NeuroImage **16,** 754–764 (2002)
doi:10.1006/nimg.2002.1092

New Robust 3-D Phase Unwrapping Algorithms: Application to Magnetic Field Mapping and Undistorting Echoplanar Images[1]

R. Cusack* and N. Papadakis†,‡

**MRC Cognition and Brain Sciences Unit, Cambridge; †Wolfson Brain Imaging Unit, Cambridge; and ‡Department of Psychiatry, University of Sheffield, United Kingdom*

Received September 7, 2001

The phase, as well as the magnitude, of MRI images can carry useful information. It may be used to encode flow or temperature, or to map the magnetic field for the undistorting of EPIs and automated shimming. In all cases, we measure the extra spin given to nuclei. Unfortunately, we can only measure the final phase of the spins: the rotation is wrapped into the range $[-\pi, +\pi]$, and to obtain a measure of the parameter of interest the missing multiples of 2π must be replaced—a process known as phase unwrapping. While simple in principle, standard phase unwrapping algorithms fail catastrophically in the presence of even small amounts of noise. Here we present a new algorithm for robust three-dimensional phase unwrapping, in which unwrapping is guided, so that it initially works on less noisy regions. We test the algorithm on simulated phase data, and on maps of magnetic field, which were then used to successfully undistort EPI images. The unwrapping algorithm could be directly applied to other kinds of phase data. © 2002 Elsevier Science (USA)

INTRODUCTION

Ideally, the magnetic field throughout the bore of an MR scanner would be homogenous in the absence of any applied gradients. Unfortunately, different materials, such as tissue, bone, and air, act to strengthen or weaken magnetic fields to different extents—a property characterised by their magnetic susceptibility. When a head is placed in the scanner, inhomogeneities are induced in the field, which can lead to loss of signal ("dropout") and image distortion. While shimming, the application of small extra gradients, can cancel out some of these inhomogeneities others remain. To counteract the problem of distortion, Jezzard and Balaban (1995) suggested measurement of the magnetic field strength across the head, which may then be used for undistorting MR images. This paper describes a new algorithm for the processing of a phase-contrast MR acquisition that can be used to measure the field.

Information in the phase, rather than the magnitude, of MR images can allow us to measure many interesting parameters, such as flow rate or magnetic field strength. Usually, two datasets are acquired, which differ only in the degree to which the phase is influenced by the parameter of interest. One dataset is used as a reference, and subtracted from the phase of the other. In this way, uninteresting phase components that are common to both images, such as the transmitter or receiver characteristics, or those of the digital filtering, are removed. Only the effect of the parameter remains.

The phase is usually proportional to the parameter of interest. However, it will have been wrapped into the range $[-\pi, +\pi]$. If the parameter we wish to measure gives phase differences that are small enough to fall into this range, then the measured phase can be used directly for parameter estimation. If, on the other hand, the magnitude of the phase difference is large enough to be outside this range, then to recover the true phase, we must restore the missing multiples of 2π. This process known as phase unwrapping. The phase of voxels in a volume may be unwrapped relative to each other, provided that the underlying phase is relatively slowly changing over space. So, for example, in Fig. 1b, working from left to right, we can identify that the change from $+\pi$ to $-\pi$ at *A* and *B* indicates that phase wrapping has occurred in the negative direction, and to unwrap we need to add 2π to subsequent values. Conversely, the change from close to $-\pi$ to $+\pi$ at position *C* and *D* indicates that wrapping has occurred in the positive direction, and we need to subtract -2π from subsequent values. Note that here, we are taking a single reference value (on the far left) and then unwrapping relative to this. We discuss later the problem of calculating the appropriate offset for the whole map.

[1] Part of this work was presented at the Applied Optics and Opto-Electronics Conference, University of Loughborough, UK, Sept 2000; and at Human Brain Mapping, Brighton, UK, June 2001.

1053-8119/02 $35.00

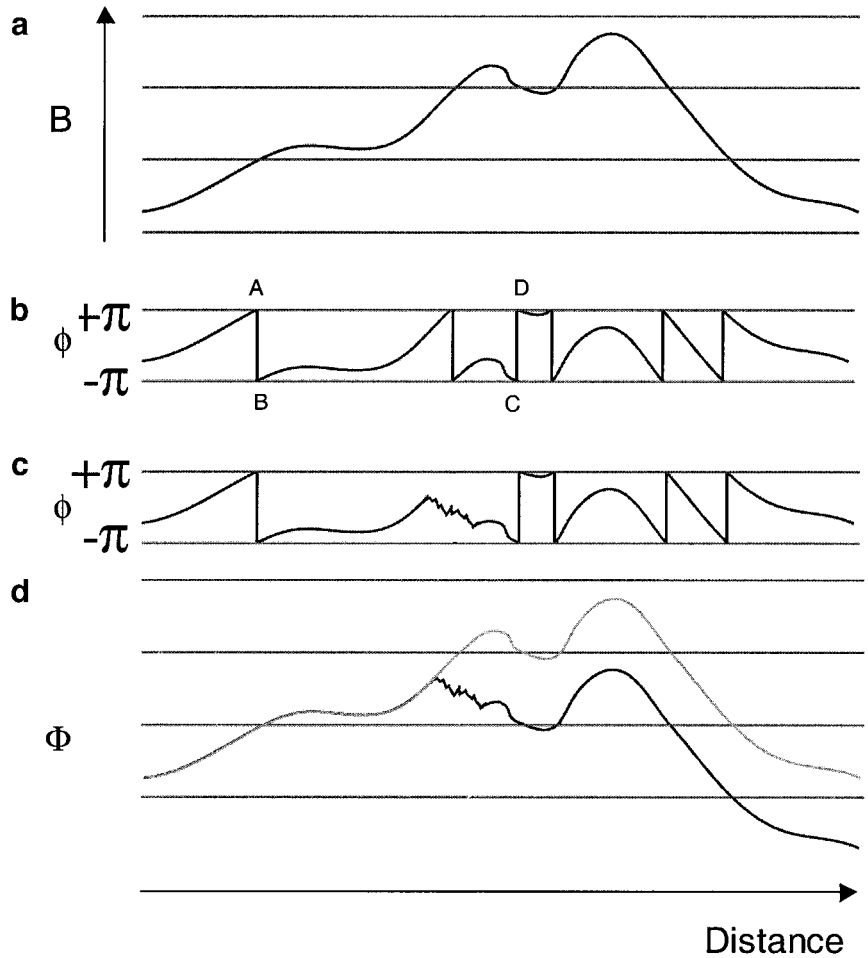

FIG. 1. The underlying B field we are trying to measure is shown in (a). The wrapped phase is shown in (b), and with the addition of some noise in (c). Unfortunately, as shown in (d) noise catastrophically interferes with unwrapping: the derived field (dark curve) differs over much of space from the true one (light curve).

Let us consider the limitations on the unwrapping procedure in the absence of noise. In theory, provided the underlying field is spatially continuous, the only discontinuities that can appear in the wrapped field would be those from the wrapping process itself, and unwrapping is straightforward. In practice, we deal with sampled fields, which yield a set of discrete measurements. To unwrap, we calculate the magnitude of the phase jump from one sample to the next adjacent one. If the magnitude of the phase difference is greater than π, phase wrapping has occurred—otherwise it has not. In the absence of noise, we can correctly unwrap the field provided there are no discontinuities between adjacent voxels in the underlying field that are greater than π.

Whilst in principle phase unwrapping is simple, in practice it is rather more difficult to implement. The problem is that it is highly sensitive to errors in regions with low signal-to-noise. These errors can propagate into areas with good signal, and lead to a catastrophic failure of the unwrapping process. So, for example, if just a few voxels in a map are noisy (Fig. 1c) and one extra or fewer 2π phase discontinuity is introduced, then all voxels downstream of those in unwrapping will have a large error of $\pm 2\pi$ (see Fig. 1d). Even small amounts of noise lead to failure of unwrapping over large portions. With one-dimensional data, it is very difficult to recover this information.

It is simple to extend phase unwrapping to two or three dimensions. Again, we first choose a reference pixel and then unwrap the others relative to this. However, note that as we unwrap from one pixel to another, there are now many different paths we could follow (see Fig. 2a). Each different path can be considered as a one-dimensional profile. Given that our unwrapped map represents a physical quantity that can only have one value at each position, the unwrapped value of a pixel cannot be dependent on

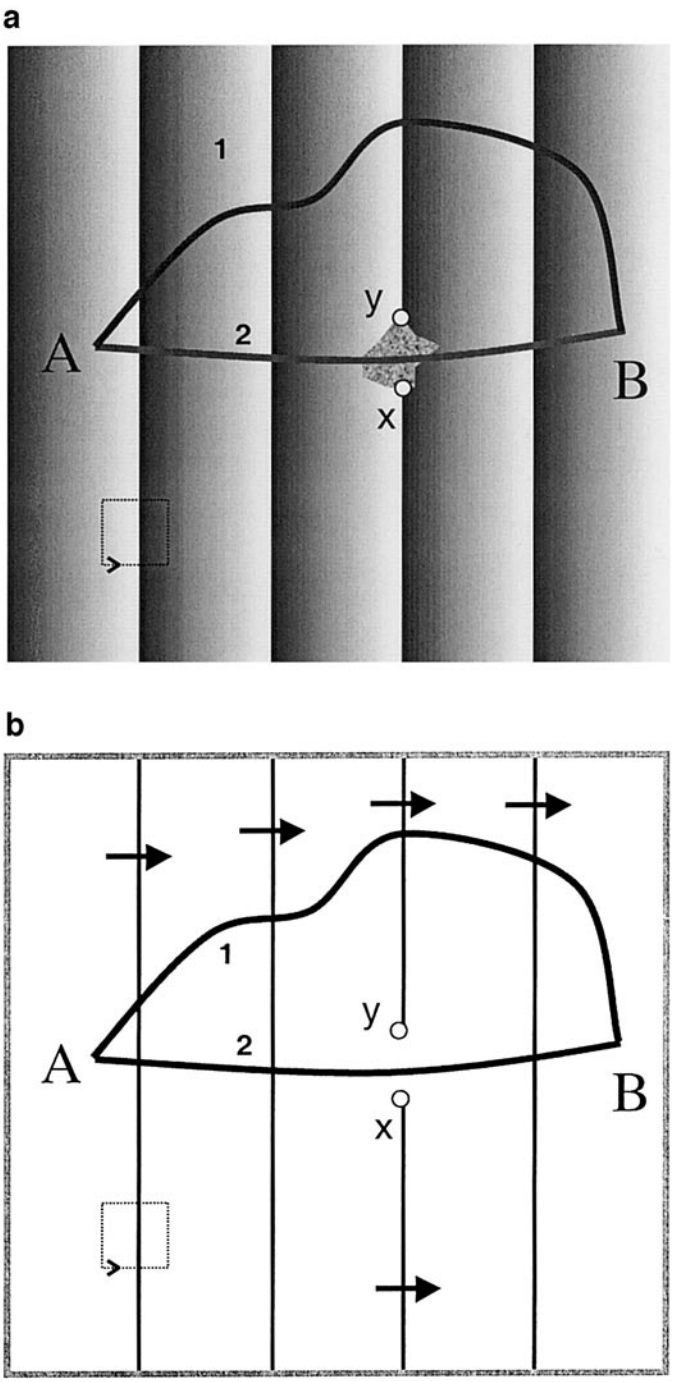

FIG. 2. The upper panel represents the phase as a function of position in two-dimensional space. The darkest shading represents $-\pi$ and the lightest shading $+\pi$. The speckled area represents a patch of noise. The sharp boundaries between dark and light are phase discontinuities: these are shown as lines in the lower panel, with the arrows indicating the direction in which 2π must be added. When unwrapping from one pixel (A) to another (B), there are many paths that can be followed (e.g., 1 or 2). Only one of these (2) passes through the noisy area. Unwrapping around a closed loop (dotted square) should give a sum of zero, but where discontinuities end (x, y) nonzero sums are found. These give a signature of noise in a region, which we use to guide unwrapping.

the unwrapping path we take to get there. It is this redundancy that helps us robustly unwrap multi-dimensional maps even when they are noisy. As in one dimension, catastrophic errors in unwrapping can propagate through the map where a 2π phase discontinuity is missed, or an extra one crossed. We can develop a robust algorithm provided that the noise, in which 2π phase discontinuities are missed, or extra ones introduced, is confined to particular regions of the image, and we unwrap these regions last. If we can deal with localized areas of noise, the process will be very much more robust. Of course, where an error spans large parts of the phase map, then information may be lost and irrecoverable by any algorithm.

Here we try to ascertain the amount of noise as function of space and unwrap the less noisy parts first. We investigate different ways of estimating noise. In one approach, we have used two measures derived from the magnitude of the signal in the field maps. A second approach is based on algorithms developed for two-dimensional phase unwrapping in optics and synthetic aperture radar. This approach uses the property that as unwrapping is path independent in a noise-free map, the sum of signed phase wraps around any closed loop should be zero (e.g., Dirac, 1931; for introduction see http://www.physicstoday.org/pt/vol-54/iss-8/captions/p27box1.html). Conversely, where different paths between the same pair of voxels give different unwrapped values, there will be loops with a nonzero sum, and by identifying these we can find noisy regions of the map. When working in two dimensions with data sampled in a rectangular lattice, the signed sum around any loop is equal to the total of the signed sum around the 2×2 pixel rectangles it contains. If the sum around the larger loop is non-zero, then at least one of the sums around the 2×2 pixel rectangles must be nonzero (Bone, 1991; Huntley, 1989). Hence, if there is noise in a region that destroys path independence, there will always be nonzero sums around at least one of the 2×2 pixel rectangles in that region. We refer to these points around which there is a nonzero sum as poles. When working in three dimensions with data sampled in a cuboid lattice, any loop can be calculated by summing 2×2 voxel rectangles oriented in three orthogonal planes that make it up. Using logic similar to that in two dimensions, the presence of poles that lead to a loss of path independence in a region will always be indicated by one or more nonzero sums around the 2×2 voxel rectangles, and we can use these to identify noisy regions. We first identify all poles in a map by calculating the sum of the phase unwraps around all possible 2×2 loops in all directions. Then, we generate a diffuse field around each, and sum these to form an estimator of the noise at each voxel in the image. This is then used to guide unwrapping.

METHOD

Unwrapping Procedure

Our approach was to identify less noisy regions of the phase image and unwrap these first. To do this, we calculated an estimate of the amount of noise at each voxel, with higher values corresponding to more noise (hereafter called the noise estimation field), and unwrapped voxels with lower values first. We compared a simple unguided unwrapping strategy to our new algorithm with different types of estimators of noise. The first noise estimators were based on the magnitude of the signal in the phase map. For simulated data, this was the magnitude of the phase difference map, negated. For the fMRI data, this was either the magnitude of the high TE field map, negated, or the minimum magnitude of the voxel across the two field maps, negated. A second type of noise estimator was derived from the phase poles. First, all of the poles in the image were identified. For each pole, for every voxel in the map, a noise estimation field value is calculated. This is high close to the pole, and lower further away. The noise fields for all of the individual poles are then added together to form the single total noise field, which is used to direct unwrapping.

For each iteration, only voxels lower than a particular threshold (the noise-field threshold) are unwrapped. On the first iteration, this threshold is low; for subsequent iterations it is gradually increased. To unwrap, first a starting seed and initial pole-field threshold were specified. All of the voxels adjacent to the seed were examined. For each, if its noise field value is lower than the current noise-field threshold, the voxel is unwrapped: if it is higher, unwrapping is deferred. In this way, the regions far away from poles (with a low noise-field value) are unwrapped first; and regions close to poles done later. When all of the voxels lower than the noise-field threshold have been unwrapped, unwrapping becomes braver (the noise-field threshold is increased) and this process repeats for another iteration.

Summary of "Standard" Algorithm

The simple unwrapping algorithm used a recursive "flood fill" technique. The stages are as follows:

(1) Set current voxel C to be a seed voxel S

(2) For each neighbor that has not yet been visited, unwrap. The neighbors of the voxel (x, y, z) were checked in the following order: $(x, y, z + 1)$, $(x, y, z - 1)$, $(x, y + 1, z)$, $(x, y - 1, z)$, $(x + 1, y, z)$, $(x - 1, y, z)$. This stage was repeated, until the whole map was unwrapped. A first-in first-out stack was used so that all of the neighbors of a voxel are checked before the neighbors of these neighbors.

Summary of New Algorithm

In stage one, the noise field that is to be used to guide the unwrapping is derived either from the magnitude of the field map or using the pole field. In stage two, the map is unwrapped, beginning in low-noise areas.

Stage 1: Calculate noise estimation field. Simulated data: Either (a) noise estimation field was magnitude of phase-map, negated so that negative values correspond to lower noise regions; or (b) identify poles in the map by calculating unwrap total around loops around 2×2 voxel squares in all three orthogonal planes, and flagging nonzero values as poles and then smoothing this in three dimensions. fMRI data: Either (a) noise estimation field to be the magnitude of long TE field map, negated; (b) noise field was minimum magnitude for a voxel in either of the acquisitions; or (c) generate pole field as described for simulated data above.

Stage 2: Iterative unwrapping of low-noise areas first. A similar procedure to the standard algorithm was used, except that unwrapping was guided by the noise estimation field.

(1) Set current voxel C to be a seed voxel S, and set noise field threshold
(2) For each neighbor of the current voxel C:
  (i) If it has a noise field value less than the current threshold, unwrap and repeat stage (2) for this new voxel.
  (ii) If it has a noise field greater than the current threshold, don't unwrap now, and only repeat stage (2) for this voxel when threshold is high enough.
(3) When all voxels encountered below current threshold are completed, increase noise field threshold, and continue with new voxels that may now be unwrapped.

Specific Details

Stage 1 with Pole Field

Pole positioning. Let us suppose we unwrap around a small loop in the xy plane with coordinates $(x, y, z) \rightarrow (x, y + 1, z) \rightarrow (x + 1, y + 1, z) \rightarrow (x + 1, y, z) \rightarrow (x, y, z)$. If we find a pole in this small square, our algorithm puts a source of the pole field at the position $(x + 0.5, y + 0.5, z)$. Similar positioning is used for loops in the xz and yz planes.

Field shape. A volume specifying the positions of all of the poles was generated. This contained one at the center coordinates of a pole and zero elsewhere. This volume was then smoothed to form the pole field. A pseudo-Gaussian smoothing filter was chosen as this is approximately separable (i.e., convolutions in 3 orthogonal planes may be performed, rather than a more time consuming 3-dimensional convolution). A kernel of five points (0.1, 0.2, 0.4, 0.2, 0.1) was applied repeatedly on

the x dimension, then on the y, then on the z. This kernel was used as it only requires divisions by whole multiples of two (and rescaling) and hence is fast. The degree of smoothing was optimised by parametrically varying the number of iterations between 0 and 5.

Unwrapping procedure. Unwrapping starts at a seed voxel, which is arbitrarily assumed to have no extra multiple of 2π. This seed voxel is chosen to be one with a low pole field. Then, in a flood-fill procedure, the six nearest neighbor voxels are all checked for their pole field. Any below the current threshold are unwrapped, and this stage recursively repeated on them. Any above the threshold are queued for unwrapping at the relevant threshold.

Pole-field threshold steps. The minimum and maximum pole field values were found. First, flood-fill unwrapping started with all voxels lower than the minimum pole-field threshold being unwrapped. When this was done, the threshold was increased by a single step and any further voxels unwrapped. This procedure was repeated until all of the steps had been done and all voxels unwrapped. For the evaluation of the different smoothing kernels, a high value (5000 steps) was chosen to ensure that the unwrapping was gradual enough to avoid an interaction. We then evaluated a range of numbers of steps with the optimum smoothing kernel. In fact, there was little computational cost to increasing the number of steps.

Stage 2: Seed voxel. This was determined by finding the centre of mass of the magnitude of the MR image corresponding to the first phase map, and then searching the neighboring voxels along 3 orthogonal directions for 16 voxels in each direction for the one with the lowest pole field value.

Test Data

Simulated data. We used both simulated noisy phase maps and acquired MRI data to test the algorithms. The advantage of the former is that it provides a known solution against which the result can be compared, and many datasets with carefully controlled amounts of noise can be generated. The latter allows confirmation that the algorithm works on real datasets.

The signal to be encoded in phase of the simulated maps was generated by creating a cylindrical linearly increasing field, superimposing some writing on it. The signal magnitude reduced in particular regions to give a number of noisy clusters. An example of a slice through a resultant wrapped volume is shown in Fig. 4a. The volumes were each 128 × 128 × 128 voxels. The cylindrical field increased at a rate of 0.5 radians per pixel. The letters had peak height of 4.85 radians but were smoothed to avoid over-abrupt changes at the edges. For each cluster of reduced signal strength, the magnitude of the signal was then reduced in a spherical gaussian form by a factor given by:

$$m(r) = 1 - e^{-0.01\,r^2}$$

$$\text{where } m(r) = \text{magnitude of noise;} \qquad (1)$$

$$r = \text{distance in voxels from centre of noise cluster}$$

Each reduced signal strength cluster was centred at a random position chosen from a uniform distribution throughout the volume. To the complex signal calculated using this phase and magnitude information was added noise of random phase and magnitude 0.1 relative to the unattenuated signal magnitude.

For the pole-field noise estimator, we first optimised the smoothing level on a high noise (100 cluster) image. This was unwrapped with no pole-field smoothing, and then 1, 2, 3, 4, or 5 iterations of smoothing. Then, with the optimum smoothing level, we compared performance of the algorithms as a function of noise level, by varying the number of clusters between 5 and 100 in steps of 5.

fMRI data. Three-dimensional phase maps were calculated by acquiring a pair of three-dimensional gradient echo data sets with all the parameters identical apart from the echo time. Two images are taken and one used as a reference image so that phase shifts due to the transmitter or receiver coils, to digital filtering, or to other constant shifts, are subtracted out. After complex reconstruction of the two datasets, the first dataset was normalised to have a magnitude of one, and then divided in complex mode into the second dataset. From Euler's formula, this is equivalent to a subtraction of their phases, and so the phase of the resultant complex map, after unwrapping and appropriate scaling, yields a map of the magnetic field. These phase maps were collected on ten subjects using a Bruker MEDSPEC 3T scanner at the Wolfson Brain Imaging Centre in Cambridge. We used a Bruker "GEFI TOMO" 3-D gradient echo sequence. We collected two phase maps (64 × 256 × 64 matrix size, 256 mm^3 FOV, TE = 7 and 16-1 ms, TR = 32 ms). Each of the phase maps took around 2 min to acquire. The difference between the TEs used was relatively long, at 9-1 ms, to ensure good signal to noise, but this causes there to be more phase wrapping, and it might be possible to reduce this while maintaining adequate signal. For the comparison of the different estimators of noise; only nine datasets were used as the raw magnitude of the individual field maps was not available for one.

To test the use of the final unwrapped data for the undistorting of EPIs, we used the seventh volume in an EPI time series (21 slices, TR = 10.029 s, TE = 27 ms, 128 × 128 voxels, FOV 250 × 250 mm, bandwidth in phase-encode direction 1563 Hz). After unwrapping, the Jezzard and Balaban (1995) undistorting algo-

rithm was applied. To undistort, we first had to rotate our phase map and EPIs into the same orientation and coregister them. As mentioned in the introduction, we cannot calculate the absolute offset of the phase maps; we can only unwrap the voxels relative to each other. However, the effect of the addition of a constant of the whole map will be to add a constant displacement to the entire image. This will then be eliminated by our coregistration stage, the exact details of which we hope to expand upon in a later publication.

Pixels were shifted in the phase-encode (y, anterior-posterior) direction. The shift was determined according to the function:

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

where $\phi(r)$ is unwrapped phase map, Δy is shift in voxels, Δt is evolution time, b_y is bandwidth in y direction, l_y is number of voxels in y direction.

After unwrapping, but before undistorting, the phase map underwent some additional processing. It was masked, to remove the areas of low signal outside the brain. We used an algorithm borrowed from Steve Smith's Brain Extraction Tool (http://www.fmrib.ox.ac.uk/fsl). The brightness level at which 2% of the voxels were dimmer was found (t_2), and the brightness level at which 98% were dimmer was found (t_{98}). Then the brightness $0.7\ t_2 + 0.3\ t_{98}$ was chosen as the threshold for masking.

Then, to guard against loss of signal within the image to be undistorted or imperfect matching between the phase map and the image, we dilated the phase map. Voxels where the phase map had been masked, but were within 10 mm of an unmasked voxel were replaced by the phase value of the closest unmasked voxel. Finally, the phase map was smoothed by a gaussian filter with FWHM 5 mm.

To resample the EPI images, we used a Hanning windowed Sinc function with a range of 9 voxels:

$$I'(x', y', z') = \frac{1}{2} \sum_{y=y'-R}^{y'+R} 1(x, y, z) \times \frac{\sin \pi(y - y')}{\pi (y - y')} \left(1 + \cos \frac{\pi (y - y')}{R + 1}\right), \qquad (3)$$

where R = range of interpolation = 9, I′ = interpolated image, I = raw image.

Note that as voxels only move in the y (anterior-posterior) direction, only one-dimensional interpolation is required.

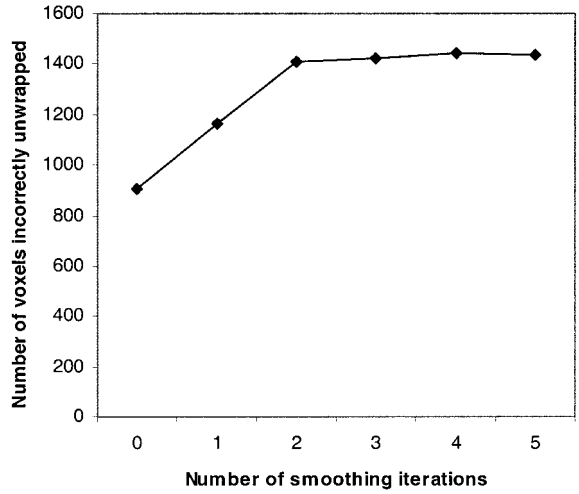

FIG. 3. The number of voxels incorrectly unwrapped using the pole-field noise estimator as a function of the number of smoothing iterations.

RESULTS

Simulated Data

We first optimized the degree of pole-field smoothing by varying the number of smoothing iterations parametrically. We unwrapped a noisy map (100 clusters) and evaluated performance by counting the number of incorrectly unwrapped voxels. The results are shown in Fig. 3. On this simulated data, optimal performance was obtained with no pole-field smoothing, and so this was used for the remaining investigations. Figure 4 shows an example of the simulated data used to test the unwrapping algorithms. As can be seen from Fig. 4b, the standard unwrapping algorithm fails when it encounters the noise clusters, and the errors then propagate throughout the map. In our new algorithm, the noisy areas are estimated either by generating the pole field (Fig. 4c) or using the negated magnitude (Fig. 4e), which then directs unwrapping. Figures 4d and 4f illustrate successful unwrapping using the new algorithm with magnitude and pole field information, respectively. To gauge the robustness at different levels of noise, we examined performance of the algorithms for different numbers of noise clusters. Performance is shown in Fig. 5. The new algorithm, using either of the noise estimators, has far fewer errors for all numbers of clusters. The difference becomes larger as the number of clusters increases. The pole-field and magnitude estimators perform similarly.

fMRI Data and Undistorting

We first optimised the degree of pole-field smoothing. As we are interested in the success of unwrapping in regions in which we have good signal, we first masked the images in the same way as we did prior to undis-

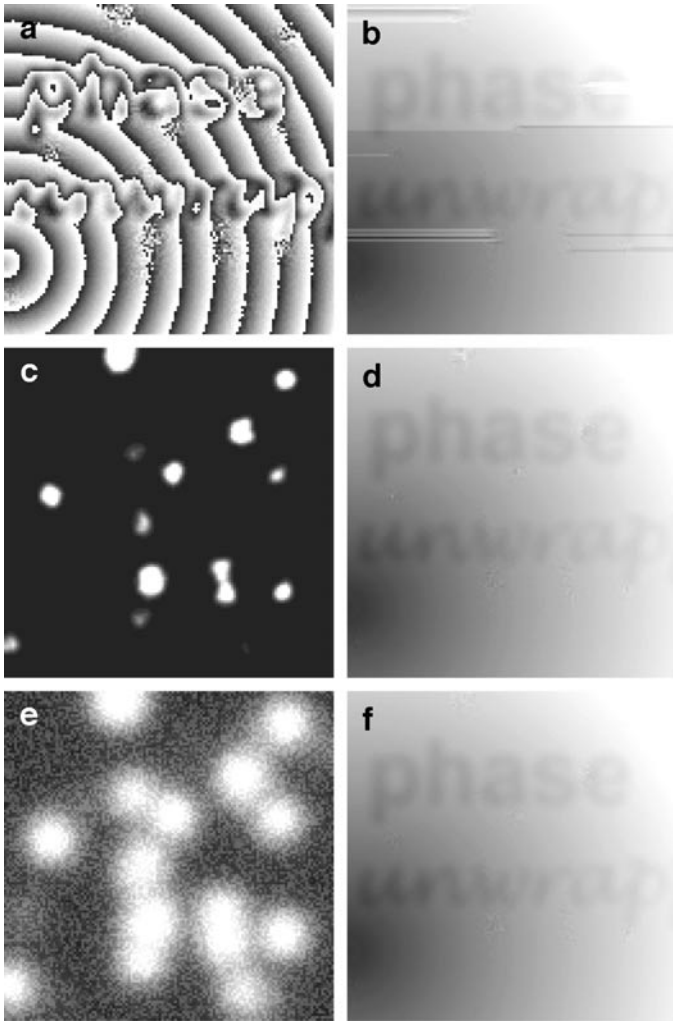

FIG. 4. An example slice through the simulated data used to test the algorithm. (a) shows the raw simulated data (see text for details); (b) the result of a simple unwrap; (c) the magnitude-field noise estimator; (d) the result of unwrapping with the new algorithm using this estimator; (e) the pole-field noise estimator and (f) the result of unwrapping using the pole-field estimator.

tortion. As we do not know the true unwrapped value for each voxel, we used a different measure to evaluate success: we counted the number of jumps between adjacent voxels greater than π in the unwrapped, masked maps. The results are shown in Fig. 6. It can be seen that optimal performance was obtained with a single smoothing iteration, and this was used for the remaining investigations. Figure 7 compares the output of the simple unwrapping procedure and the new algorithm. It can be seen that where the simple algorithm fails catastrophically in many areas, no errors can be seen in the output from the new method. More than 50 maps have been unwrapped in our laboratory, and no algorithm failures have been identified. As an illustration, Fig. 8 shows sagittal slices near midline of raw and unwrapped, masked phase maps for 10 subjects. Note

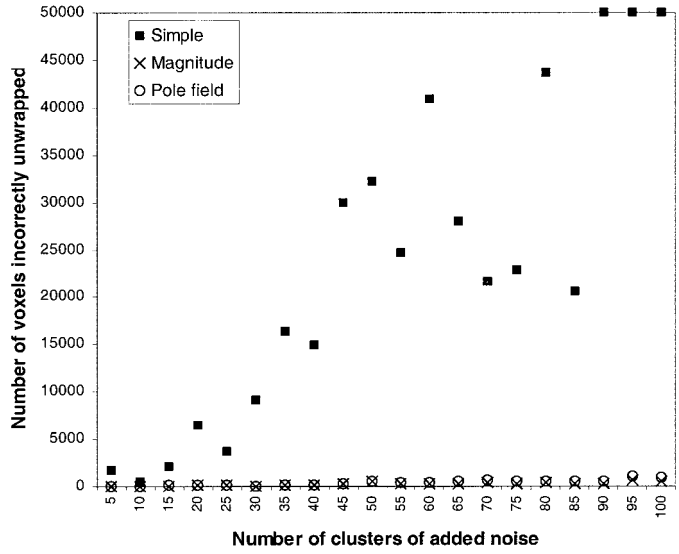

FIG. 5. The number of voxels unwrapped incorrectly with the standard simple unwrap, and the new algorithm with two different noise estimators, as a function of the degree of noise added to the simulated data.

the similarity in the pattern of distortions for the different subjects.

We also investigated the effect of varying the number of steps in the unwrapping process, which will affect how gradually unwrapping approaches noise. We used the same number-of-jumps measure to evaluate performance as was used for assessing the effect of smoothing. Figure 9 shows that the results change little after a few tens of iterations in the unwrapping process, although in our implementation, there was little computational cost in having a larger number of steps. Figure 10 illustrates successful undistorting of

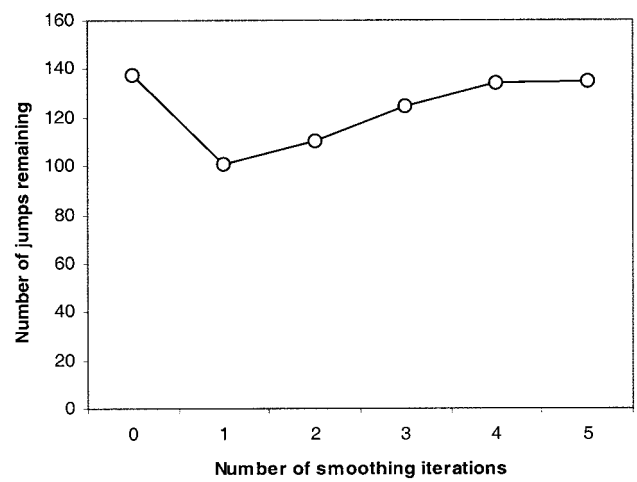

FIG. 6. The effect of varying the number of smoothing iterations of the pole-field noise estimator on the number of jumps remaining in the unwrapped map.

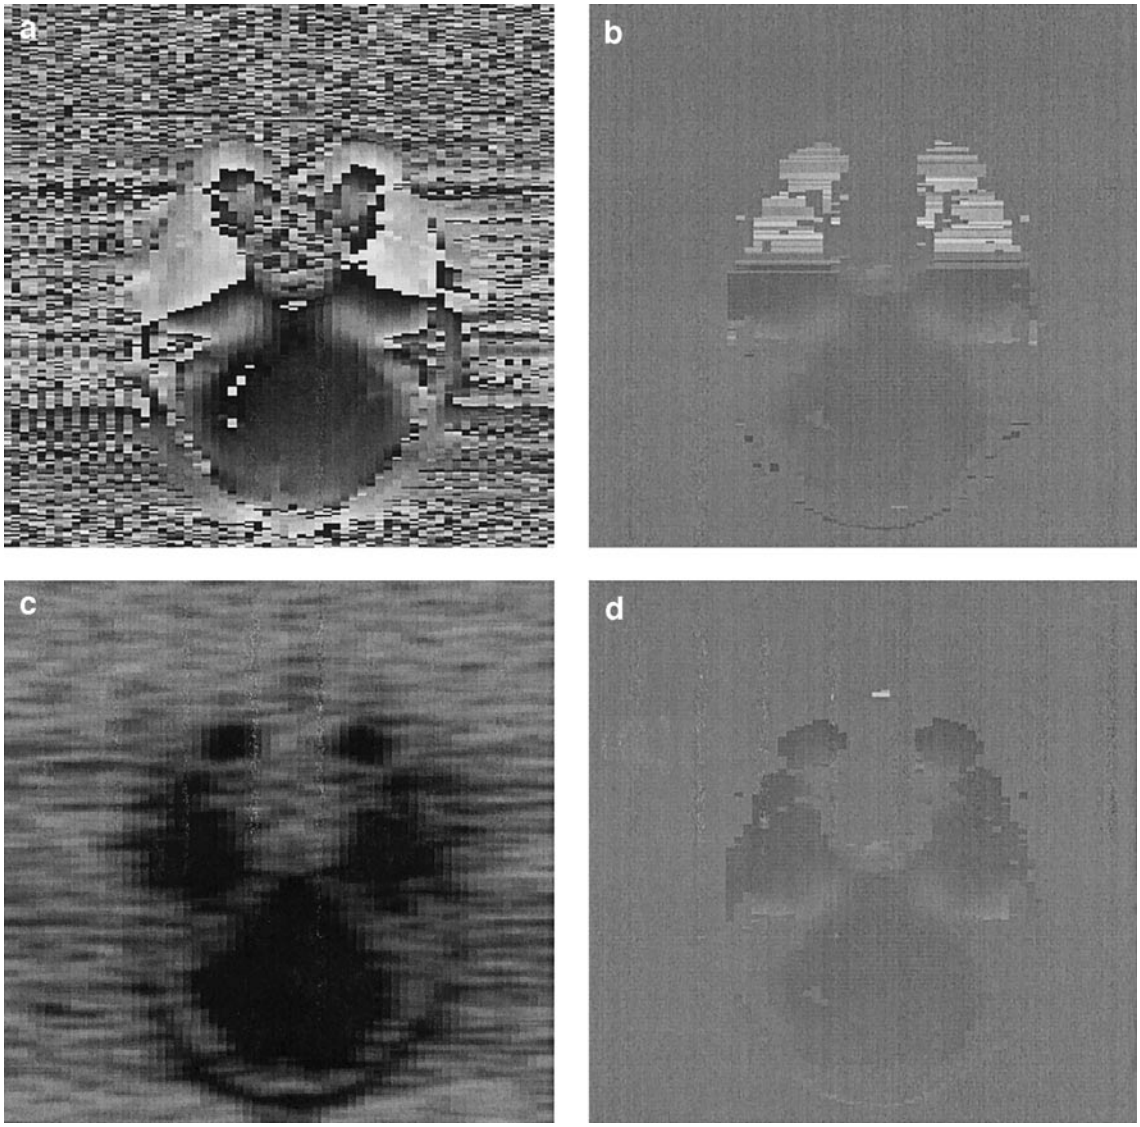

FIG. 7. Comparison of: (a) signal magnitude; (b) phase unwrapped with the standard algorithm; (c) the pole field; and (d) phase unwrapped with the new method.

an EPI image. Inspection shows the undistorting leads to an image more similar in shape to the structural.

DISCUSSION

A new algorithm has been presented for three-dimensional phase unwrapping. It robustly unwraps simulated and measured MRI data. It is fast, even on large datasets. We have used it to obtain magnetic field maps and shown subsequent successful undistorting of MRI images using the pixel-shift method.

Other noise-immune unwrapping methods are also well established in fields such as optics and synthetic aperture radar reconstruction and might be applied to three dimensions (Ghiglia and Pritt, 1998). For example, one strategy might be to divide the raw phase map into many distinct regions, and then try to minimize a cost function at the boundaries between the regions. Another is to find a function that can model the unwrapped map and then optimize the parameters that define this by minimizing the squared difference between this model when wrapped and the raw map. While both of these strategies might be successfully applied to three dimensions, there are problems associated with them. Tiling strategies can be slow: the number of regions, and interactions between regions, increases rapidly as maps become larger. Least-squared methods are systematically biased, underestimating the total phase change over a map (Loffeld *et al.*, 1996). However, both may have applications in particular circumstances, where time is not a problem, or the data can be precisely modeled.

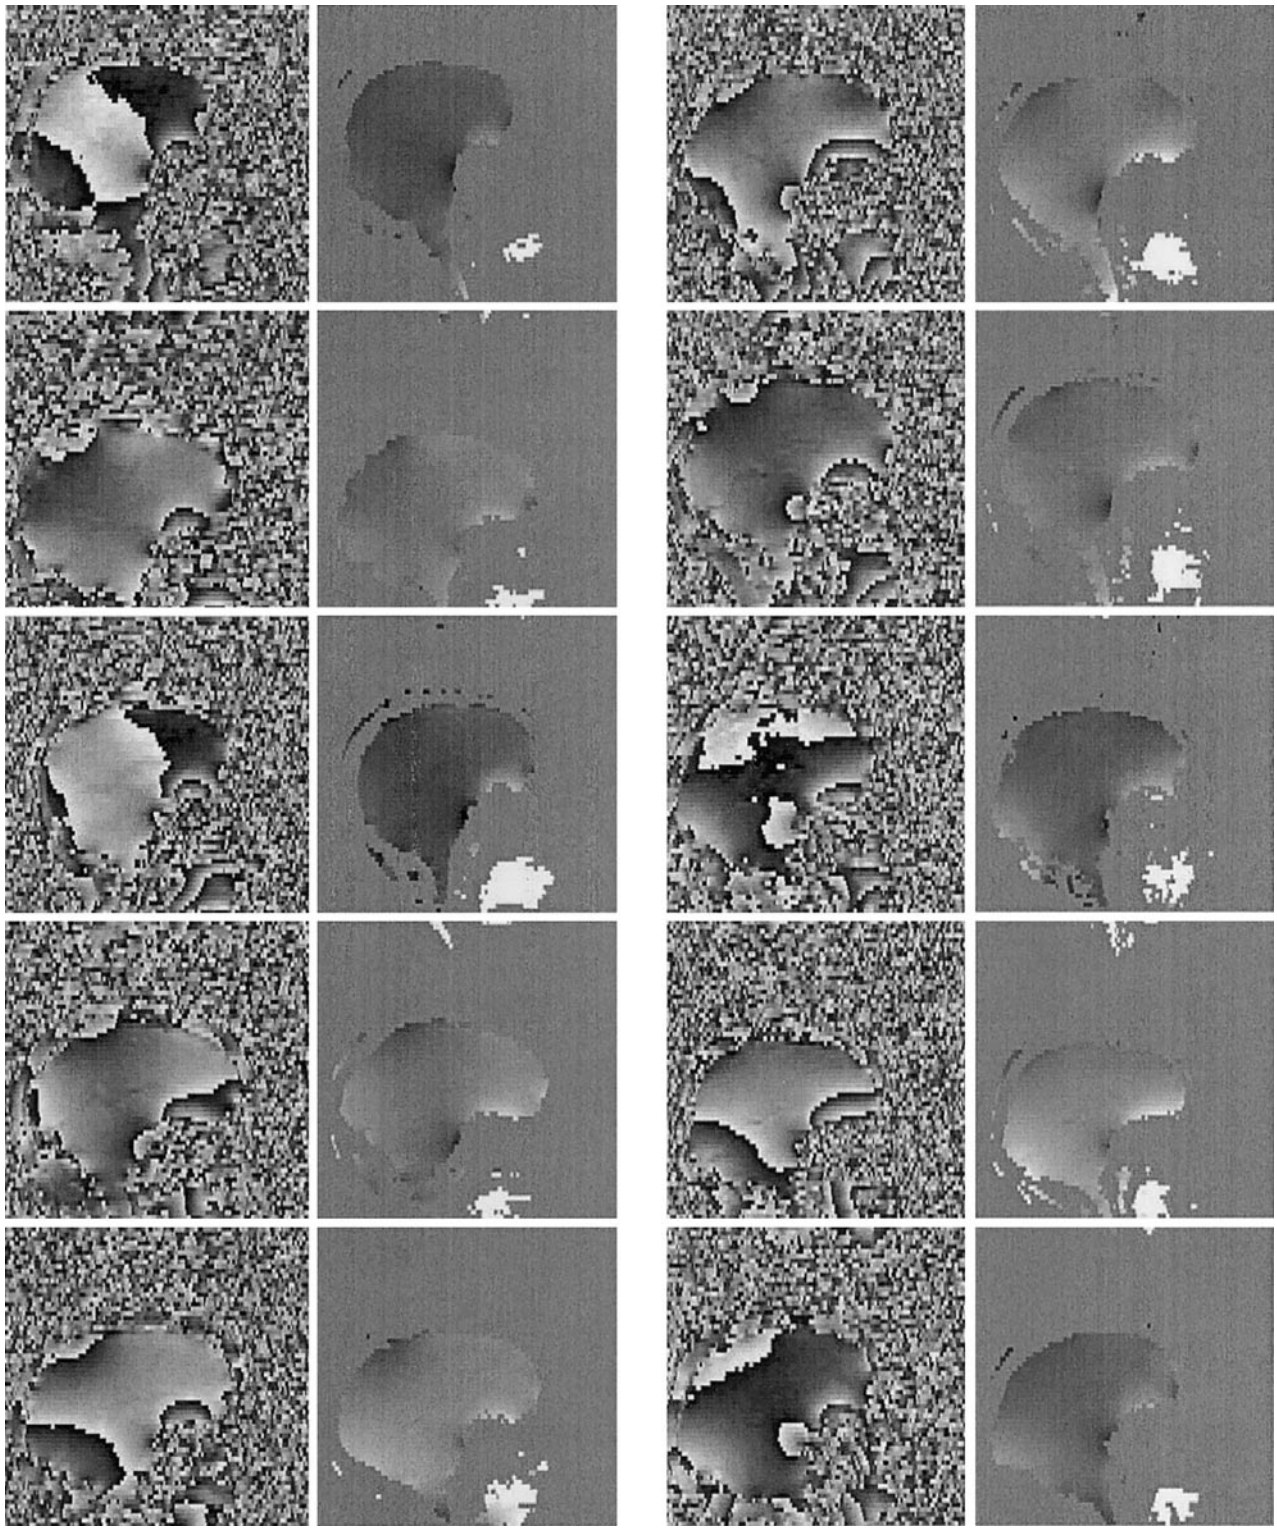

FIG. 8. Sagital sections through the phase maps of ten subjects. The first and third columns show raw unwrapped phase maps; the second and fourth unwrapped, masked maps.

As mentioned in the introduction, we can acquire inhomogeneity maps using EPI or non-EPI acquisition sequences. We believe it is more convenient to use non-EPI sequences for two reasons. First, if EPIs are highly compressively distorted in a particular region, then if we collect phase maps using an EPI sequence,

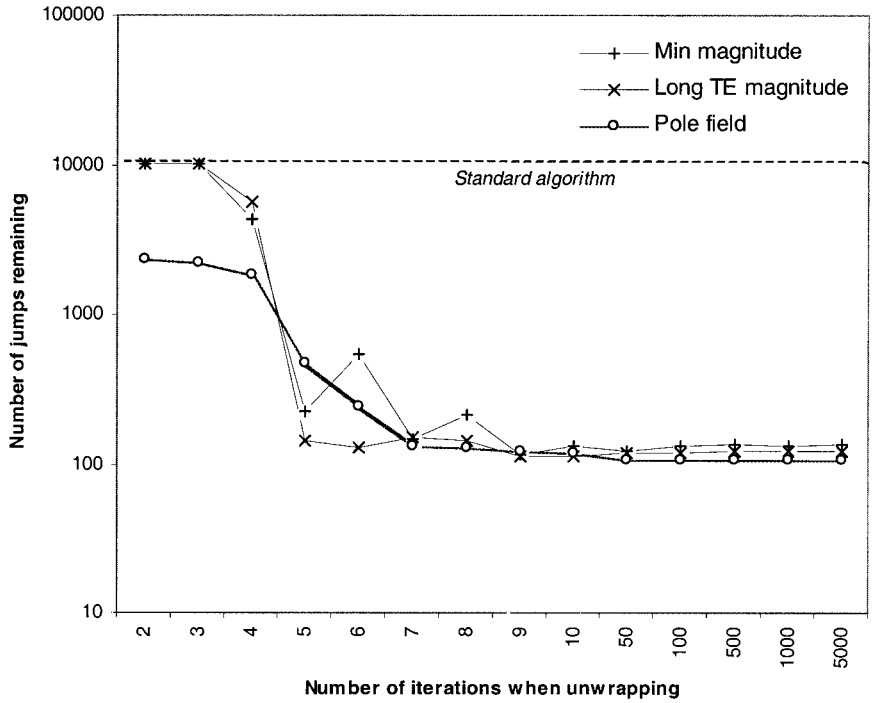

FIG. 9. Mean numbers of jumps remaining as a function of the number of noise-estimation field iterations for three different estimators. Standard errors over subjects, which are not shown on the graph for clarity, decreased with increasing numbers of iterations, and were similar for the three noise estimation methods.

these will also be compressed and least information will be available for undistorting in one of the regions where it is most needed. Second, if the distortion field is measured in undistorted space, undistorting is a pull problem: for each voxel in the destination undistorted space, we can look up the appropriate voxel in the warped EPI map. This contrasts with the more difficult push problem, where for each voxel in distorted space, we know where it moves to in undistorted space. This latter is unlikely to give us complete coverage in the undistorted space. We either need to do some kind of fit and then invert the transformation or interpolate the final dataset. However, while non-EPI sequences do have these advantages, as they take longer to acquire, they are more susceptible to subject movement. Also, EPI sequences are fast enough to allow an alternative to spatial unwrapping, where images with many different TEs can be acquired (Chen and Wyrwicz, 1999), and unwrapping is performed over this dimension.

As discussed in the Introduction, there are theoretical limitations on what type of images may be spatially unwrapped. In particular, if the field that we are attempting to measure contains true phase discontinuities, these will cause unwrapping errors unless invalid paths crossing them can be eliminated. True phase discontinuities may occur when measuring magnetic fields that change very rapidly over space (e.g., from the B1 field of an RF coil), or when using phase-encoding to measure other quantities, such as the displacement of the myocardium or very rapid flow. However, from our measurements on many brains with the parameters described, the changes of the B0 field over space are gradual enough to avoid true phase discontinuities.

In our undistorting procedure, we did not correct for changes in intensity that may have been caused by compression or expansion of the images. If the raw images are undistorted, and then passed on to a stage of parametric statistical analysis in which a *t* statistic is calculated, a uniform scaling of the signal will not have an effect. If movement artefacts interfere with this process, it might be better to try to correct the intensity in the images. This can be attempted by the linear approximation of dividing by the degree of compression or expansion (Jezzard and Balaban, 1995) or by using an alternative to the pixel shift method (e.g., conjugate phase, Weisskoff and conjugate gradient methods; Munger *et al.*, 2000). However, recent work by Jenkinson (2001) has suggested that intensity corrections can induce an unacceptable degree of noise.

Another recent approach to tackling geometric distortions in EP images has been taken by Andersson *et al.* (2001). Their algorithm uses the changes in image shape that occur as an object moves in a distortion field. It would be interesting to compare their measurements with those obtained from direct measurement of the distortion field using a technique such as that used here. When mapping magnetic fields, the phase differ-

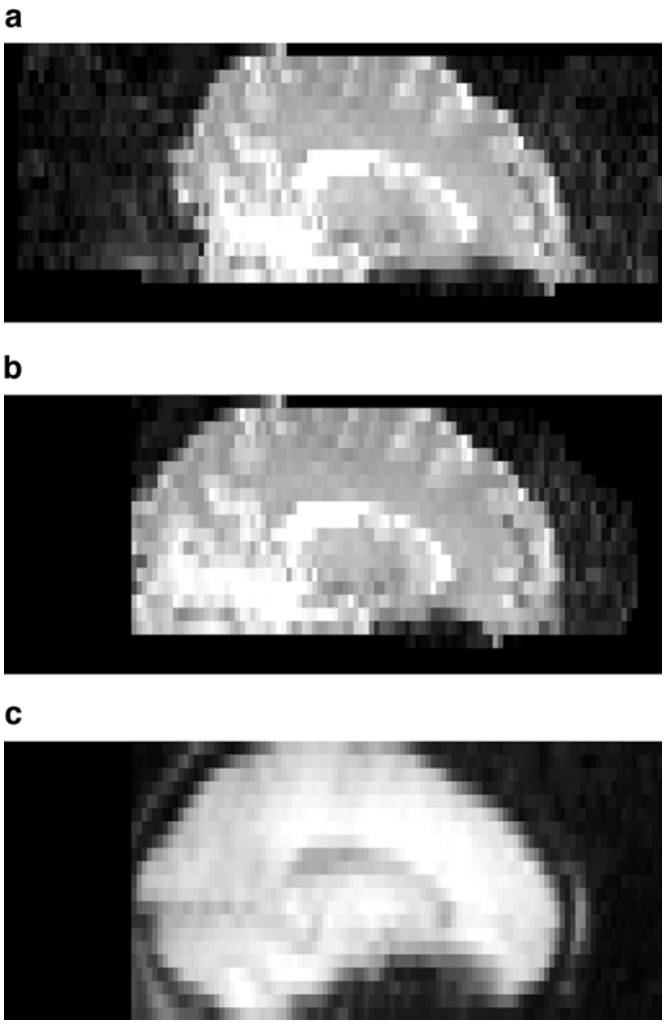

FIG. 10. Result of undistorting a single sagital EPI slice. The top panel shows the raw EPI, the middle panel the same slice undistorted, and the bottom panel the magnitude of the phase image, which is collected with a non-EPI sequence and hence is not distorted.

ence we measure between two scans will be proportional both to the time between the scans and the magnetic field strength. For a fixed difference in time, the phase difference will be proportional to the magnetic field. Hence, at higher field strengths, there will be greater phase wrapping and a robust algorithm even more important.

In conclusion, we present a new algorithm for performing three-dimensional phase unwrapping by identifying noisy regions in phase maps and unwrapping these last. The algorithm is robust on simulated and MRI data. Furthermore, as a general three-dimensional phase unwrapping method, it could be applied to phase data from MR that represents quantities other than magnetic field strength, such as velocity, or to data from other fields such as optics.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the help of Adrian Carpenter, Emma Williams, and Kay Martin at the Wolfson Brain Imaging Centre, Cambridge, for help in acquiring the phase maps; Matthew Brett and Ingrid Johnsrude for their help with the tools of the trade; and Ian Nimmo-Smith for his helpful comments on this manuscript.

NOTE

Compiled software and the source code to the algorithms used here are available. Please contact the author.

REFERENCES

Andersson, J. L. R., Hutton, C., Ashburner, J., Turner, R., and Friston, K. 2001. Modelling geometric deformations in EPI time series. *NeuroImage* **13:** 903–919.

Bone, D. J. 1991. Fourier fringe analysis: The two-dimensional phase unwrapping problem. *Appl. Optics* **30:** 3627–3632.

Chen, N. K., and Wyrwicz, A. M. 1999. Correction for EPI distortions using multi-echo gradient-echo imaging. *Magn. Reson. Med.* **45:** 525–528.

Cusack, R., Huntley, J. M., and Goldrein, H. T. 1995. Improved noise-immune phase-unwrapping algorithm. *Appl. Optics* **34:** 781–789.

Dirac, P. A. M. 1931. Quantised singualarities in the elctromagnetic field. *Proc. R. Soc. London A* **133:** 60–72.

Ghiglia, D. C., and Pritt, M. D. 1998. *Two-Dimensional Phase Unwrapping: Theory, Algorithms and Software.* Wiley, Chichester, New York.

Huntley, J. M. 1989. Noise immune phase unwrapping algorithm. *Appl. Optics* **28:** 3268–3270.

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.

Loffeld, O., Arndt, C., and Hein, A. 1996. Estimating the derivative of modulo-mapped phases. ESA Workshop on Applications of ERS SAR Interferometry, Zurich. http://www.geo.unizh.ch/rsl/fringe96/papers/loffeld-et-al/.

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