

# Non-rigid Motion Estimation and Spatio-temporal Realignment in FMRI

Peter R Bannister, J Michael Brady  
Robotics Research Group, Department of Engineering Science  
University of Oxford, Parks Road, Oxford, OX1 3PJ  
{prb,jmb}@robots.ox.ac.uk

Mark Jenkinson  
Oxford Centre for Functional Magnetic Resonance Imaging of the Brain  
John Radcliffe Hospital, Oxford, OX3 9DU, UK  
mark@fmrib.ox.ac.uk

## Abstract

*Existing approaches to the problem of subject motion artefacts in FMRI data have applied rigid-body registration techniques to what is a non-rigid problem. We propose a model which can account for the non-linear characteristics of movement effects, known to result from the acquisition methods used to form these images. The model also facilitates the proper application of temporal corrections which are needed to compensate for acquisition delays. Results of an implementation based on this model reveal that it is possible to correct for these effects, leading to accurate realignment and timing correction.*

## 1 Introduction


Functional Magnetic Resonance Imaging (FMRI) is a relatively recent development of Magnetic Resonance Imaging (MRI) and is a non-invasive technique which can be used to form images of neural activation. The technique can answer questions about the way in which the brain works and can also characterise deficiencies due to illness or injury. The modality is based on measurements of the magnetic behaviour associated with blood flow change due to metabolic activity which can be observed using an MRI scanner. Images are acquired using a multi-slice Echo Planar Imaging (EPI) protocol which raster scans k-space. While this achieves the necessary acquisition speeds for individual slices which are grouped along the z-axis to form volumes (typically a volume must be acquired every 3 seconds so that the temporal dynamics of the haemodynamic response can be captured), the signal-to-noise ratio (SNR) suffers and as a consequence image resolution must be lowered. This makes it much harder to reliably detect activation

and dynamic behaviour in a subject's brain. Also, because this detection is based on the statistical comparison of many scans of the same sections of the brain, the situation is further confounded by the adverse effect of subject motion at even very low levels of a few millimetres or less. This situation is exacerbated by the fact that clinical patients will generally move far more than cooperative volunteers while in the scanner. Echo Volumetric Imaging is the full extension of EPI to 3-D which means that the full volumetric image is acquired in a single shot. It is extremely difficult to carry out in practice and results in a very low resolution image (currently, about  $32 \times 32 \times 16$  voxels) with large levels of distortion due to an even lower SNR than is typical of multi-slice EPI.

Given that SNR considerations limit FMRI to stacked-slice acquisitions in practice, the problem of temporal offsets within a volume due to the successive acquisition of slices remains a significant confound to motion correction. Previous attempts to correct FMRI data for artefacts introduced during acquisition have considered spatial realignment and slice-timing correction as two distinct and separate stages in the processing chain. We discuss the problems inherent in such an approach in section 2, and go on to propose a new method, which we have called Temporally Integrated Geometric EPI Re-alignment (TIGER). We present the initial results from an implementation based on this model and conclude by describing a more sophisticated realisation based on this model which may be applied to real FMRI studies.

## 2 Existing Approaches To Motion Correction In FMRI

While the need for slice-timing is acknowledged widely, these corrections are not always performed in practice on



**Figure 1. In this figure the shaded bars, fixed to the same reference frame as the brain, indicate the acquisition timings associated within each slice. Even in the case of rigid-body motion, spatial realignment (in this case correcting a simple pitch of the head) will lead to a breakdown in the correlation between slice location used to determine slice-timing correction (marked with dotted lines) and associated acquisition timings (shaded bars).**


FMRI data. Applying the two corrections separately is convenient, not least because it facilitates the use of existing tools such as MCFLIRT [3] and SPM [2] for rigid-body motion correction along with separate temporal interpolation of each voxel time-course to re-shift slice-timings. There are fundamental errors in the assumptions underlying this distinction, however, regardless of the order in which the two steps are performed. Clearly, if no subject motion has occurred, it is sufficient simply to apply slice-timing correction as a series of temporal interpolations over each voxel time-course in turn, where the amount of shift is proportional to the temporal offset associated with the slice containing the voxel being considered. However a complete lack of subject motion is unlikely to occur in real data, so the interaction between motion and acquisition delays must be modelled in order to correct fully for the resulting artefacts in the data.

Assuming that motion correction is carried out before any temporal corrections, data which may not correspond to acquisition at a consistent point in time will be co-registered. If slice-timing correction is applied after the initial realignment, the corrected images will contain data from several discrete sample times within individual slices. This is because, in the general case of through-plane motion, spatial registration will realign the data so that intensity values from individual slices in scanner space are distributed across several slice locations in the corrected data. This is illustrated in figure 1. Specifically, if rigid-body realignment is performed, subsequent slice-timing will make the incorrect assumption that data within individual slices will have been acquired at the same time-point. In this situation, it is necessary to keep a record of the slice in which


the data were originally acquired, and then apply the appropriate timing correction, a step which is usually omitted. It might, therefore, seem obvious that temporal re-sampling should be carried out before motion correction. An obstacle to such a re-sampling is that in order to carry out slice-timing correction by temporal interpolation of a particular voxel, the time-course of that voxel must be known. If the subject has moved, there is no guarantee that a voxel in object space will be in constant alignment with a voxel in scanner coordinates. This creates a cyclic problem where motion correction is needed in order to determine slice-timing before motion correction. Recognising this interdependency is by no means original; but prior to this work there have been no attempts to utilise knowledge about the acquisition sequence in order to correct the situation.

The situation is worse when considering a voxel on an intensity boundary (for example on the interface between two tissue types or on the perimeter of an activating region), since in such a case the voxel described by a particular set of scanner coordinates may rapidly switch between two different intensity regions in object space, thus creating a physically implausible (uncorrected) time-course on which to base temporal interpolation.

In general, if a rigid-body model is assumed, the motion correction stage will also ignore the fact that there may not be a parallel correspondence between the slices from different volumes if motion has occurred during or between activations (That is, it is no longer correct to assume that the relative movement between slices is consistent across different volumes), such as the example in figure 2. Thus a rigid-body spatial realignment will inevitably attempt to compute pair-wise voxel comparisons on data which originates from different spatial locations in the subject's brain. It is also possible that the slices will remain unaligned, even when the volumetric optimisation has reached a minimum. In conclusion, movement throughout a scan will lead to different displacements in individual slices. This is ignored by volumetric corrections which assume a rigid-body transformation over the entire volume. These observations show that the separate application of slice-timing and motion correction cannot accurately account for motion artefacts in FMRI. For this reason, we propose an integrated approach to these two corrections that is able to cope with the potential spatial non-linearities in the data. If the subject has moved during acquisition of that volume, a quite likely possibility, as previously noted, is that there will be redundancy at some locations yet sparseness and/or voids at other points of the volume. This is due to a particular location in the object having moved sufficiently quickly between slice acquisitions to have been captured either more than once or not at all, as shown in figure 3. This observation is not new in the FMRI literature: incorrect assumptions about spatial distribution of slices have already been identified as the cause of



**Figure 2. Local distribution of slices within different volumes may force the motion correction problem to require a non-rigid solution.**



**Figure 3. Movement during the acquisition of a volume (in this case a simple nodding motion is depicted) may cause acquisition redundancy and voids.**

errors in the realignment of FMRI data [6, 1]. At a more basic model level, because subject motion may occur more markedly during some volume acquisitions than in others, the initial motion correction stage may be trying to infer an affine relationship between two volumes which do not have the same overall morphology (such as the example shown in figure 2). Another flaw in applying standard intensity-based volumetric registration in this situation is that each volume will contain a unique spatio-temporal distribution of voxels. While all the voxels within each scanner slice are temporally aligned, it is not possible to say before motion estimation which physical locations in object space these voxels relate to or whether this object-to-scanner correspondence is consistent across all volumes. A method that registers individual slices to a single anatomical volume has been proposed in [4] but this does not take into account the timing issues and requires a cross-modal registration. The latter condition is particularly confounding as any 2-D to 3-D registration of this kind is bound to be poorly constrained

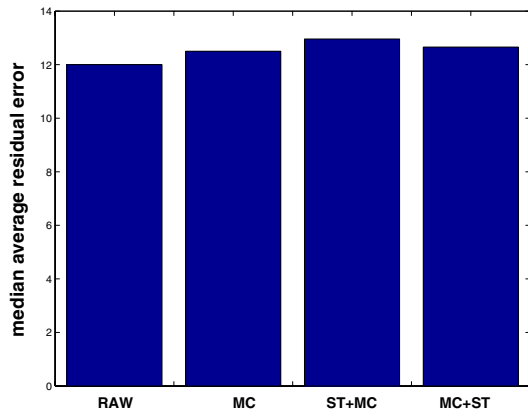
(the most pathological example of this would be to trying to match a circle to a unique plane within a sphere) and even more so in a cross-modal application where one of the modalities is low resolution EPI, with all its inherent distortions.

In order to demonstrate the potentially large inaccuracies resulting from the use of separate spatial and temporal realignment of FMRI data, a sample data-set was constructed according to a known motion design. The images simulated subject motion in a ‘nodding’ pattern where each volume contained rotational motion about the y-axis (left to right) projected through the centre (Taken to be the centre of gravity computed over the intensity values of the template image) of the image and ranging from  $-2$  to  $+2$  degrees over each volume.

Using a single high-resolution EPI image (containing  $192 \times 256 \times 128$  voxels, representing a FOV of  $192 \times 256 \times 163.8$  millimetres), this motion was applied incrementally to each set of 6 slices to simulate the intra-volume motion found in many real data-sets. In other words, if the motion over a volume was determined to be 2 degrees, the first 6 slices were pitched forward by 0.095 degrees, the second set by 0.19 degrees and so on up to a rotation of 2 degrees for the final ( $21^{st}$ ) set of 6 slices.

Even if the nature of the motion were known accurately, it would no longer be possible to adopt the parallel spatial correction approach that could be used for images containing inter-slice motion only, since the number and source of the voxels intersecting each slice of the volume change between time-points. Figure 4 shows the effects of such corrections. Using the median residual intensity error measured between the original un-corrupted data and the corrected data, results show that for through-plane motion, not only does rigid-body motion correction and slice-timing fail to correct the motion-related distortions, the results of the corrections are actually *worse* than no correction at all. Clearly, both in the case of through-plane and in the earlier example of in-plane motion, standard motion correction techniques using rigid-body registration and simple voxel-wise temporal interpolation for slice-timing cannot accurately compensate for the types of motion typically encountered in FMRI. Based on the earlier observations regarding the need for integrated spatio-temporal motion correction, and the clear demonstration of the shortcomings of existing techniques, a new model for subject motion in FMRI is required.

Estimation of slice-timing and motion inevitably relies on voxel intensities acquired through a relatively low resolution modality (including activation-related signal changes of interest), so a spatio-temporal scheme offers the additional advantage that it may be possible to combine information over a large number of images to give more robust conclusions than those normally available from pair-wise



**Figure 4. Comparison of relative accuracies, measured as median average residual variance, of separate spatial and temporal corrections when applied to data known to contain inter-slice movement within individual volumes. From left to right: Uncorrected data (RAW), MCFLIRT rigid-body motion correction (MC), slice-timing correction then MCFLIRT (ST+MC), MCFLIRT then slice-timing correction (MC+ST)**

realignment of EPI volumes.

### 3 Spatio-temporal Model

The new model described in this section is intended to bypass the shortcomings of the primarily sequential methods described above. The acquisition process is modelled as a set of discrete time-points within each volume, one corresponding to each slice, so as to allow changes in the orientation of the head for each slice. A practical approximation to this process is to express the orientation of each slice as a fraction of the estimated gross motion contained in the volume. This is done by estimating the motion between two volumes and decomposing this motion into a set of evenly-spaced rotations about a common axis [7]. It should be noted that the associated increase in DOF is only a total of 6 additional parameters for the whole time-series. Thus the assumption of an even distribution of slices along some postulated rotational axis is the key to keeping the subsequent optimisation problem manageable. This is discussed in more detail below. Due to the increase in the number of DOF for the registration problem, however, it is necessary to devise a method which is able to efficiently cope with the permitted flexibility in the data and which applies constraints and some global cost invariant over the data as a whole.

The proposed method can be divided into 3 separate sections: the initialisation stage, which transforms the images from volumetric data into distributed, discrete slices; the cost function optimisation, which serves to refine the initial image pose estimates by exploiting spatio-temporal similarities in the data; and, finally, the temporal interpolation which implements spatially-integrated slice-timing corrections and includes the estimation of intensity values associated with voxels which are deemed not to have been acquired as a result of inter-slice movement.

#### 3.1 Slice Pose Estimation

Given two volumes which have been acquired one after the other in an FMRI time-series, it is possible to estimate the total motion of the subject’s head within the first volume  $V_n$ , as the rigid-body transformation  $\mathbf{T}_n$  which relates  $V_n$  to the second volume,  $V_{n+1}$ . This is the approach adopted by current rigid-body motion correction schemes. Although this estimate can only be an approximation in the case of an independent-slice model, since the movement is intra-scan rather than volumetric, it provides a useful initial approximation which can then be refined progressively throughout a subsequent optimisation process. It is assumed that  $T_n$  is smoothly varying (sudden intra-scan movements within the volume are possible; but in this first study the model will avoid this additional level of complexity in the interests of testing the underlying hypothesis of significant inter-slice motion) so that the transformation can be decomposed into a rotation  $\theta$  about some common axis [7]  $\mathbf{q} = q_x\mathbf{i} + q_y\mathbf{j} + q_z\mathbf{k}$ . This formulation, or ‘screw decomposition’, makes it possible to construct a set of incremental rotations about  $\mathbf{q}$  by allowing the angle of rotation to range as  $\theta_s = s \cdot \Delta\theta$  where  $\Delta\theta = \frac{\theta}{S}$ , where  $s$  is the current slice number and  $S$  is the number of slices. Similarly the translational component  $\mathbf{t} = t_x\mathbf{i} + t_y\mathbf{j} + t_z\mathbf{k}$  can be applied in increasing fractions up to the full (rigid-body) value.

These slice-based transformations can be written as  $\mathbf{T}_{n,s}$  where  $s = 1, \dots, S$ . By applying each incremental transformation to the corresponding plane in scanner space it is possible to form an independent-slice representation of the volume. This is illustrated in figure 5. This decomposition provides a more realistic interpretation of the data in each slice with respect to the subject’s head than the standard stacked-slice model [4, 5]. Furthermore, because the distribution of slices is determined by a limited number of parameters (specifically  $\mathbf{q}$ ,  $\theta$  and  $\mathbf{t}$  - all of which are derived from the rigid-body matrix  $\mathbf{T}_n$ ) it is possible to update the orientation of the slices by varying only  $\mathbf{T}_n$ , thus limiting the complexity of a correction scheme based on this model. This is an important aspect of the model, given the dimensionality of the spatio-temporal data and the number of sample points.

An important corollary to this observation is that by altering  $\mathbf{T}_n$ , and thus the slice distribution within the test volume  $V_n$ , the distribution of slices within the adjacent reference volume  $V_{n+1}$  will also be affected. This may be thought of as an ‘accordion’ effect where the first slice of the test volume and the final slice of the reference volume are constrained while the midpoint slice is varied, thus altering the spacing of slices within each volume. An illustration of this relationship between volumes is given in figure 6. For example, if  $\mathbf{T}_n$  relates slice 1 of  $V_n$  to slice 1 of  $V_{n+1}$  and  $\mathbf{T}_{n+1}$  relates slice 1 of  $V_{n+1}$  to slice 1 of  $V_{n+2}$ , then if slice 1 of  $V_{n+1}$  moves by  $\mathbf{P}$ , the transformation  $\mathbf{T}_n$  followed by  $\mathbf{P}$  will map slice 1 of  $V_n$  to the (now perturbed) slice 1 of  $V_{n+1}$ . If the updated rigid-body transformation associated with  $V_n$  is written

$$\mathbf{T}_n^{\mathbf{P}} = \mathbf{P} * \mathbf{T}_n \quad (1)$$

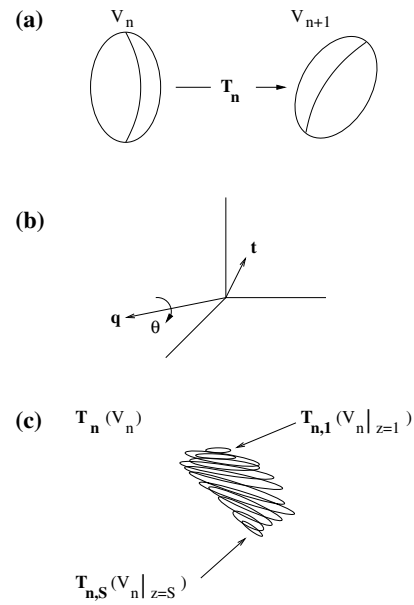
where  $\mathbf{P}$  describes the parameter updates determined by optimisation (the *perturbation matrix*), then the rigid-body transformation  $\mathbf{T}_{n+1}$  describing the orientation of the next volume  $V_{n+1}$  (which is based on its relative position with respect to  $V_n$ ) must be adjusted. To understand this, consider the relationship between the newly perturbed slice 1 of  $V_{n+1}$  and its original orientation in the example above, which equals  $\mathbf{P}^{-1}$ . Therefore, the updated  $\mathbf{T}_{n+1}$  becomes  $\mathbf{P}^{-1}$  followed by the original  $\mathbf{T}_{n+1}$ . According to the relationship between  $V_n$  and  $V_{n+1}$  described in section 3.1, this update can be written:

$$\mathbf{T}_{n+1}^{\mathbf{P}^{-1}} = \mathbf{T}_{n+1} * \mathbf{P}^{-1}. \quad (2)$$

This ensures that the product  $\mathbf{T}_{n+1} * \mathbf{T}_n$  is the same, independent of the application of  $\mathbf{P}$ , which is necessary as the product of the two transformations describes the mapping between slice 1 of  $V_n$  and slice 1 of  $V_{n+2}$ , neither of which have moved. By performing this slice-wise decomposition for every image in the series, it is possible to obtain an initial approximation to the data as a more realistic representation of the subject’s head in scanner coordinates, rather than a volume corrupted by a single volumetric transformation (illustrated in figures 2 and 3).


### 3.2 Cost Function Optimisation

Having constructed the initialised distributed-slice time-series by applying the model from section 3.1 to each volume, the next stage in correcting the data is to construct a cost function which attains a minimum when the 4-D time-series is aligned. The assumption underlying this stage, as with conventional intensity-based rigid-body schemes designed to work on fMRI data, is that there will be a small level of intensity variability restricted to a small number



**Figure 5. By expressing the rigid-body transformation  $\mathbf{T}_n$  relating volume  $V_n$  to volume  $V_{n+1}$ , shown in (a), as a rotation  $\theta$  about a single axis  $\mathbf{q}$  and a translation  $\mathbf{t}$ , depicted in (b), it is possible to determine a set of transformations  $\mathbf{T}_{n,s}$  which describe the local movement of  $n$  individual slices. These can be applied to each slice of volume  $V_n$ , written  $V_n|_{z=s}$  where  $s = 1, \dots, S$ , to give volume  $\mathbf{T}_n(V_n)$  shown in (c), where each slice is transformed by progressively greater fractions of the original transformation,  $\mathbf{T}_n$ . This provides a more accurate model of the data acquired on a slice-by-slice basis within each volume.**

of voxels. These voxels are assumed to correspond either to regions of the subject’s brain which have activated under the stimulus of the experiment or those areas which have been affected by an acquisition artefact. The intensity of the remaining voxels should remain broadly constant in object coordinates so the correction should seek a realignment over the 4-D data which minimises this variance. A Least Squares metric is used to reduce the likelihood of local movement patterns dominating the evaluation of a minimum. An example of such a situation would be where there is no appreciable motion throughout two adjacent volumes but where those two volumes are significantly offset from the rest of the time-series. By generating a reference image from the entire remainder of the time-series, there is a much smaller possibility of the presence of an adjacent (similarly misaligned) volume driving the evaluation of the cost function. When interpolating intensity values after reg-



**Figure 6. At each optimisation step for a volume  $V_n$ , a perturbation  $P$  is applied to  $T_n$  and the new slice distribution  $T_{n,s}^P(V_n)$  tested. Throughout this optimisation step, the end slices of the volume pair  $[V_n, V_{n+1}]$  shown in light grey are kept fixed. In effect what  $P$  does is to vary the mid-point slice between these two volumes, shown in dark grey. Thus, for every perturbation applied to  $V_n$  it is necessary to apply the inverse  $P^{-1}$  to  $T_{n+1}$  (giving  $T_{n+1}^{P^{-1}}(V_{n+1})$ ) to maintain the constraint on the end (light grey) slices.**

istration, sinc interpolation is applied to the original data. Currently multi-dimensional optimisation using the golden section method is used to search for cost function minima, with the order of rotation searches being permuted to avoid local minima traps. More sophisticated optimisation strategies are considered in section 5.

In order to enforce the anatomical constraint, that is, that the realigned data should represent the data in object space, not simply in a pose which satisfies the cost function, it is necessary to include at least one volume which contains little or no inter-slice movement. While this may seem to be a very specific case of the likely input data, it is relatively easy to select such a volume from within the time-series by examining the relative RMS difference between each volume, as estimated by a rigid-body scheme. In the case where the relative movement between two volumes is very low, an assumption can be made that little or no inter-slice movement occurred over the course of the acquisition of either. Thus the images will hopefully contain a physically-consistent representation of the brain, even if its global pose is not entirely co-incident with the coordinate axes.

Either of these two internally static images can then be used as a canonical reference volume for the spatio-temporal realignment. An appropriate modification to the realignment scheme is to select the static volume using the procedure described above and to then allow the adjacent volume to perturb according to the model shown in figure 6. Once this second volume has been realigned with respect to the first static volume, it is then possible to use both these volumes as a reference image in the variance cost for the realignment of the third volume. In this way, the reference image is built up iteratively but anatomical constraints

are enforced by virtue of starting with a volume which has been measured as the most likely to contain little in the way of inter-slice movement. In addition to the intermediate interpolation described above, care must also be taken when performing the temporal re-sampling applied to each voxel time-course after spatial realignment. The underlying acquisition process should ensure that once optimisation is complete, all the voxels within each slice of scanner space should contain intensity values acquired at the same time. Unlike existing rigid-body schemes which are inappropriately used with slice-timing, the spatial realignment of the data resulting from the proposed slice-based model allows the application of appropriate temporal interpolation. However, care must be taken to deal with situations where there may be no acquisitions for a given voxel at a particular volume's time-point. In order to apply a shift to reflect the temporal offset of each slice it is necessary to have at least one sample for each voxel location and time-point. If reconstruction of the data has concluded that the particular voxel was 'missed' during acquisition as a result of subject motion, its intensity value is estimated using Hermite splines. By maintaining a record of the slice from which each transformed intensity value originated, it is then trivial to apply the appropriate temporal shift using sinc interpolation.

## 4 Testing and Results

The method described above has been implemented and applied to synthetically-generated test data, based on the high resolution EPI image (described in section 2) to minimise the effect of interpolation artefacts when re-aligning the data to simulate motion. Two datasets were constructed which contained a static initial volume followed by 20 volumes which contained motion. The first dataset simulated a nodding motion distributed over the individual slices of each volume up to a maximum of 2 degrees, while the second dataset contained a nodding design of the same magnitude, centered on the centre of mass of the original image.

An initial analysis of the cost function behaviour on this test data revealed that the cost function exhibited slightly offset minima, even when the images were perfectly realigned using the independent slice model. We believe that this discrepancy is due to the extremely high levels of accuracy which are required for the spatial interpolation of sparse data. Testing the spatio-temporal optimisation on the shaking dataset revealed that the realignment scheme was able to reduce the average RMS error within each volume from an un-corrected value of 1.9331 millimetres to 0.2829 millimetres. The deviations from ground truth in the parameter estimates were noted to consistently follow the observations made in the initial cost function analysis, verifying the optimisation.

In order to test the accuracy of the spline interpolation,

	Uncorrected	Post-realignment	With Splines
Shake	14.2082	7.2197	7.1796
Nod	16.0322	12.3656	12.2599

**Figure 7. The table shows the decreased residual error (measured as the median residual error between the reference image and the re-aligned data) which results from time-course correction using splines when compared to spatial realignment alone**

it is necessary to compare the median residual error for the data before realignment (computed over all intensity values), after spatial realignment but before spline reconstruction, and finally again after the spline reconstruction. For the nodding and shaking data-sets, the results are given in figure 7.

Qualitative examination of the realigned data revealed that very little perceptible global motion remained within each slice. Of potentially greater interest was the presence of local detail variation between time-points. This reflects the changing role played by interpolation across a range of slice poses and further reinforces the earlier hypothesis that interpolation is responsible for much of the error with respect to zero motion observed in the cost plots. While deviations from zero were still observable in the parameter estimates for the nodding data, the realignment was able to reduce the average RMS error from 2.1863 millimetres to 0.7451 millimetres. The current implementation takes approximately 72 hours to process a 20 volume time-series on a 64-bit ES40 667MHz Alpha platform, approximately 80% of which is dedicated to the additional parameter permutation searches.

## 5 Future Work

Given that the dominant source of error in the above implementation was interpolation accuracy, an improvement in the spatial estimation of the TIGER framework is likely to result from the inclusion of a more sophisticated interpolation technique. The issue of accuracy of interpolation has been highlighted repeatedly as a fundamental factor in the reliability of motion correction, even more so than in conventional medical image registration [3]. It is possible that advanced forms of spatial interpolation, such as Markov Random Fields could be adapted to include the robust estimation of inhomogeneous sampling rather than having to use a separate spline interpolation stage. An alternative approach to the problems caused by inaccurate interpolation might alternatively involve deriving the interpolated values from the original k-space data. By interpolating in k-space,

before the data is Fourier transformed into cartesian image space, it might be possible to achieve more accurate interpolation by avoiding the errors associated with interpolating truncated data in the form of images.

We plan to test the improved implementation on data which contains realistic levels of noise added to the motion corrupted images. Pending this confirmation of the method's robustness, testing will be extended to real FMRI data which is known to contain intra-slice motion.

## 6 Acknowledgements

The authors would like to thank Christian Beckmann, Stuart Clare, Stephen Smith and Mark Woolrich for input and the MRC and EPSRC for funding this research.

## References

- [1] R. Cox. Motion and Functional MRI: Informal notes for the Boston '96 workshop on functional MRI, 1996.
- [2] K. Friston, J. Ashburner, C. Frith, J.-B. Poline, J. Heather, and R. Frackowiak. Spatial registration and normalization of images. *Human Brain Mapping*, 2:165–189, 1995.
- [3] M. Jenkinson, P. Bannister, J. Brady, and S. Smith. Improved optimisation for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2):825–841, 2002.
- [4] B. Kim, J. Boes, P. Bland, T. Chenevert, and C. Meyer. Motion correction in fMRI via registration of individual slices into an anatomical volume. *Magnetic Resonance in Medicine*, 41:964–972, 1999.
- [5] L. Muresan, R. Renken, J. Roerdink, and H. Duifhuis. Position-history and spin-history artifacts in fMRI time-series. In A. V. Clough and C.-T. Chen, editors, *Proc. Medical Imaging 2002: Physiology and Function from Multidimensional Image*, volume 4683, pages 444–451. Proc. SPIE, 2002.
- [6] D. Noll, F. Boada, and W. Eddy. Movement correction in fMRI: The impact of slice profile and slice spacing. *International Soc. of Magnetic Resonance in Medicine*, page 1677, 1997.
- [7] R. Paul. *Robot manipulators*. Series in Artificial Intelligence. The MIT Press, 1981.