A Fast, Automated, N-Dimensional Phase Unwrapping Algorithm

FMRIB Technical Report TR01MJ1

Mark Jenkinson

Oxford Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB),
Department of Clinical Neurology, University of Oxford, John Radcliffe Hospital,
Headley Way, Headington, Oxford, UK

Abstract

This report investigates the general problem of phase unwrapping for arbitrary n-dimensional phase maps. A cost function-based approach is outlined which leads to an integer programming problem. To solve this problem a best-pair-first region merging approach is adopted as the optimisation method. The algorithm has been implemented and tested with 3D MRI medical data for fMRI applications in EPI unwarping and rapid, automated shimming.

Keywords: phase unwrapping; echo planar imaging (EPI); field mapping.

1 Introduction

Phase unwrapping is a problem that occurs in several fields as diverse as Synthetic Aperture Radar and MR Angiography. In all cases the problem is that the measured phase signal can only take on values in a 2π range, whilst the original phase signal can take on any value. That is, $\theta_{measured} = \theta_{original} \bmod 2\pi$. This "wrapping" leads to artificial phase jumps being introduced near boundaries. To recover the original, smooth phase (up to some global 2π offset) requires an *unwrapping* algorithm.

Medical MR images require phase unwrapping for many applications such as field mapping (often used to remove geometric distortion in EPI [14, 10, 20]), chemical shift mapping [6] and velocity measurement [15] (used in MR angiography). Also, in medical imaging applications the phase maps are typically 3D, rather than the 2D maps acquired in applications such as optical and radar imaging. It is therefore important to have a method which takes into account the 3D nature and constraints of the problem. A 2D method may be applied on a slice-by-slice basis, but can result in phase wraps existing in the through-slice direction after unwrapping, which must be fixed by some other post-processing option. Furthermore, 3D images usually contain many more voxels than typical 2D images, and so speed is an issue, especially in applications such as automatic-shimming where rapid, automatic and robust methods for phase unwrapping are essential.

Furthermore, in 3D MRI, a large proportion of the voxels do not contain any interesting signal (e.g. air, bone, *etc.*) and so the signal to noise ratio (SNR) in these voxels is extremely low. Consequently the true phase information is swamped by noise. By disregarding such voxels when unwrapping, both robustness and speed can be greatly improved.

Existing unwrapping techniques can be categorised according to their:

- Dimensionality (1D, 2D, 3D, *etc.*);
- Application (Synthetic Aperture Radar, general optical interferometry, MR angiography, MR chemical shift mapping or MR field mapping);
- Approach (fitting functions, cost function optimisation, filtering, region growing / merging).

Methods based on fitting functions (usually truncated Taylor series [13, 5] or polynomials [10, 6]) can be easily generalised to work with data of any dimension. However, these methods impose considerable smoothness on the phase since the functions can not vary too rapidly, as otherwise no unwrapping will be performed. This is undesirable for mapping applications where there are small, rapidly changing regions that are not smooth and hence not estimated correctly.

Generalising other methods can be difficult because they use application specific knowledge, such as the coherence function in Synthetic Aperture Radar (SAR) imaging [4, 11, 12] or multiple MR images for chemical shift mapping [19] or prior expectations about expected velocities [9, 21, 3]. In addition difficulties can arise because the underlying method does not generalise naturally, such as for the filtering methods (e.g. 1D Markov [7] or 2D Kalman [12]).

Many of the existing unwrapping techniques proposed for MR applications were only written to deal with 2D slices. These methods are based on various approaches including region growing [16, 18, 1], cost function optimisation [17], neighbourhood voting/filtering [2] and phase path integration [8]. Although these approaches all generalise straightforwardly to 3D images they are generally inefficient and consequently unacceptably slow for many applications, especially rapid, automated shimming.

This report presents a fast, robust method for unwrapping phase maps that is general and can be applied to a single "image" of any dimension. The method uses a region-merging approach to find the minimum of a cost function that penalises phase differences across boundaries. It has been implemented specifically for 2D and 3D medical images (and is available at www.fmrib.ox.ac.uk/fsl), but the principles apply equally well to phase unwrapping problems of any dimension. In fact, the existing implementation would only require minimal modification to tackle any particular phase unwrapping problem.

2 Theory

2.1 Problem Definition

Let the measured (wrapped) phase be ϕ_j where j is an n-dimensional index specifying the spatial location (or voxel). The objective is to find the unwrapped phase denoted θ_j such that:

$$\theta_j = \phi_j + 2\pi m_j \tag{1}$$

where the integer value m_j at each voxel specifies the corrective offset required. The unwrapped phase should be smooth and not contain the 2π jumps that are typically seen in the wrapped phase.

2.2 Cost Function Formulation

The approach taken here to solve this problem starts by dividing the whole (n-dimensional) volume into N regions, each of which contains no phase wrappings. Each region is then treated as a single entity by the algorithm, with a single offset. As the number of regions to be dealt with is much less than the number of voxels this provides a considerable speed increase.

Now consider the interface, or boundary, between two regions, A and B.

The unwrapped phase, θ in region A will be related to the wrapped phase, ϕ by: $\theta_{Aj} = \phi_{Aj} + 2\pi M_A$ where M_A is an integer offset that is the same for all voxels in region A.

A cost function suitable for unwrapping is one that penalises phase differences along this interface. One good candidate is the sum of squared difference in phase. That is:

$$C_{AB} = \sum_{j,k \in \mathcal{N}(j)} (\theta_{Aj} - \theta_{Bk})^2 \tag{2}$$

$$= \sum_{j,k \in \mathcal{N}(j)} (\phi_{Aj} - \phi_{Bk} + 2\pi M_{AB})^2 \tag{3}$$

where $M_{AB} = M_A - M_B$ and M_A is the integer offset applied to region A. The summation is taken over two indices, j and k, where j is the index of a voxel in region A, while k is the index of a voxel in region B that is also in the (simply connected) neighbourhood of the voxel with index j (*i.e.* $j \in \mathcal{I}(A)$ and $k \in \mathcal{N}(j) \cap \mathcal{I}(B)$).

The total cost over the whole volume is the sum over all the interfaces:

$$C = \sum_{A,B} C_{AB}. \tag{4}$$

Differentiating with respect to the individual parameters, M_A, gives

$$\frac{\partial C_{AB}}{\partial M_A} = 8\pi^2 N_{AB} M_{AB} + 4\pi P_{AB} = 0 \tag{5}$$

which yields the equations for the minimum cost solution:

$$M_{AB} = M_A - M_B = \frac{-P_{AB}}{2\pi N_{AB}}. \tag{6}$$

where N_{AB} is the number of interfacing voxel pairs and $P_{AB} = \sum_{j,k \in \mathcal{N}(j)} (\phi_{Aj} - \phi_{Bk})$.

However, not only does this tend to generate more equations (equal to the number of distinct interfaces) but, more importantly, only *integer* values can be taken by the parameters M_A. Therefore, equation 6 cannot be solved exactly as it is an integer programming problem. Finding the optimal solution (minimum cost) for such a problem is a difficult and often time-consuming task. For example, there are 2^N "neighbouring" solutions of the continuous solution (where N is the number of regions).

2.3 The Region Merging Algorithm

An alternative solution to the problem is to treat it as a set of potentially inconsistent constraints:

$$M_{AB} = \text{round}\left(\frac{-P_{AB}}{2\pi N_{AB}}\right). \tag{7}$$

The problem is then one of finding the set of consistent constraints which lead to a low (ideally minimal) cost.

Let $K_{AB} = -P_{AB}/(2\pi N_{AB})$.

Consider the difference in cost between $M_{AB} = L_{AB} = \text{round}(K_{AB})$ and $M_{AB} = L_{AB} \pm 1$. The cost difference is:

$$\Delta C_{AB} = 8\pi^2 N_{AB}\left(\frac{1}{2} \pm (K_{AB} - L_{AB})\right) \tag{8}$$

Since, by definition of rounding, $-\frac{1}{2} \leq K_{AB} - L_{AB} \leq \frac{1}{2}$, then $\Delta C_{AB} \geq 0$ for both cases, confirming that $M_{AB} = L_{AB}$ is (locally) the best solution.

The smallest cost difference is:

$$\Delta C_{AB} = 8\pi^2 N_{AB}\left(\frac{1}{2} - |K_{AB} - L_{AB}|\right). \tag{9}$$

Choose the interface where this cost difference is the largest. That is, choose the pair of regions where getting the offset "wrong" is the most disastrous.

Therefore, select the pair of regions according to:

$$AB = \arg\max_{RS} \Delta C_{RS}. \tag{10}$$

Once this pair is selected, merge them into a new, single region Q using the "optimal" offset L_{AB}. Update the statistics for all interfaces to this new region. That is, for all regions W, the new statistics are:

$$N_{QW} = N_{AW} + N_{BW} \tag{11}$$

$$P_{QW} = P_{AW} + P_{BW} \tag{12}$$

From these, calculate K_{QW}, L_{QW} and then ΔC_{QW} and then repeat the selection process to choose a new pair of regions to merge.

Keep merging regions until there are no more interfaces between regions (usually when only one region remains if there is a single connected set of regions initially).

As each iteration of this algorithm merges two regions, the number of iterations required is $N - 1$. Within each iteration a search for the maximum ΔC value must take place, followed by the updating or deletion of all constraints related to the two selected regions.

2.4 Creating the Initial Regions

Many different approaches could be taken to determine the initial regions. The one which we have taken is to split the volume up into connected components, inside which the phase remains within given bounds. That is, a phase partition is initially chosen — say $\{[-\pi, -2\pi/3), [-2\pi/3, -\pi/3), [-\pi/3, 0), \ldots, [2\pi/3, \pi]\}$. Then, for each partition, create a mask of voxels where the wrapped phase is within the partition. Determine the connected components of the mask and assign each a unique region number. Repeat for all other partitions, assigning different region numbers.

This method ensures that the phase difference between two voxels *within* an region must be less than the size of an individual phase partition. Hence, the smaller the partition, the less likely it is that a connected component will contain a phase wrap.

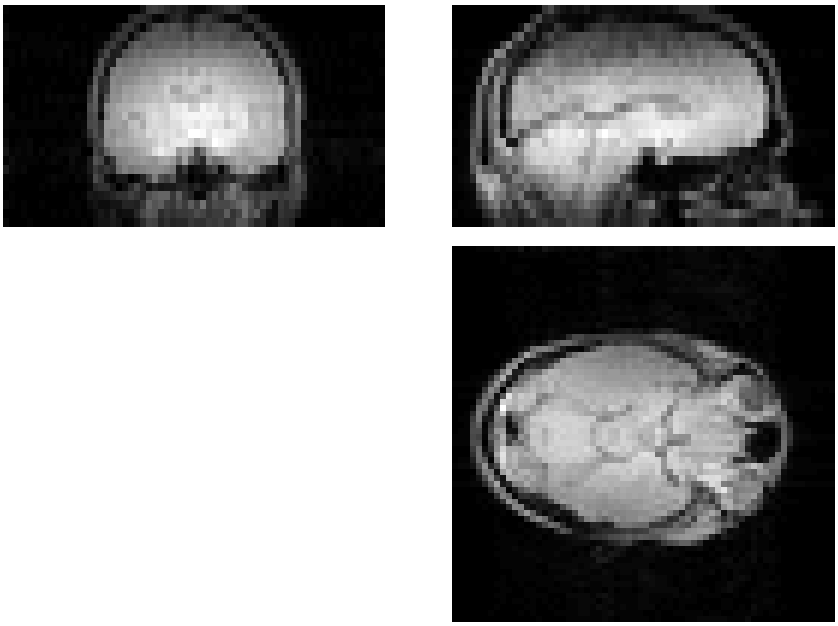

Figure 1: Magnitude Image

3 Materials and Methods

The above unwrapping algorithm has been implemented in C++ for 3D MR images. This program, called *PRELUDE* (Phase Region Expanding Labeller for Unwrapping Discrete Estimates), is available as part of the FMRIB Software Library (FSL) at www.fmrib.ox.ac.uk/fsl.

In order for the implementation to be fast it is necessary to use data structures that are efficient for this algorithm. We have used the Standard Template Library data structure, *map*, to store the constraints between pairs of regions. This has the advantage of constant insertion and deletion complexity, plus logarithmic search complexity when using the region numbers as the key. The only step where the map is inefficient is for finding the maximum ΔC value, which is $O(N_C)$, where N_C is the number of constraints. However, as this is only done at most N times (once per iteration), it is a small penalty to pay for having very efficient insertion, deletion and search operations, since these are used several times (on average) per iteration. Consequently the overall complexity (worst case) is $O(N^2 \log(N_C) + N.N_C)$, since each iteration has a complexity of $O(N_C) + O(N \log(N_C))$, representing the search and updating/deletion operations respectively.

4 Results

4.1 Real Data Examples

In this section some example images are shown to illustrate how the algorithm works.

Consider the corresponding magnitude and phase images shown in figures 1 and 2. By thresholding the magnitude image at a robust threshold[1] a mask is calculated, as shown in figure 3. Regions are then created and labelled, excluding the masked out areas, as shown in figure 4. Finally the region merging algorithm is used to unwrap the phase, with the result shown in figure 5. Note that for this image there were 1920 regions created (using 6 phase partitions) which is substantially less than 102400 voxels in the image or 31566 voxels within the mask. Furthermore, of the regions created 1102 were single voxels, while only 163 were larger than 27 voxels. The total time for unwrapping was 1.7 seconds on a desktop PC.

The next example shows the performance on high resolution images. This venogram image is 512 $\times$ 512 $\times$ 64 with voxel dimensions of 0.38mm $\times$ 0.50mm $\times$ 1.12mm. Figures 6, 7, 8 and 9 show the magnitude, wrapped phase, mask and unwrapped phase images respectively. Note that for this large image the regions were restricted to be planar (in order to that reduce the chance of two areas of differing phase offset being

[1]The robust threshold used here is set as $0.1 * (I_{98} - I_2) + I_2$ where I_n is the nth percentile of the image I.

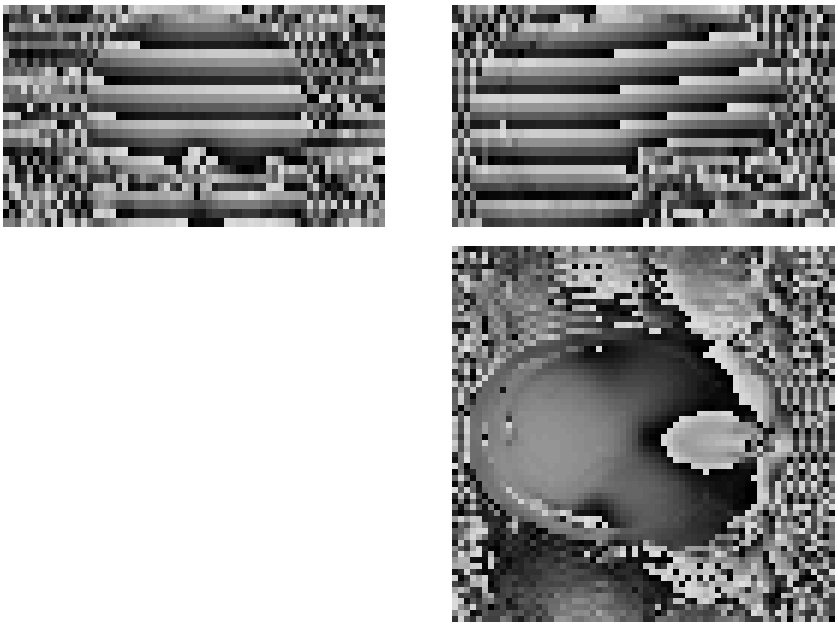

Figure 2: Measured (wrapped) Phase Image

Figure 3: Mask Image

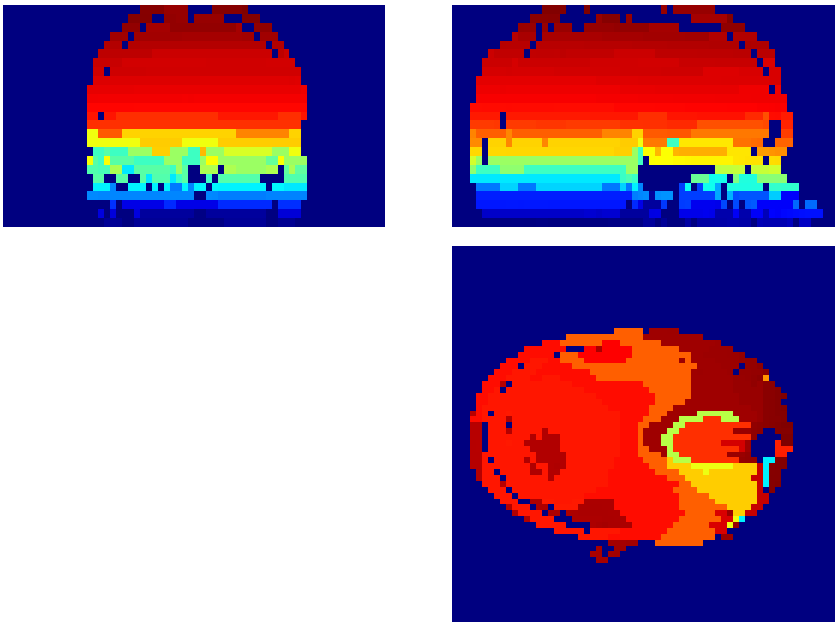

Figure 4: Region Label Image

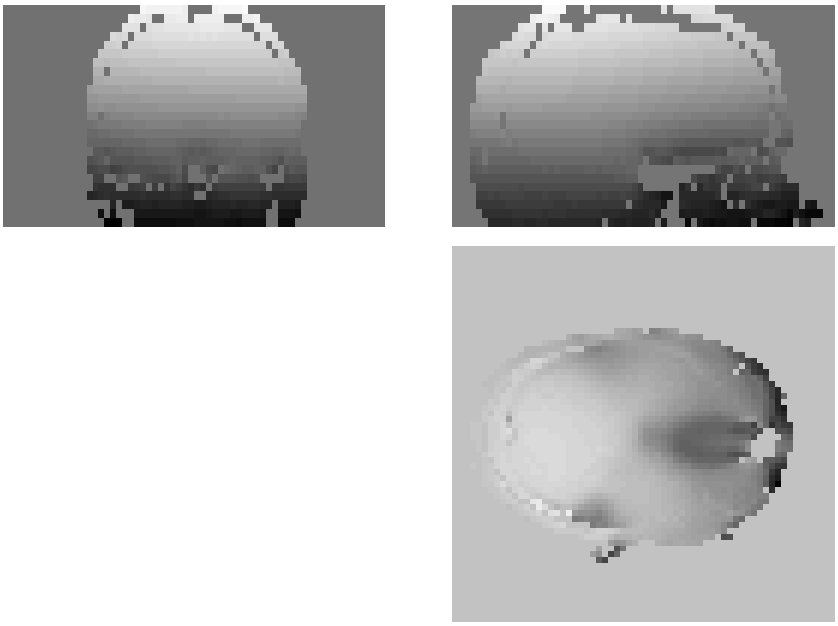

Figure 5: Resulting Unwrapped Phase Image

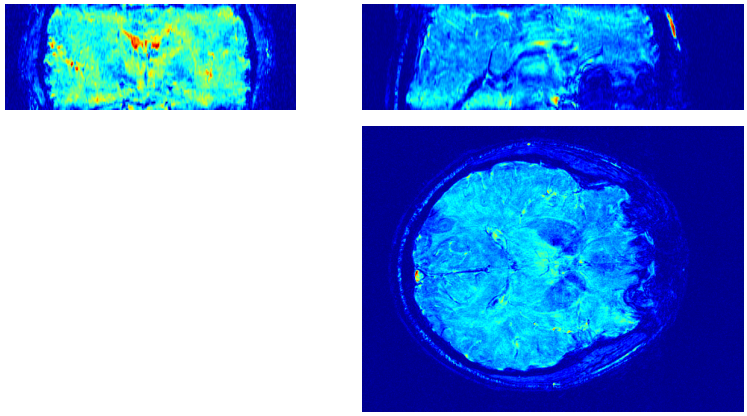

Figure 6: Magnitude Image

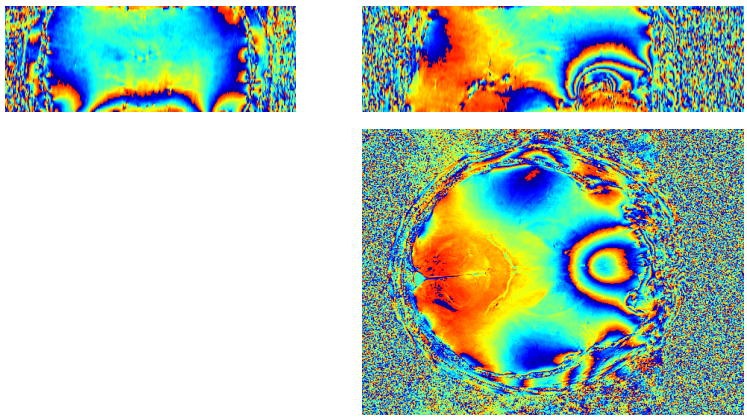

Figure 7: Measured (wrapped) Phase Image

connected together by small chains of noisy voxels — see discussion) and the processing time was slightly over an hour.

4.2 Quantitative Validation

To quantitatively measure the performance of the phase unwrapping algorithm with respect to signal to noise ratio (SNR), a set of simulated images were unwrapped. The results were then compared with the original (ground truth) phase and the mis-classification ratio (MCR) calculated. In this context the MCR is the number of voxels that were incorrectly unwrapped (that is, had the incorrect multiple of 2π added to them) divided by the total number of voxels.

The base image used was a $64 \times 64 \times 25$ complex image with unit magnitude at all voxels and a quadratic phase function centred in the middle of the image (see figure 10). The maximum phase gradient in the base image was π radians between neighbouring voxels, which occurs at the outermost edge of the image. In the interior of the image the phase gradient changes linearly with position, being zero at the centre of the image.

Various amounts of Gaussian noise was added to the real and imaginary parts of the base image to generate a set of synthetic images, each with a different SNR. That is, $I_{synth} = I_{base} + \sigma(n_R + jn_I)$, where n_R and n_I are independent, unit variance, Gaussian noise images. The SNR of this image is given by the signal amplitude (unity in this case) divided by the total noise amplitude, giving SNR $= 1/\sqrt{2}\sigma$. Some of these synthetic images are shown in figure 11. The added noise induces phase errors which can be measured by comparing the phase of I_{synth} with the phase of I_{base}, the true phase, thus allowing the standard deviation of the induced phase change to be calculated.

The unwrapping results are shown in table 1 as well as some examples which are displayed in figure 12. It can be seen that the the algorithm was extremely accurate and robust, with zero MCR for SNR greater than 5, and only a single voxel erroneously unwrapped with an SNR of 5. It is only at very low SNR (1 and 2),

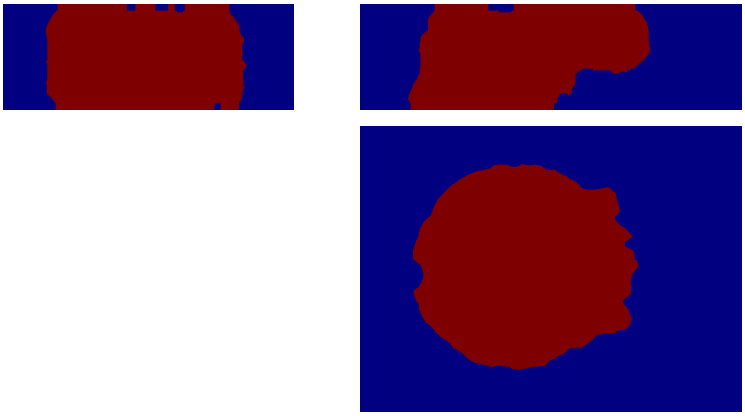

Figure 8: Mask Image

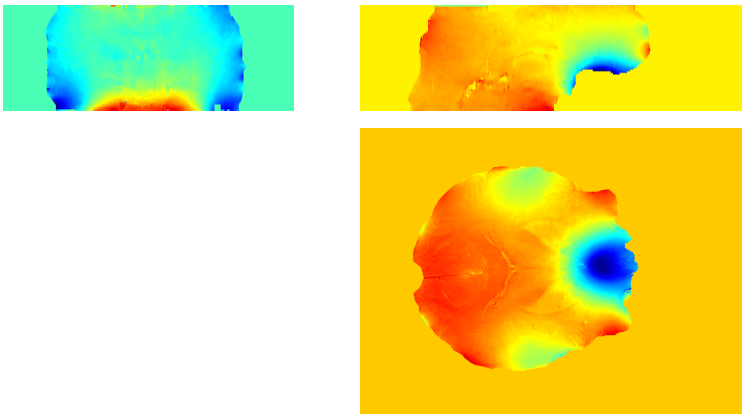

Figure 9: Resulting Unwrapped Phase Image

SNR	Noise level	std. dev. phase	MCR %
1	1	50.2°	80.0
2	0.5	22.5°	7.3
5	0.2	8.18°	0.001
10	0.1	4.05°	0.000
20	0.05	2.03°	0.000
50	0.02	0.81°	0.000
100	0.01	0.40°	0.000
200	0.005	0.20°	0.000
500	0.002	0.081°	0.000
1000	0.001	0.040°	0.000

Table 1: Results of simulated tests for phase unwrapping

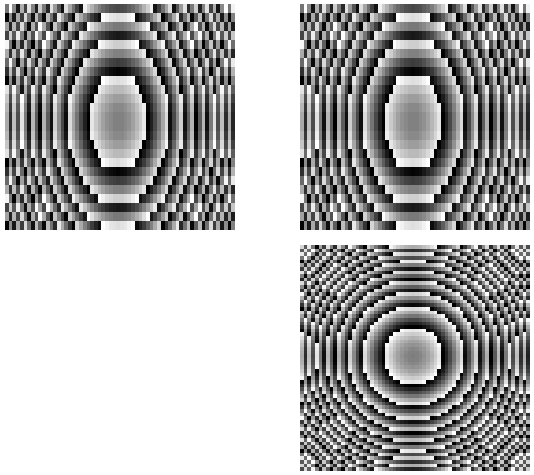

Figure 10: Synthetic phase image

where the standard deviation of the phase exceeds $20°$, that a significant number of voxels were incorrectly unwrapped. In fact, even with a SNR of 1, the unwrapping was largely successful in the central part of the image (where the phase gradient was less than $\pi/2$ radians per voxel).

For most MR imaging methods the SNR is significantly better than this. For example, a SNR of 50 (equivalent to a phase error of approximately $1°$) a B0 mapping sequence is typical at 3T. This is an order of magnitude better than required for this algorithm to work well and therefore this method should be both accurate and robust for phase images acquired with typical MR sequences.

5 Discussion and Conclusion

This report has presented a fast, fully automated, robust phase unwrapping algorithm that can be applied for phase maps of any dimension. The implementation of this algorithm for 2D and 3D MRI images has been demonstrated and the code is currently being used for EPI unwarping and rapid, automated shimming applications in fMRI.

There are two potential weaknesses in the current approach. The first is during the creation of the initial regions. If a single region includes two areas where the original phase differs by more than 2π then the algorithm can never successfully recover this original phase difference. Using the method of phase partitioning outlined in section 2.4 this is unlikely to happen in low resolution images (typical EPI images), although it has been noted to happen in high resolution images (voxel sizes of 1mm^3 or less) where small tracks of closely related phase voxels can be found in noisy areas, linking up larger, stable areas which should not be grouped in the same region. In practice we have overcome this by limiting the initial regions to be 2D (within plane only) for high resolution images. This dramatically reduces the likelihood of finding connecting tracks within noisy areas. In fact, we have not found any phase maps where two sizeable areas have been incorrectly connected

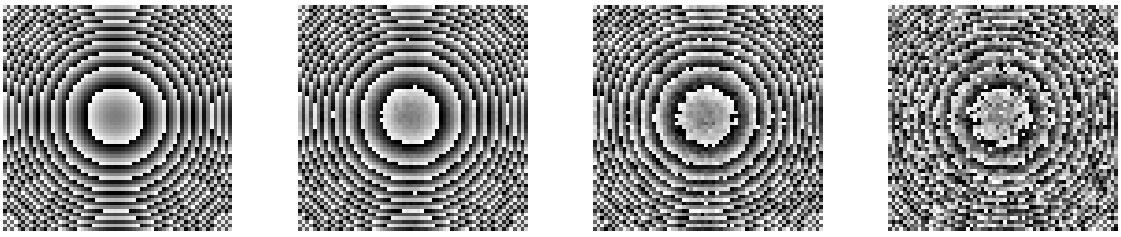

SNR = ∞ SNR = 5 SNR = 2 SNR = 1

Figure 11: Some of the synthetic phase images with added Gaussian noise (central slice shown).

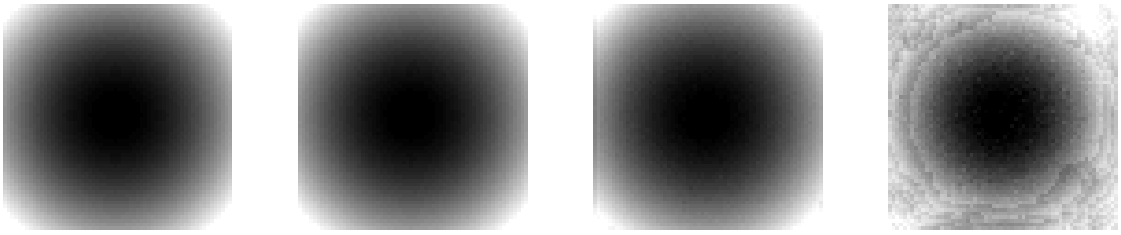

True Phase SNR = 5 SNR = 2 SNR = 1

Figure 12: Some of the unwrapping results of the synthetic phase images with added Gaussian noise (central slice shown). Note that the leftmost image, included for comparison, is the true phase which was used to generate the synthetic images.

using this method. Other methods, such as splitting regions into convex components (possibly by morphological operations) could also be pursued. However, the greater the number of initial regions, N, the longer the unwrapping takes to perform. In fact, applying a pre-processing stage which merged very small (single voxel or so) regions could result in a considerable speed-up for larger images.

The second potential weakness of the method is that the region merging algorithm attempts to solve the integer programming problem by "greedy" optimisation. That is, always choosing the best (least opportunity cost) candidate pair at each stage, and so can get stuck in local minima. Other optimisation methods could be used to solve the cost function minimisation instead, but although this method does not guarantee the optimal solution, it is very fast and we believe that it is a good compromise between speed and guaranteed accuracy. Moreover, unlike region-growing, the results of the region-merging algorithm do not depend on the location of any initial seed point.

Testing of the accuracy and robustness of the method was performed using a set of simulated images. These images contained various amounts of added Gaussian noise so that the performance of the algorithm at different SNRs could be determined. The results, quantified using the Mis-Classification Ratio, showed that the algorithm was extremely accurate and robust, for SNRs greater than 5. For most MR imaging methods the SNR is significantly better than this (our B0 mapping sequence had a SNR of 50), and therefore this method should be able to perform well with most MR phase data.

The software implementation of this phase unwrapping method is freely available (from www.fmrib.ox.ac.uk/fsl as a part (PRELUDE) of the EPI unwarping tool (FUGUE)) in both binary and source forms. It is hoped that this will enable more users in the MR community to develop applications where fast, automated phase unwrapping is necessary, and to participate in the evaluation and improvement of the current algorithm.

6 Acknowledgements

The author wishes to thank the European MICRODAB project for supporting this work.

References

[1] Li An, Qing-San Xiang, and Sofia Chavez. A fast implementation of the minimum spanning tree method for phase unwrapping. *IEEE Trans. on Medical Imaging*, 19(8):805–808, 2000.

[2] Leon Axel and Daniel Morton. Correction of phase wrapping in magnetic resonance imaging. *Medical Physics*, 16(2):284–287, 1989.

[3] A. Bhalerao, C-F Westin, and R. Kikinis. Unwrapping phase artifacts in 3d phase contrast angiograms. In *Proceedings of 1st Joint Meeting of CVRMED/MRCAS '97*, Grenoble, France, 1997.

[4] Curtis W. Chen and Howard A. Zebker. Two-dimensional phase unwrapping with use of statistical models for cost functions in nonlinear optimization. *Journal of the Optical Society of America A: Optical Image Science and Visualisation*, 18(2):338–351, 2001.

[5] Neng H. Ching, Dov Rosenfeld, and Michael Braun. Two-dimensional phase unwrapping using a minimum spanning tree algorithm. *IEEE Trans. on Image Processing*, 1(3):355–365, 1992.

[6] G.H. Glover and E. Schneider. Three-point dixon technique for true water/fat decomposition with b0 inhomogeneity correction. *Magnetic Resonance in Medicine*, 18(371–383), 1991.

[7] Igor P. Gurov and Denis V. Sheynihovich. Interferometric data analysis based on markov nonlinear filtering methodology. *Journal of the Optical Society of America A: Optical Image Science and Visualisation*, 17(1):21–27, 2000.

[8] Mark Hedley and Dov Rosenfeld. A new two-dimensional phase unwrapping algorithm for mri images. *Magnetic Resonance in Medicine*, 24(1):177–181, 1992.

[9] A. Herment, O. Mousseaux, E. ad Jolivet, A. DeCesare, F. Frouin, A. Todd-Pokrpoek, and J. Bittoun. Improved estimation of velocity and flow rate using regularized three-point phase-contrast velocimetry. *Magnetic Resonance in Medicine*, 44(1):122–128, 2000.

[10] P. Jezzard and R.S. Balaban. Correction for geometric distortion in echo planar images from B0 field variations. *Magnetic Resonance in Medicine*, 34:65–73, 1995.

[11] T. R. Judge and P. J. Bryanston-Cross. A review of phase unwrapping techniques in fringe analysis. *Optics and Lasers in Engineering*, 21:199–239, 1994.

[12] Rainer Krämer and Otmar Loffeld. Phase unwrapping for sar interferometry with kalman filters. In *Proceedings of the EUSAR'96 Conference*, pages 165–169, 1996.

[13] Zhi-Pei Liang. A model-based method for phase unwrapping. *IEEE Trans. on Medical Imaging*, 15(6):893–897, 1996.

[14] P. Munger, G.R. Crelier, and T.M. Peters. An inverse problem approach to the correction of distortion in epi images. *IEEE Trans. on Medical Imaging*, 19(7):681–689, 2000.

[15] G.L. Nayler, D. B. Firmin, and D.B. Longmore. Blood flow imaging by CINE magnetic resonance. *Journal of Computer Assisted Tomography*, 10:715–722, 1986.

[16] G.S. Sobering, Y. Shiferaw, P. van Gelderen, and C.T.W. Moonen. Quantitative measurement of the phase error for a simple and rapid phase-unwrapping algorithm. In *Proc. International Soc. of Magnetic Resonance in Medicine*, 1995.

[17] S. Moon-Ho Song, Sandy Napel, Norbert J. Pelc, and Gary H. Glover. Phase unwrapping of MR phase images using poisson equation. *IEEE Trans. on Image Processing*, 4(5):667–676, 1995.

[18] Jerzy Szumowski, William R. Coshow, Fang Li, and Stephen F. WQuinn. Phase unwrapping in the three-point dixon method for fat suppression mr imaging. *Radiology*, 192(2):555–561, 1994.

[19] Yi Wang, Debiao Li, and Mark Haacke. A new phase unwrapping method for water and fat separation in inhomogeneous fields. In *Proc. International Soc. of Magnetic Resonance in Medicine*, 1995.

[20] R. M. Weisskoff and T. L. Davis. Correcting gross distotion on echo planar images. In *Proc. International Soc. of Magnetic Resonance in Medicine*, page 4515, Berlin, 1992.

[21] Q.S. Xiang. Temporal phase unwrapping for CINE velocity imaging. *Magnetic Resonance in Medicine*, 5(5):529–534, 1995.