the same shape as the EPIs, and we can coregister them.

To forward distort the field map magnitude we used a pixel shift approach. For each voxel in the magnitude of the field map, we shifted it by the value of the distortion field at that voxel. A resultant image is shown in Fig. 2b. Note that as this is a "push" problem—we know the forward transformation—where the distortion is expanding the image, not all voxels in the destination image have been covered, and black lines result. These were removed by taking each voxel with no value in turn and performing a weighted average of all the values in a Gaussian kernel around it (similar to a smooth) that did have a value. Figure 2c shows the result. This smoothed distorted magnitude image was then coregistered with the EPIs using SPMs "mutual information" coregistration, set to allow only translations. The same translations that put the distorted magnitude image in registration were then applied to the distortion field derived from the field map. This was also resliced to the same matrix size as the EPIs for later processing.

This coregistration stage has the additional advantage that it makes the undistortion procedure robust to offset displacements affecting the whole image, which might be introduced as a result of wrapping and then unwrapping (see *Phase Unwrapping* above). If a fixed offset is added to the whole field map image, then this will result in a fixed translation of the entire distorted magnitude image in the phase-encode direction. The EPIs are then coregistered with this, and when undistortion is performed, the translation will be reversed. Provided that no part of the distorted magnitude image is translated to a point outside the field of view, then any fixed offset to the whole distortion field will have no effect on the final image. To reduce the chance of any part of the image leaving the field of view, our field maps were adjusted by adding or subtracting the appropriate whole multiples of 2π so that the mean phase of the unwrapped image fell into the range $[-\pi, \pi]$. In practice, none of the distorted magnitude images came close to leaving the FOV.

Undistortion

Once the distortion field and the EPIs are in registration, the EPIs were undistorted. This was done using the pixel shift method (Jezzard and Balaban, 1995) without intensity correction. Resampling was then done using Hanning windowed one-dimensional sinc interpolation with a range of 10 voxels in each direction. In experiment 5 we provide three different ways of looking at the patterns of distortion applied. A video showing the pattern of undistortion is also available on our website (http://www.mrc-cbu.cam.ac.uk/Imaging/fieldmap_undistort).

Motion and Distortion

It is established that head movement leads to some change in the pattern of field inhomgeneities and hence pattern of distortion (e.g., Andersson *et al.*, 2001; Hutton *et al.*, 2002). However, it may still be that a substantial component of the distortion is constant within a subject across time and may be usefully corrected using a single high-resolution field map. As a first assessment, the pattern of distortions was examined by eye at different points in the time series. A typical example is shown in Fig. 3. By eye, it appears that the pattern of distortions is largely similar over time when the movements are in this typical range. Furthermore, in experiment 1 we explicitly tested for the validity of our approximation—but it should be noted that for movements that are substantially larger than those in our studies, this approximation may be violated, and the procedure would be expected to be less powerful.

It is not clear *a priori* whether it is better to perform motion correction before or after undistortion. It might be that to a first approximation, field inhomogeneities remain fixed with respect to the reference frame of the head or alternatively that they remain fixed with respect to the scanner. Which of these approximations is most accurate was studied in pilot work (to be reported elsewhere) that found that it was better to perform undistortion after realignment, and so this is what we have done here.

EXPERIMENT 1

Aim and Method

There were two aims to this experiment. First, we wished to test whether the undistortion process made the shape of the EPIs more like the true shape of the brain, as measured using the non-EPI structural scan. Second, we wanted to examine the importance of the movement by distortion interaction. To do this, we tested whether undistortion performed better when applied to an EPI volume which was acquired closer in time to the acquisition of the field maps than for one which was collected further away in time. It is usual that volumes acquired at a very different time within the same

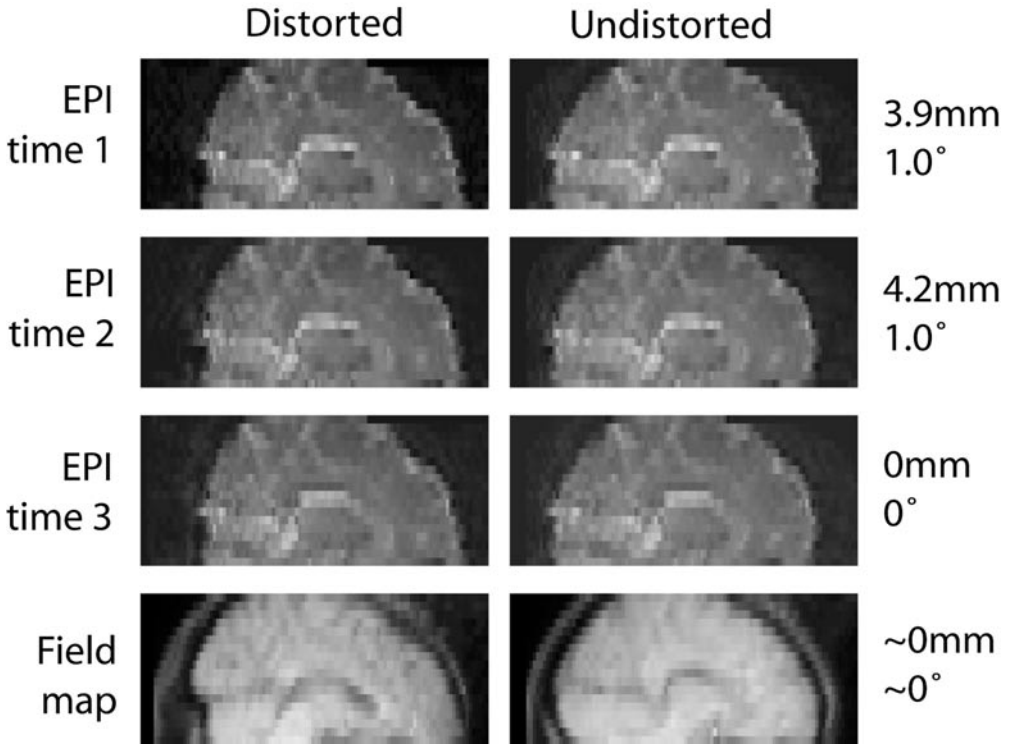

FIG. 3. The top three rows show the a representative sagittal slices from the raw (left column) and undistorted (right column) EPIs at three different time points. The bottom row shows the magnitude of the field map as acquired without distortions (right column) and after distortion (left column). The numbers on the far right show the movement parameters relative to the third time point, summarized by the square root of the sum of the squares of the translations and rotations. The field map was acquired soon after the third time point and so the head positions should be similar.

session are displaced due to gradual subject movement, and undistortion might then be less efficient.

Echo-planar data from 9 subjects were acquired using the 128 × 128 EPI parameters and 12 using the 64 × 64 EPI parameters. On all subjects, we acquired a T1-weighted high-resolution structural image. We also acquired the field maps, as described above.

As a first step, we examined the change in shape by eye. To do this, for each subject we used SPM to coregister the mean EPI (as produced by realignment) with the structural, using the mutual information cost function. The undistorted mean EPI and structural were also coregistered in the same way. To assess quantitatively the change in registration with undistortion, we used the mutual information measure in a similar way to Hutton *et al.* (2002). This measure, which has its origins in information theory, expresses the degree of shared information in two images and is commonly used as a cost function for the optimization of matching of images with different tissue contrasts (Collignon *et al.*, 1995). We obtained the measure by recording the negated final value of the cost function displayed by SPM during mutual information coregistration. The higher this mutual information measure, the better the match between the volumes.

For comparison of the 64 × 64 EPIs, we used the 10th image in the sequence, to allow the spin history to reach a steady state. For the 128 × 128 EPIs, because of the long TR, we used the second image. We also analyzed the last image of each.

Results and Discussion

Figure 4 shows sagittal slices through the high-resolution structural, the mean EPI, and the mean undistorted EPI for the 9 subjects for which 128 × 128 EPIs were acquired. Note the general improvement in shape visible by eye in several of the subjects. Analysis of the regions of the brain in which undistortion has the most effect is given in experiment 5. Then, as a quantitative measure, Fig. 5 shows the effect of undistortion on the mutual information measure between the sample EPI and the T1-weighted structural scan for 21 subjects—9 of which were scanned at 128 × 128 (subjects 1–9) and 12 scanned at 64 × 64 (subjects 10–21). It can be seen that according to this measure, all but one of the subjects showed an improved match after undistortion. A repeated measures ANOVA confirmed a significant effect of undistortion ($F(1,20) = 59.0$, $P < 0.0001$).

Table 1 shows the movement parameters as produced by the realignment algorithm for the nine subjects on which the 128 × 128 acquisition was used. It can be seen that there was substantial movement between the second and last EPIs. The field maps were acquired after the final EPI, and so the head should have been in a similar position. To allow us to compare the effect of movement on the quality of undistortion, we compared the degree of improvement following undistortion in coregistration between the second EPIs and the structurals to the improvement following undistortion for the final EPI and the structurals. As Fig. 6 shows, there was a similar

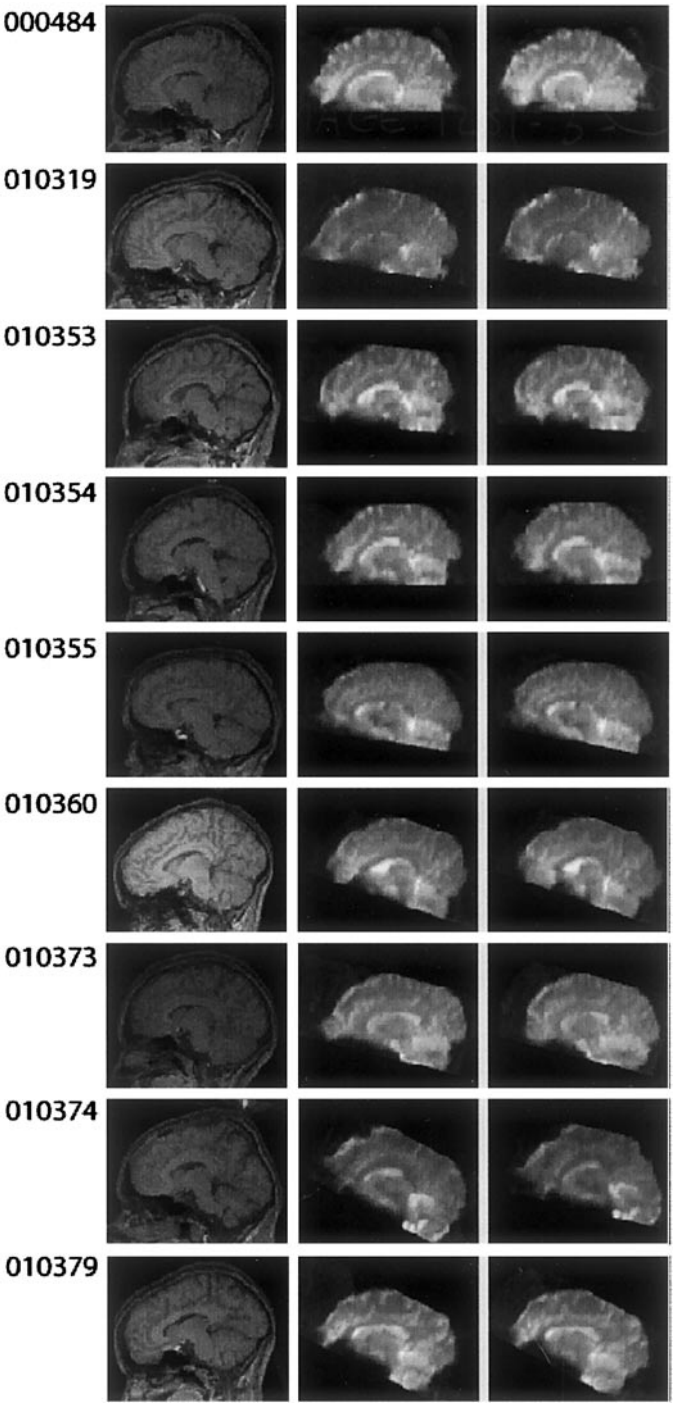

FIG. 4. For nine subjects (one per row), sagittal slices through high-resolution structural images (first column); mean EPI before undistortion (second column); and the mean EPI after undistortion (third column) are shown. Phase encoding was in the anterior–posterior direction.

improvement for the second EPI as for the last. This suggests that the second EPI has been undistorted just as well as the last and hence that the movement by distortion interaction is small by comparison with the improvement due to the undistortion.

Finally, we guard against a possible alternative explanation for the results. It is possible that the resampling procedure that forms part of undistortion will have added some slight smoothing to the image and hence increased its mutual information match. To check this, we took the final volumes before undistortion from the nine subjects from the 128 × 128 acquisitions and smoothed them using a 3-mm FWHM Gaussian kernel. We then recalculated the mutual information difference between the undistorted (and unsmoothed) relative to this new baseline. As can be seen by eye from Fig. 7, even this small kernel leads to substantially more smoothing than the resampling. However, as can be seen from the third column in Fig. 6, even relative to this baseline there is a substantial increase in the mutual information measure. This was confirmed using a repeated measures ANOVA ($F(1,8) = 18.6$; $P < 0.005$).

EXPERIMENT 2

Aim

In experiment 1, we used the mutual information measure to demonstrate that after undistortion, the EPIs were more similar to the true brain shape of each subject. In this experiment, we used a different measure to demonstrate that in general the EPIs become more similar to the standard EPI template image supplied with SPM 99. We used the SPM normalization routines, which apply affine and nonlinear transformations to optimize the match between an object and target image. The matching process, described in detail by Friston *et al.* (1995), Ashburner *et al.* (1997), and Ashburner and Friston (1999), simultaneously optimizes the match between the warped object and the likelihood of a particular set of warps being required.

We normalized both the mean of the EPIs before undistortion (as generated by realignment) and the undistorted mean to the standard EPI template. We then used the deformations toolbox (J. Ashburner, 2001, http://www.jiscmail.ac.uk/

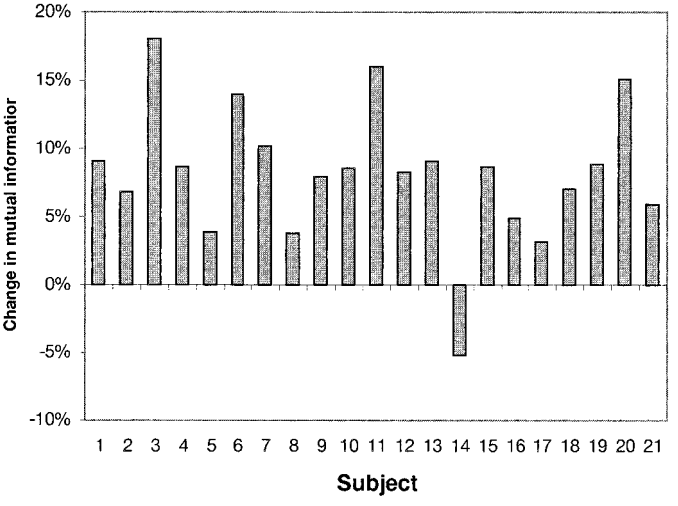

FIG. 5. Change in mutual information between an EPI and structural image when undistortion is applied.

TABLE 1

Displacement of the Second EPI Volume Relative to the Final One Summarized over Nine Subjects in Experiment 1

	Translations (mm)			Rotations around axis (degrees)		
	x	*y*	*z*	*x*	*y*	*z*
RMS	0.38	0.49	1.21	0.68	0.49	0.31
Minimum	−1.96	−1.11	−4.69	−1.05	−2.20	−0.89
Maximum	0.79	2.41	4.20	5.67	0.74	1.45

cgi-bin/wa.exe?A2=ind0107&L=spm&P=R7813) to examine the nature of the warping fields required. It was expected that less warping would be required when undistortion had been performed and, more specifically, that normalization would have to perform less warping in the phase-encode direction.

Method

We used data from 14 subjects acquired using the 64 × 64 EPI protocol. The mean of 202 volumes produced by realignment, and the same means after undistortion, were normalized to the standard EPI template supplied with SPM 99. The warping fields were then examined using SPM in two ways. First, a mask was made from the standard EPI template so that only the warping of voxels within the brain would count. This was done by applying the Brain Extraction Tool (Steve Smith, http://www.fmrib.ox.ac.uk/analysis/techrep/) to the EPI template, which finds the edge of the brain and then sets all voxels outside it to zero. Later statistics were only calculated for voxels that had nonzero values in this masked template. To assess warping, the Jacobean determinant of the strain field was calculated for each voxel that fell inside the mask. This gives a measure of the change in volume of each voxel. We wanted to focus upon warping, rather than whole brain scaling, so we summarized the degree of variability in the volume changes by calculating the standard deviations of these volume changes. A high variability across the brain indicates high warping as different parts of the brain are being stretched to different degrees.

We also examined the strain fields separately for the *x*, *y*, and *z* dimensions. For each voxel, the *x* strain field is a measure of how much each point in the image is being stretched in the *x* dimension and similarly for *y* and *z*. Again, as we were interested in warping rather than whole brain scaling, we calculated the variability (standard deviation) of the stretches in each direction. Field map undistortion is intended to correct the anterior–posterior (*y*) displacement due to magnetic field inhomogeneities. Hence, it might be expected that after undistortion, normalization would have less work to do (require less warping) in the *y* direction specifically.

Results and Discussion

Figure 8 shows the change in the standard deviation of the volume change over voxels required after undistortion. It can be seen that for 12 of the 14 subjects, substantially less warping was required after undistortion. An overall effect was confirmed using a repeated measures ANOVA ($F(1,13) = 28.5$, $P < 0.001$). It is not known why two of the subjects did not follow the trend—a possibility is that if they had unusually shaped brains, the distortions were actually making the brains of those subjects more like the standard template. Alternatively, perhaps movement during the field map acquisition disrupted it, or movement through the session changed the pattern of distortions, although from examination of the movement parameters, these were not subjects that moved more than the others. To further examine any possible relation between movement and the effectiveness of undistortion, two summary measures of the movement were calculated. One was the square root of the sum of the squares of the differences between the maximum and minimum translations over the entire session and the other the square root of the sum of the squares of the differences between the maximum and minimum rotations. We then examined the correlation between each of these measures and the percentage of change in warping after undistortion—but no substantial correlation was found (translations, $r = 0.09$; rotations, $r = 0.03$). It seems that the small to moderate movements observed in our study do not measurably alter the effectiveness of undistortion, although it is almost certain that large movements would do (Hutton *et al.*, 2002).

Figure 9 shows the change in the standard deviation of the strains for the *x*, *y*, and *z* dimensions separately. For the *y* dimension, in all subjects, the degree of variation in the warping applied to voxels across the brain reduced with the application of undistortion. For the *x* and *z* dimensions, most subjects showed a reduced variation in warping, although to a lesser extent than for the *y* dimension. Again, it is not know why a few subjects showed a small increase in the warping required. To assess these changes statistically, a repeated measures ANOVA was performed on the standard deviation of the strains. This had two factors, the first with two levels specifying whether data were undistorted (raw vs undistorted) and the second with three levels (*x*, *y*, or *z*) specifying the dimension. A significant main effect was found of undis-

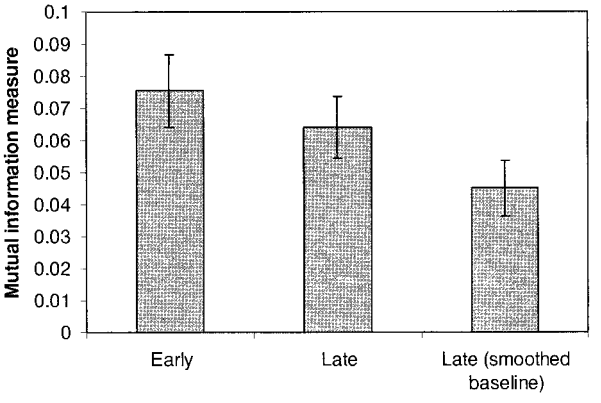

FIG. 6. Change in mutual information with undistortion for an EPI early in the sequence, later in it, and the later image compared to an image not undistorted but smoothed with a Gaussian filter of 4 mm FWHM. The means over nine subjects with 1 SD are shown.

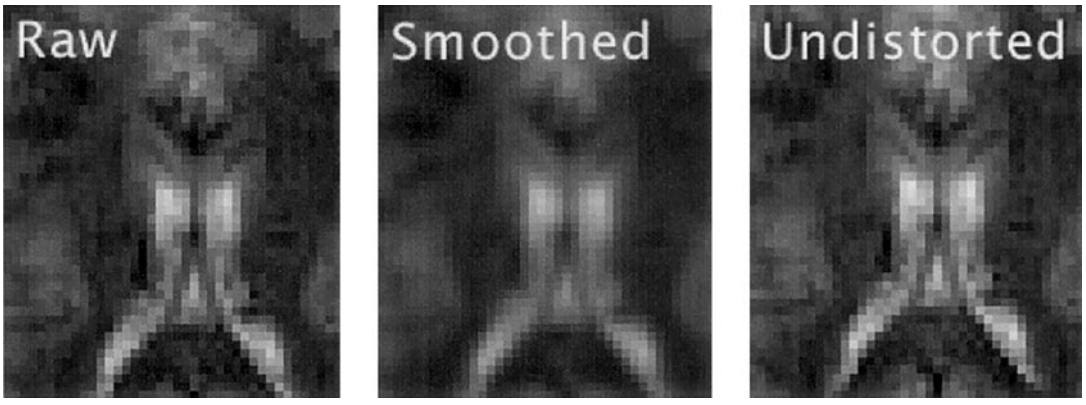

FIG. 7. A comparison of the smoothness of the images following smoothing with a filter of 4 mm FWHM and undistortion. A small portion from an axial slice covering the ventricles is shown on a representative subject.

tortion ($F(1,13) = 59.6$; $P < 0.001$), indicating that overall there was a reduction in the range of strains when undistortion was applied. There was a difference in the total amount of strain in each dimension ($F(2,26) = 13.3$; $P < 0.001$). More interestingly, there was an interaction between the effect of undistortion and the dimension ($F(2,26) = 43.2$; $P < 0.001$), which is probably a reflection of the greater reduction in the ranges of strains in the y dimension than in x and z. To confirm this, we performed a further repeated measures ANOVA on difference scores between the strains, which for each subject was the standard deviation of the strains after undistortion minus the standard deviation of the strains before. In this ANOVA, there was only one factor (dimension) with three levels (x, y, or z) allowing the application of post-hoc pairwise comparisons (LSD). These revealed that the x and the y and the z and the y dimensions showed significantly different effects of undistortion but the x and z dimensions were similar (95% confidence intervals for difference and significance for $x - y$, $z - y$, and $x - z$: 0.12–0.24, $P < 0.001$; 0.12–0.25; $P < 0.001$; 0.049–0.051, NS).

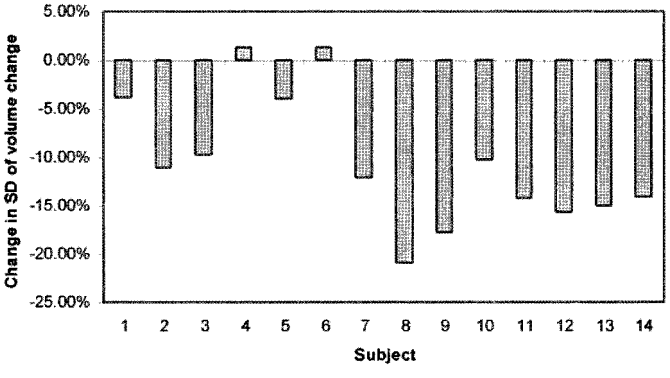

FIG. 8. The amount of warping required to perform normalization if undistortion is done first minus the amount required if normalization is performed without prior undistortion. Warping was measured by the standard deviation of voxel volume change.

EXPERIMENT 3

Aim

In this experiment, we examined the effect of undistortion on the statistical power of a random effects group study. We might expect that the result of nonlinear warping to standard space (spatial normalization) would be better if undistortion is applied first. By reversing the distortions that have affected the image, the images are made more similar to the true shape of the brain (experiment 1). Experiment 2 showed that less warping was then required to transform the subjects' brains to standard space, and as a result, better convergence might be expected during this normalization process. Indeed, it might be impossible for normalization to generate the correct warping fields if undistortion is not applied first. For robustness, SPM 99 (and other normalization algorithms) use a limited set of basis functions, which may not be able to generate the rapid changes in the distortion field required as a result of field inhomogeneities. The application of Bayesian priors also constrains the warping applied to be fairly smooth.

If spatial normalization is more accurate following undistortion, then this might lead to a better overlap of the same brain regions across subjects. When a limited brain region is then activated in an experiment, power to detect the functional activation in an across-subjects random effects analysis might then be increased. In this experiment, we compared the statistical power of a group study using a random effects analysis when the data are undistorted before normalization to when they are not.

Method

Auditory task. On 10 subjects, we acquired two sessions each of 98 volumes each of the 128 × 128 EPIs. Auditory stimuli were presented for 20 s (2 TRs), followed by 10 s (1 TR) of silence. In the current study, we just examined the sound versus silence contrast; subsequent analysis looked for other effects in the data. Hence, while for completeness we describe the stimuli here, their exact form is not relevant to the current study. They were sequences of diotic pure tones

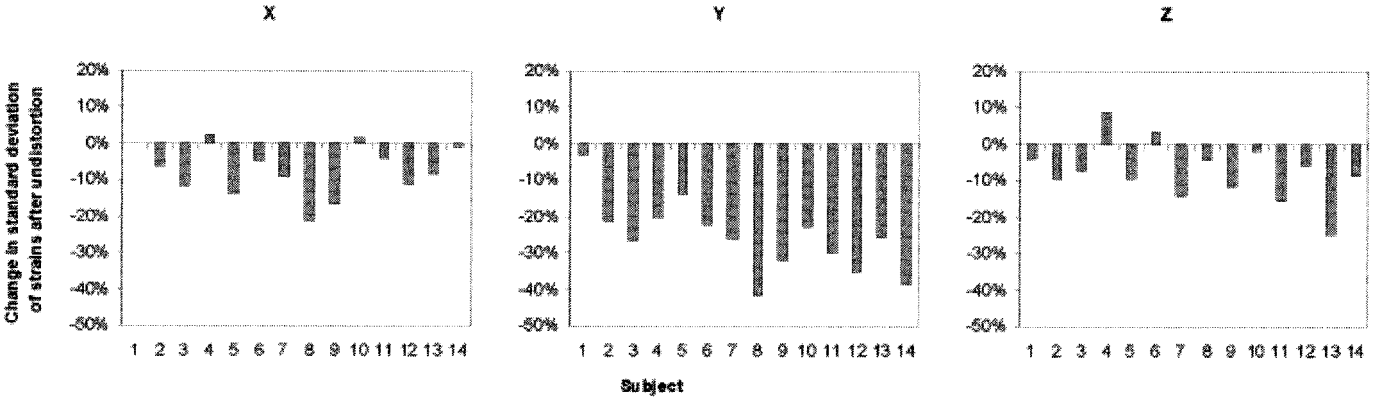

FIG. 9. The amount of warping required to perform normalization if undistortion is done first minus the amount required if normalization is performed without prior undistortion. Warping was measured by the standard deviation of voxel strains in *x*, *y*, and *z* directions.

in a repeating pattern ABA- ABA-, where A represents a tone of one frequency, B a tone of another frequency, and the hyphen a gap. The tones, and the gaps, were all 125 ms in duration. Within a 20-s block, the A and the B tones were fixed in frequency. In half the A tones were higher in frequency, and in the other half the B tones were higher in frequency. The A and the B tones were separated by 1, 3, 5, or 7 semitones. The geometric mean frequency of the tones was chosen randomly from a wide frequency range (0, 3, 6, 9, 12, 15, 18, or 21 semitones above 369 Hz). The stimuli were presented via electrostatic headphones built into a pair of industrial ear defenders using a sound delivery system provided by the MRC Institute of Hearing Research (Nottingham, UK) at 95 dB SPL. The subject wore flat frequency response earplugs (E-A-R Ultratech Hifi) with a quoted attentuation of 22 dB at 1 kHz. It should also be noted for completeness that subjects were performing a task on the sounds, and the presence or the absence of sounds had also been associated with key pressing. As such, the activation in our "auditory" contrast probably also included some motor component.

Motor task. Thirteen subjects took part in an Eriksen flanker task (Eriksen, 1968), while we acquired 202 EPI volumes of the 64 × 64 EPIs. A fixation point was presented for 500 ms, followed by a stimulus for 500 ms. The stimuli comprised a target and two identical flankers. Responses were made to the central stimulus only. The stimulus was replaced by the fixation point for a variable period. On average the intertrial interval was 2.5 s. A response was made with the index finger of either the left or right hand on a button box, depending on the color of the stimulus presented (two alternative choice task with one color mapped to each finger and a third color indicating no-go trials). On 120 of 180 trials an index finger movement was required, about half for each index finger (depending on how many errors a subject made, though this never exceeded 3% of all required responses). As for the purpose of this analysis we were interested only in primary motor activation the event types shall not be discussed further.

Data analysis. The data were initially processed in two alternative ways, with or without undistortion. The auditory data were either (1) motion corrected, normalized to a standard EPI template (supplied with SPM 99), and smoothed (8-mm FWHM) to accommodate individual differences in anatomy or (2) processed in the same way but undistorted after motion correction but before normalization and smoothing. No slice-timing correction was applied as the use of sparse imaging with a long 10-s TR filled with the same stimuli mean that a steady state should be reached before each acquisition and that temporal interpolation is inappropriate. The motor task data were either (1) corrected for slice-time acquisition differences using sinc interpolation, motion corrected, normalized to a standard EPI template (supplied with SPM 99), and smoothed (8-mm FWHM) to accommodate individual differences in anatomy or (2) processed in the same way but undistorted after motion correction but before normalization and smoothing.

Due to the presence of signal dropout in regions where the field is very inhomogenous, cost-function masking was used during normalization. The template to which we were normalizing was acquired on a different machine at lower field strength and did not have the same pattern of dropout. If cost-function masking is not applied, then to fill the space where dropout has occurred to make the object brain match the template, normalization pulls down parts of the frontal and temporal cortex. This strong incorrect warping distorts the normalization process. A better strategy, therefore, is to prevent the matching process in the regions of dropout by providing a cost-function mask to the normalization algorithm. On the mean of each subject under each processing condition, a mask was drawn using Chris Rorden's MRIcro (http://www.psychology.nottingham.ac.uk/staff/cr1/mricro.html) by eye, with areas in the lower slices where there was signal dropout included in the mask. This mask was then smoothed with an 8-mm FWHM filter and then thresholded at 0.001—following a procedure as described in Brett *et al.* (2001), who applied it to the masking of lesions when normalizing damaged brains. This mask, which was 1 in nonmasked areas and 0 within the mask, was supplied to SPM as an "object mask" to be used during normalization.

Following normalization, the data were smoothed using an 8-mm FWHM Gaussian filter. As the auditory data were

acquired using a sparse imaging protocol with a long TR, they were modeled using a boxcar function that was 1 for all scans that were preceded by sound and 0 for all scans preceded by silence. The motor task data were acquired using a more rapid TR and hence were modeled using an event-related design. A single column in the design was generated from recorded response times modeled button presses by the left hand and another column for button presses by the right hand. No-go trials were not modeled. To both the auditory and motor models, the six columns of movement parameters generated during realignment were added in an attempt to remove movement-correlated artifacts.

Results and Discussion

Figure 10 shows activation in the auditory and motor tasks as analyzed with and without undistortion. These analyses were conducted using a random effects model in SPM 99 in which the contrasts for each voxel in each subject were entered into a second level analysis, and then a one-sample t test was performed. The activation, shown superimposed on the T1 template, is shown in the coronal slices crossing auditory cortex ($y = 12$) for the auditory task and crossing primary motor cortex ($y = -14$) for the motor tasks. Activations that satisfied the FDR-corrected threshold of $P < 0.02$ are shown (Genovese, Lazar and Nichols, submitted for publication; http://www.sph.umich.edu/~nichols/FDR). The peak T values across the whole volume are shown in the top left corner of each image. In the auditory task, if undistortion is applied power is increased, as is reflected by the increase in peak T and the greater extent of activation. In the motor tasks, no significant activation was seen without the application of undistortion, but after its application substantial activation was seen in primary motor areas.

These measurements show that undistorting the EPIs before normalization increases the statistical power of group studies.

EXPERIMENT 4

Method

In experiment 4, we measured directly whether undistorting EPIs improves the overlap in normalized activations from different subjects. We used data from the 13 subjects who had done the motor task analyzed in experiment 3. For both the undistorted and pre-undistorted data, *for each subject*, we used SPM to find voxels that were activated at a level of 0.001 uncorrected in response to left and right movements.

We first compared whether the application of undistortion made a difference in the total number of voxels activated in each subject. Then, we looked at the degree of overlap between different subjects. We did this in two ways. First, for visualization, we calculated the overlap maps of the number of subjects that showed activation at each voxel. Second, for every possible pair of subjects, we calculated the overlap between their two activation patterns using a measure of *the number of voxels overlapping* divided by the *lowest number of activated voxels across the two individual maps*. In this way, for each pair of regions, the degree of overlap was summarized as a number between 0 (no overlap) and 1 (one a subset of the other).

Results and Discussion

Figure 11 shows the number of subjects showing activation at each point in space, for left and right movements, and for data before or after undistortion. Also shown (white dotted region) is the primary motor cortex, Brodmann Area 4. It can be seen that when the EPIs were undistorted, the activation observed in different subjects showed much greater overlap. Furthermore, it can be seen that when undistortion was performed, the activation was accurately localized to the expected anatomical region.

There was no increase in the total number of active voxels within subjects (left press, six subjects showed more, six showed fewer, one no change, $t(12) = 0.865$; right press, seven showed more, six showed fewer, $t(12) = 0.315$). However, the numerical analysis confirmed that the degree of overlap changed substantially. A sign test was used as the data were not normally distributed (left press, 23% of pairwise comparisons showed less overlap, 67% more overlap, $P < 0.01$; right press, 28% less overlap, 62% more overlap, $P < 0.01$). This experiment demonstrates that when EPIs were undistorted prior to normalization, activation in different subjects was in better registration, which then led to an increase in statistical power, as observed in experiment 3.

EXPERIMENT 5

Aim and Method

The aim of the final investigation was to ascertain which brain regions are most affected by the distortion that can be corrected using our procedure. To do this, we compared the warping if only normalization was applied to that if undistortion and then normalization were applied. We present a novel way of visualizing the two patterns of distortions in a typical subject. We also performed a quantitative evaluation of the distortions over a group of subjects. Two possible effects on group studies of incorrect spatial normalization are considered. It might be that if undistortion is not performed before normalization, then in all of the subjects, the signal from a particular brain region is displaced by a similar spatial vector. In this case, activation in that region will still be seen in a group analysis, but in the wrong place. Alternatively, the signal from different subjects may be displaced in different directions and/or by different distances, in which case the activation in that region will no longer overlap across subjects and there will be reduced power. We evaluated these two possible errors separately and present maps of the magnitude of each of the problems in a number of slices.

The two measures were calculated as follows. Let the displacement vector at voxel r, when undistortion and normalization are applied for subject s, be $D_{\mathrm{UN}}(s, r)$ and the vector field when normalization is applied without undistortion $D_{\mathrm{N}}(s, r)$. Then the difference in displacement when undistortion isn't applied is simply

$$E(s, r) = D_{\mathrm{N}}(s, r) - D_{\mathrm{UN}}(s, r) .$$

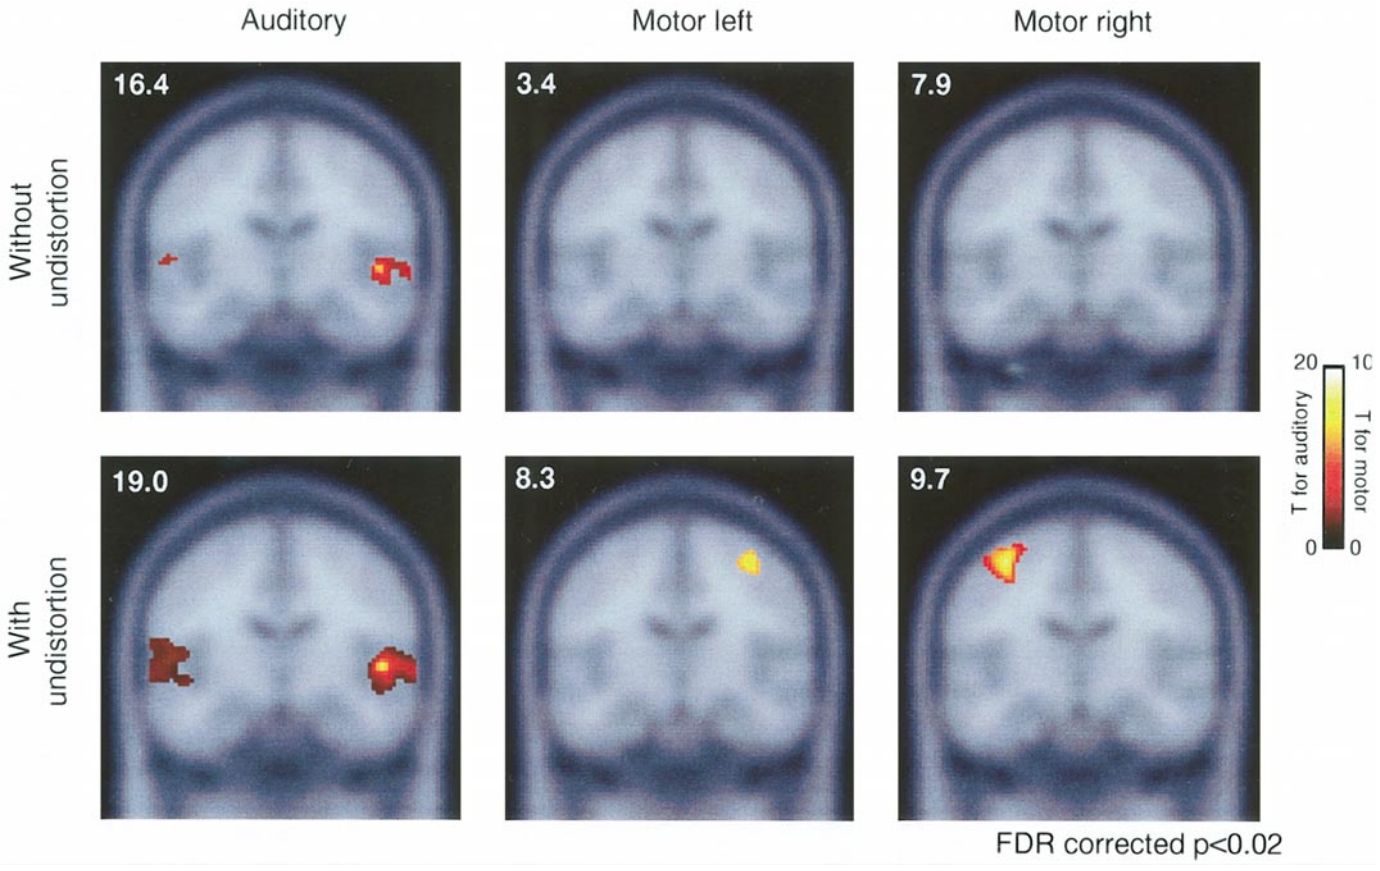

FIG. 10. Activation detected in random effects analysis of 9 subjects (auditory task) and 13 subjects (motor task) with and without undistortion.

Also, let the mean of this vector over subjects be

$$\bar{E}(r) = \frac{1}{S} \sum_{s=1}^{S} E(s, r),$$

where S is the total number of subjects.

Distortions can then be summarized in two ways. If for a particular voxel in the brain, there is on average an offset of the source of a particular voxel, but this offset is similar for all subjects, then this will lead to mislocalization, but not a reduction in power. We can examine the magnitude of this effect by calculating

$$m(r) = |\bar{E}(r)|.$$

If, on the other hand, the difference vector is different for each subject, then we will be sampling from different parts of the brain in different subjects, and hence we are likely to see a reduction in power for that voxel. This can be summarized by examining the standard deviation of the displacement in each subject relative to the mean displacement over all subjects.

$$\sigma(r) = \sqrt{\frac{\sum_{s=1}^{S} |E(s, r) - \bar{E}(r)|^2}{S}}.$$

To further help us visualize how the deformations differ as a result of undistortion being applied or not, we also generated a volume with a simple checkerboard across all sagittal planes, comprising 8-mm squares alternating between 1 and 0. For a single representative subject, this battenberg cake[1] image was then warped in the same way as EPIs would be. It was then masked for areas outside the template brain using the same mask as generated for masking the calculations described above.

Results and Discussion

The top panel of Fig. 12 shows the measure of the degree of error in localization that might result at each voxel when undistortion is not performed. The bottom panel shows at

[1] A checkered two-colored sponge cake, popular in the United Kingdom, named in honor of the marriage of Queen Victoria's granddaughter to Prince Louis of Battenberg, 1884.

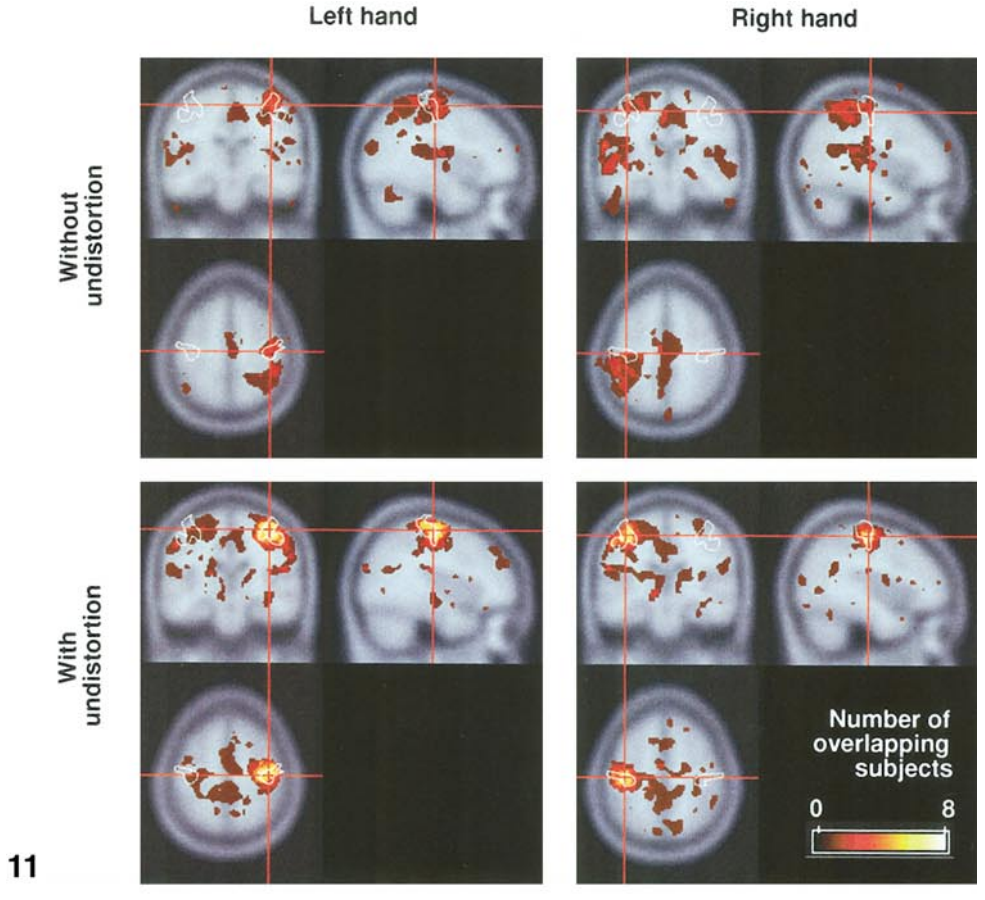

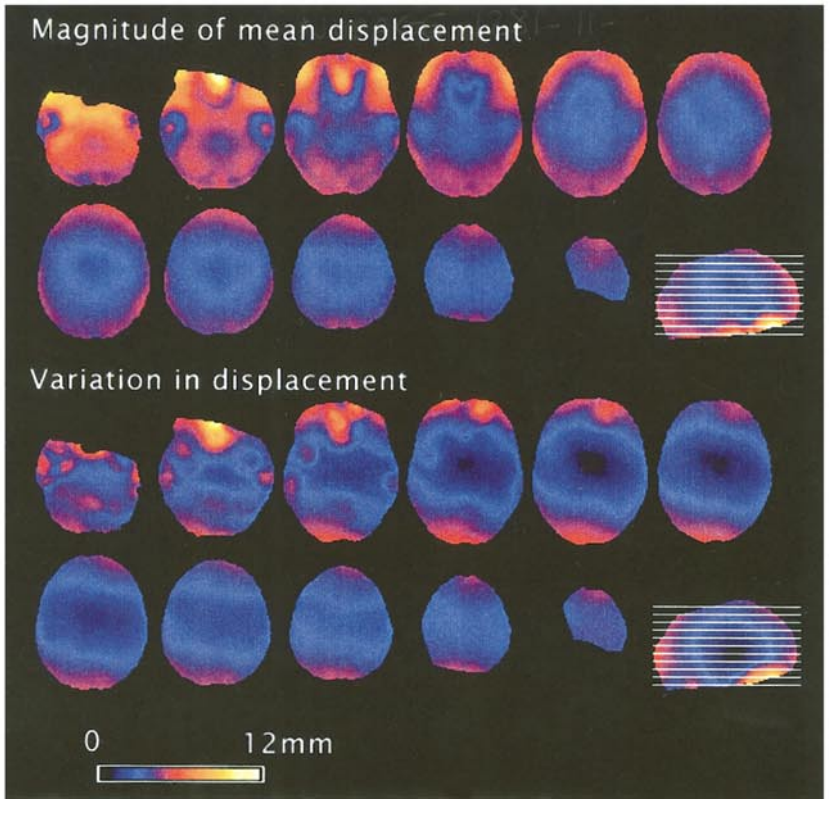

FIG. 11. Overlap between activation of different subjects with and without undistortion.

FIG. 12. Two measures of the difference between applying normalization alone compared to applying both undistortion and normalization. The top panel shows a measure of possible mislocalization of activation, described in detail in the text. The bottom panel shows a measure of variation in spatial displacement across subjects. Larger variation across subjects would be expected to lead to to a reduction in the power of group studies.

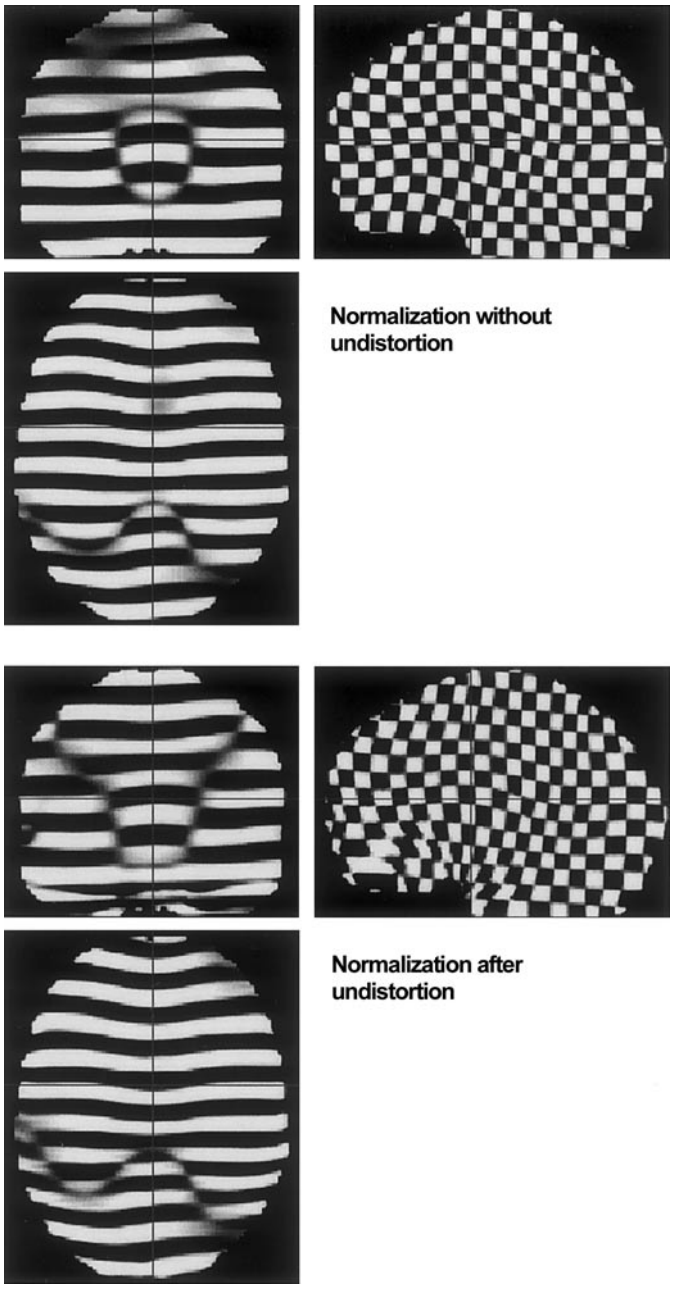

FIG. 13. Distortions applied by normalization alone (top) and undistortion followed by normalization (bottom).

each voxel the measure of the reduction in power that might be expected.

As might be expected, the largest distortions are in the inferior frontal region, with mean shifts of around 12 mm and a standard deviation of around 10 mm. However, there are also higher values at the front and back of the brain, with mean displacement differences of around 4 mm and root mean square variation of around 5 mm.

Figure 13 illustrates the warping fields for the two different procedures. As would be expected, the largest differences are near the frontal and temporal poles. However, note subtle differences elsewhere—for example, the superior and anterior frontal lobe appears to be stretched upward when undistortion is not done, but has been pulled forward in the undistorted normalized image.

CONCLUSIONS

We have presented a procedure for undistorting EPIs. We have shown that the prior application of undistortion yields images that are more similar in shape to the true shape of the brain, as measured using a non-EPI sequence (experiment 1). We have also shown that applying undistortion generally makes the EPIs from individual subjects more similar to the standard EPI template, requiring less warping to match it, particularly in the phase-encode direction (experiment 2). Furthermore, we have shown that it can increase the statistical power of group studies using a random effects analysis (experiment 3) and that the origin of this increase in power is a greater overlap between activations of different subjects (experiment 4). Finally, in experiment 5, we measured the areas of the brain for which undistortion might have the most effect.

In the current study, we collected a single field map at the end of the session. There are some advantages in doing this over collecting EPI phase maps throughout. The field map collected has a higher resolution and a higher signal-to-noise ratio. Second, and importantly, as we collect the field information in undistorted space, the undistortion process becomes a "pull" problem, rather than a "push" problem. In a pull problem, we can step through each of the voxels in the destination space, look at the undistortion field to see where the signal will have been shifted, and then sample from this location in the EPIs. If the undistortion field is also distorted, then this procedure becomes much more complicated (a push problem), requiring inversion of the field or some kind of interpolation to clean up the images. We had to solve this inverse problem in the generation of our distorted field map magnitude image (see Figs. 2b and 2c), but this image was only used in the process of registering our field maps with EPIs. Hence, the precise form of the approximation to the inverse problem will not have had a direct effect on our EPIs, as would be the case if field maps were acquired in distorted space. A final advantage of not collecting field maps in distorted space is that if there is compressive distortion, then there will be the least information on how the data should be undistorted in a region where we most need it. The more compressed a part of a field map is, the larger the distance in real space between adjacent samples of the field.

A disadvantage of collecting a single field map is that, as mentioned in the discussion of experiment 1, this procedure cannot correct for any changes in the pattern of distortion due to movement over the course of a session. The specific test in experiment 1 showed that, for our data, there was little difference between the improvement due to undistortion of an EPI from the beginning or from the end of a session. However, it might be that a combination of field map undistortion and Andersson *et al.*'s (2001) algorithm for the post-hoc removal of distortion by time interactions would lead to a additional improvement in power. The most pronounced im-

provement would be expected in the inferior frontal regions that Hutton *et al.* (2002) have recently shown to be substantially affected by changing patterns of distortion with subject motion.

Another possible way of tackling the effect of movement by distortion interactions would be to try to derive the pattern of magnetic susceptibility in a particular subject's head. This could be done either by just using the information contained in a structural image or by combining information from both the single field map and the structural image. Then, the distortions to the field that would be generated by the head in any position could be calculated using algorithms for deriving fields from susceptibility maps, such as those by Balac and Caloz (2000). These could then be used for undistortion. However, it is clear that this strategy would require a formidable amount of effort to develop and computing power to implement.

Where the EPI image is compressively distorted, the intensity of the image will increase, as a greater volume of brain will be contributing to each voxel. Conversely, where it is expansively distorted, a smaller volume will contribute and intensity will decrease. Jezzard and Balaban (1995) suggested that to a first approximation, we might correct for this by dividing by the differential of the displacement in the phase-encode direction. However, recent studies by Jenkinson (2001) have suggested that the derivative of the field map is too noisy to perform the correction accurately. We did not use the correction here, but future studies might investigate the effect of including this correction on image registration. However, for the purpose of applying statistics, as we are really interested in changes over time within a voxel—the comparison of changes that are correlated with an experimental manipulation with random fluctuation over time—it would not be expected to have any effect on final measurements. Alternatives to the pixel shift method described by Jezzard and Balaban (1995) have been proposed by Munger *et al.* (2000). Theoretically, these methods allow us to correct for more subtle effects of the field distortions on the point spread functions of each voxel and would simultaneously correct both distortion and intensity. However, evaluations by Jenkinson (2001) have suggested that these perform no better than the standard method.

Future work might further optimize some of the parameters used in our study, such as the size of the smoothing kernel for the field maps or the method of interpolation used. It has been suggested that smoothing field maps too much removes abrupt changes in field information (Hutton *et al.*, 2002) and smaller kernels might still yield sufficient signal-to-noise. Ideally, the resampling stages applied at motion correction, undistortion, and normalization would all be combined into a single resample. It would also be interesting to compare the results of group studies using the procedure we applied, normalizing directly from the mean EPI to the EPI template, with the procedure where the EPIs are first coregistered to the high-resolution structural image, and then this used to estimate normalization parameters. However, as there will still be regions of dropout in the EPIs, it might be necessary to apply cost-function masking when performing rigid body coregistration between the echo-planar and structural images. This would make an interesting future study.

In summary, the application of field map undistortion has been validated in a number of ways. Its application has been shown to improve the localization of activations and increase the statistical power of group studies.

ACKNOWLEDGMENTS

The authors thank two anonymous reviewers for their helpful comments; those at the CBU and in the wider brain imaging community in Cambridge who have provided helpful comments and suggestions throughout this project; the radiographers at the Wolfson Brain Imaging Centre for data acquisition; and Nikos Papadakis, Emma Williams, Kay Martin, and Adrian Carpenter for their help in setting up the field-map sequences. Katja Osswald is funded by Medical Research Council (UK) and Gottlieb Daimler- und Karl Benz-Stiftung; the other authors by the Medical Research Council (UK).

REFERENCES

Andersson, J. L. R., Hutton, C., Ashburner, J., Turner, R., and Friston, K. J. 2001. Modelling Geometric Deformations in EPI Time Series. *NeuroImage* **13:** 903–919.

Ashburner, J., and Friston, K. J. 1999. Nonlinear spatial normalization using basis functions. *Hum. Brain Mapp.* **7:** 254–256.

Ashburner, J., Neelin, P., Collins, D. L., Evans, A. C., and Friston, K. J. 1997. Incorporating prior knowledge into image registration. *NeuroImage* **6:** 344–352.

Balac, S., and Caloz, G. 2000. Mathematical modeling and numerical simulation of magnetic susceptibility artifacts in magnetic resonance imaging. *Comput. Methods Biomech. Biomed. Eng.* **3:** 335–349.

Brett, M., Leff, A. P., Rorden, C., and Ashburner, J. 2001. Spatial normalization of brain images with focal lesions using cost function masking. *NeuroImage* **14:** 486–500.

Chen, N. K., and Wyrwicz, A. M. 1999a. Correction for EPI distortions using multi-echo gradient-echo imaging. *Magn. Reson. Med.* **41:** 1206–1213.

Chen, N. K., and Wyrwicz, A. M. 1999b. Removal of intravoxel dephasing artifact in gradient-echo images using a field-map based RF refocusing technique. *Magn. Reson. Med.* **42:** 807–812.

Collignon, A., Maes, F., Delaere, D., Vandermeulen, D., Suetens, P., and Marchal, G. 1995. Automated multi-modality image registration based on information theory. In *The Proceedings of Information Processing in Medical Imaging* (Y. Bizais, C. Barillot, and R. Di Paola, Eds.), pp. 263–274. Kluwer Academic, Dordrecht.

Cusack, R., and Papadakis, N. 2002. New robust 3D phase unwrapping algorithm: Application to magnetic field mapping and undistorting echo-planar images. *NeuroImage,* **16:** 754–764.

Deichmann, R., Josephs, O., Hutton, C., Corfield, D. R., and Turner, R. 2002. Compensation of susceptibility-induced BOLD sensitivity losses in echo-planar fMRI imaging. *NeuroImage* **15:** 120–135.

Devlin, J. T., Russell, R. P., Davis, M. H., Price, C. J., Wilson, J., Moss, H. E., Matthews, P. M., and Tyler, L. K. 2000. Susceptibility-induced loss of signal: Comparing PET and fMRI on a semantic task. *NeuroImage* **11:** 589–600.

Eriksen, C. W. 1995. The flanker task and response competition: A useful tool for investigating a variety of cognitive problems. *Vis. Cogn.* **2:** 101–118.

Friston, K. J., Ashburner, J., Frith, C. D., Poline, J. B., Heather, J. D., and Frackowiak, R. S. J. 1995. Spatial registration and normalization of images. *Hum. Brain Mapp.* **2:** 165–189.

Hall, D. A., Haggard, M. P., Akeroyd, M. A., Palmer, A. R., Summerfield, A. Q., Elliott, M. R, Gurney, E. M., and Bowtell, R. W. 1999. 'Sparse' temporal sampling in auditory fMRI. *Hum. Brain Mapp.* **7:** 213–223.

Hennel, F. 1999. Two-dimensional deghosting for EPI. *MAGMA* **9:** 134–137.

Jenkinson, M. 2001. Improved unwarping of EPI images using regularised B0 maps. *NeuroImage* **13:** S165.

Jezzard, P., and Balaban, R. S. 1995. Correction for geometric distortion in echo planar images from B_0 field variations. *Magn. Reson. Med.* **34:** 65–73.

Jezzard, P., and Clare, S. 1999. Sources of distortion in functional MRI data. *Hum. Brain Mapp.* **8:** 80–85.

Kybic, J., Thevenaz, P., Nirkko, A., and Unser, M. 2000. Unwarping of unidirectionally distorted EPI images. *IEEE Trans. Med. Imaging* **19:** 80–93.

Munger, P., Crelier, G. R., and Peters, T. M. 2000. An inverse problem approach to the correction of distortion in EPI images. *IEEE Trans. Med. Imaging* **19:** 681–689.

Ojemann, J. G., Akbudak, E., Snyder, A. Z., McKinstry, R. C., Raichle, M. E., and Conturo, T. E. 1997. Anatomic localization and quantitative analysis of gradient refocused echo-planar fMRI susceptibility artifacts. *Neuroimage* **6:** 156–167.

Reber, P. J., Wong, E. C., Buxton, R. B., and Frank, L. R. 1998. Correction of off resonance-related distortion in echo-planar imaging using EPI-based field maps. *Magn. Reson. Med.* **39:** 328–330.

Weisskoff, R. M., and Kiihne S. 1992. MRI susceptometry: Image-based measurement of absolute susceptibility of MR contrast agents and human blood. *Magn. Reson. Med.* **24:** 375–383.