Pre-reconstruction Rigid Body Registration for Positron Emission Tomography

Peter Nordberg[1]
phgn@robots.ox.ac.uk
Jerome Declerck[2]
jerome.declerck@siemens.com
Michael Brady[1]
jmb@robots.ox.ac.uk

[1] Wolfson Medical Vision Lab
Department of Engineering Science
University of Oxford

[2] Siemens Molecular Imaging
23-38 Hythe Bridge St.,
Oxford, OX1 2EP

Abstract

Abrupt motions pose particular problems for Positron Emission Tomography because any mismatch between the subject's position during the attenuation correction scan and the PET acquisition causes the attenuation correction step during reconstruction to introduce artefacts; aquisitions with abrupt motions are often discarded. This paper adaptats a rigid body registration algorithm from the CT literature, expanding upon the details of implementation and showing its applicability to PET using realistic simulated data. The method is of special interest as it operates on the projection (sinogram) data, thus avoiding the need to reconstruct images. Given a scan with a change of position at a known time the motion can be estimated, and corrected for, before reconstruction.

1 Introduction

PET is an increasingly important imaging modality because it images specific aspects of physiology and metabolism in vivo. The quality of the images continues to improve with better hardware and reconstruction methods but, as they do so, the effects of abrupt motion have become a major limitation on the quality of the information ultimately available. Motion causes two problems for PET: the first is the blurring of the image, the second follows from the need to perform attenuation correction. With the advent of hybrid PET-CT machines attenuation is estimated from the CT scan acquired just before the PET scan. Any motion after the CT scan means that emissions from some regions will not undergo the estimated level of attenuation and the correction step will introduce artefacts [3].

This paper presents an adaptation to PET of a rigid body registration method from the work of Fitchard *et al.* [1, 2] for CT and expands on the details of implementation. The method is of interest because it allows for rigid body registration directly from the projection data (the sinogram), without the need for reconstruction – figure 1 (a). As well as reducing computational cost, this avoids the need to make choices regarding the reconstruction algorithm that is known to influence the registration. Furthermore, the method naturally allows for correction before reconstruction, thereby minimising attenuation correction errors once aligned to the same position as the attenuation scan. It is important to perform any registration method for PET on non-attenuation corrected data, lest the registration be influenced by the artefacts.

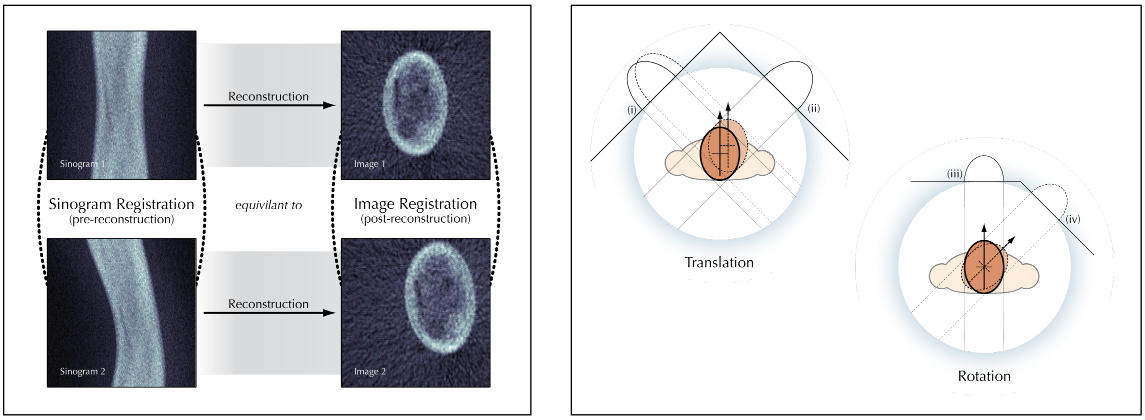

(a) (b)

Figure 1: Illustration of (a) the concept of pre-reconstruction registration and (b) how rigid body motions impact the projections that make up the sinogram: (i) & (ii) the extremes of translation, (iii) & (iv) rotation.

2 Method

If an image is rotated relative to its reference position, a rotational cross-correlation of the template (i.e. rotated) image with the original (reference) image attains a maximum that corresponds to the angle of rotation. Doing the rotational cross-correlation after performing a Fourier transform makes this shift invariant, since shifting an image does not change its frequency content. More precisely: rotating an image changes the distributions of horizontal and vertical frequency content such that the 2D frequency spectrum is rotated by an angle equal to the image's rotation and the rotational cross correlation of two frequency spectra corresponds to the cross correlation of the two images, ignoring any translation.

The need to reconstruct an image can be avoided through use of the Fourier central slice theorem. Each row of a sinogram (each projection angle) is Fourier transformed separately and stacked up in the same order to produce a new array, in which rows correspond to angles and columns to frequency components. By the Fourier central slice theorem, these columns correspond to rings in the 2D frequency spectrum of the image that would result from reconstruction; the 2D rotational cross-correlation can be achieved by cross-correlating these columns made from the reference and template sinograms.

The rotation can be removed by re-indexing the rows of the sinogram, in effect redefining which row corresponds to the zero-angle projection.

After the rotation is removed the translation can be estimated. This is done by comparing the reference sinogram and the rotation-registered template sinogram (the 'rotated sinogram'), again by cross-correlation. If an object is translated in a given direction then – under parallel beam projection, as is the case for PET – the projection in that direction will not change. The perpendicular projection will be shifted by an amount equal to the displacement and, between the two, the shift will vary sinusoidally with projection angle (figure 1).

If each row of the rotated sinogram is cross-correlated with the corresponding row of the reference sinogram, the shift that best matches the two rows should vary sinusoidally; the phase (relative to projection angle) and magnitude are determined by the direction and magnitude of the translation. As pointed out in the original papers on CT, the frequency of the sinusoid is the fundamental. Therefore, taking a Fourier transform of the estimates allows the phase and amplitude to be recovered easily; doing so also rejects much of noise and errors in the individual cross-correlation maxima.

Finally, the motion is removed by applying the reverse shifts. These should be from the sinusoidal pattern so the correction corresponds to a consistent rigid body motion. The result is a registered sinogram that, once reconstructed, will produce an image in the same position as that of the reference sinogram. The registered and reference data can be combined and a single image reconstructed using all the data.

It should be noted that there is no reason why the two segments of data being registered need to be of the same duration: the location of the maximum of a cross-correlation function depend only on the pattern of the two input functions, not their relative scale.

3 Implementation Details

While estimating the rotation, all the columns being cross-correlated should, in theory, yield the same result. In practice, the different frequency components contain different information, with the low frequency components being more reliable as they contain the information about large scale structures. For higher frequency components, the rotation estimated becomes unstable (figure 2). For this reason the rotation estimate used is based on an average of the stable ones (ignoring the DC and fundamental).

The cut-off between the stable and unstable estimates is chosen by considering the variance of the stable and unstable estimates. As there are two obvious populations characterized by their consistency (measured by variance^{-1}) and inconsistency (equated to the variance) the best cut-off will be the one that makes each population most like itself. This is the point at which the ratio of the variance of estimates based on higher frequency components to those from lower frequencies is maximal (figure 2). A minimum of 6 estimates are always averaged; this constraint is usually met by the chosen cut-off frequency anyway.

Should the mean variance of the higher frequency estimates not be higher than the mean variance of the lower frequency estimates then there is no clearly defined difference between stable and unstable estimates: this suggests that all the estimates are just noise and so the rotation estimate is set to zero.

The implementation of the translation step is simpler, the only addition to the method in [1] has been a validity check, mirroring that for the rotation. As the only frequency component (ignoring windowing effects) corresponding to the translation is the fundamental, if its magnitude is not significantly greater than the others it is assumed better to ignore any estimated motion. ‘Significantly greater’ is taken to be three standard deviations of the other components above their mean. This assumes a normal distribution – tenuous given the magnitudes cannot be negative. However, given that this step is only a fail safe it is deemed to be an acceptable approximation.

4 Demonstration

PET-SORTEO [4], a Monte-Carlo based software simulator was used to create sinograms with realistic properties, emulating those from an ECAT Exact HR+ scanner (Siemens Medical Solutions, Knoxville, Tennessee, USA), but with known positions. Several in-plane slices were combined to form a single 2D sinogram with the desired number of counts corresponding to a cross section through the centre of the phantoms (chosen to have a constant cross section). The phantom was approximately 24cm across as figure 3, which shows the

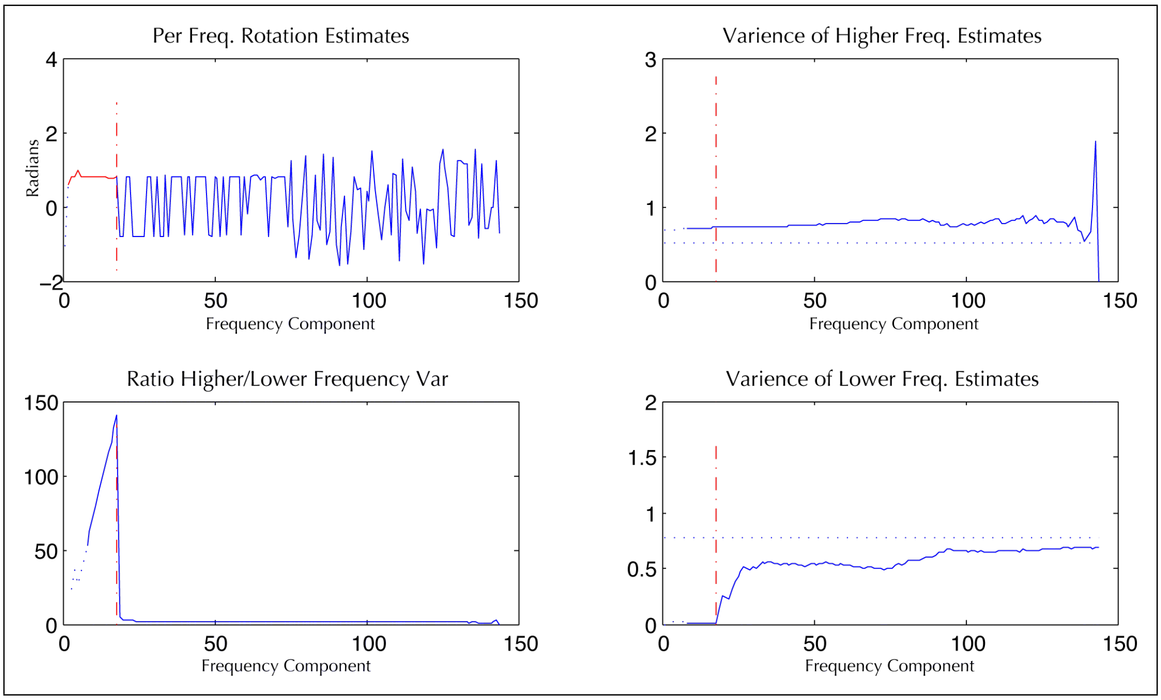

Figure 2: *Top Left:* The estimated rotation angle using each frequency component, with the selected cut-off between stable and unstable estimate. *Top Right:* The variance of the higher frequency components - for each point on the x-axis the value is the variance of all estimates from higher frequencies. Also shown is the cut-off ultimately chosen (vertical line) and the mean of the lower frequency variances (horizontal line). *Bottom Right:* Likewise, but for the variances of frequency lower than the value on the x-axis. *Bottom Left:* The ratio of variance for each possible choice of cut-off.

impact of the registration on the reconstruction. The angular resolution of the sinogram was 1.25 degrees and, after reconstruction, the voxel spacing is 2.25mm.

Using 12 motions with translations ranging from 0 to 8 cm in various directions and rotations between 0 and 45 degrees the average magnitude of errors – of Euclidean position in mm and of rotation in degrees — of the estimated motions were: 1.22 & 1.54 (with an average of 56 thousand counts per frame, before attenuation correction), 1.31 & 2.16 (80 Kcounts) , 1.02 & 1.32 (133 Kcounts), 1.09 & 0.74 (241 Kcounts), 1.01 & 0.56 (471 Kcounts).

5 Discussion

The work presented in this paper has been intended as a demonstration of the method with realistic simulated PET data and it works well, even with low count numbers. Validation and comparison against other methods, especially those operating post-reconstruction, remains to be done, but the absence of any parameters to tune illustrates the advantage of performing analysis pre-reconstruction.

The method's inherent limitation is that it is only applicable to rigid body motion – however, for neurological studies this would not be a problem and this has been seen as the main application of this method during this work. Although a 2D implementation has been used in this paper, as discussed in [2] it could be straightforwardly extended to 3D. There is also the assumption that the cross-correlation will produce the correct maximum. It is

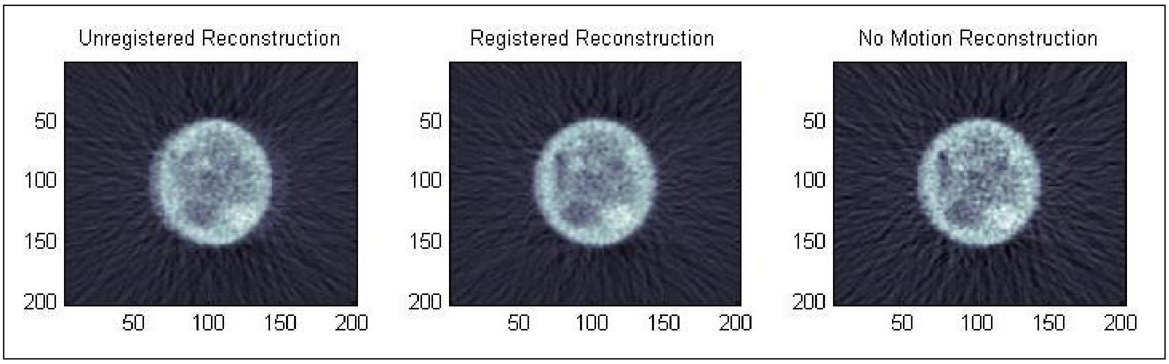

Figure 3: Illustration of the effect of registration on the images. Reconstruction is by filtered back-projection, a delayed window is subtracted to remove randoms, but no scatter correction is done and additional smoothing is applied. Here there is a 4cm shift to the right with about 940 thousand counts overall. The colour range is the same for all 3 images.

is stated that this is valid for typical CT images in Fitchard's work it, but this should be investigated again for PET.

Lastly an elegant means of deducing when the subject moved remains an open question – but given the speed and resilience to low count data it is not inconceivable to use this method prospectively, blindly dividing a scan into shorter frames and registering them together.

References

[1] E E Fitchard, J S Aldridge, P J Reckwerdt, and T R Mackie. Registration of synthetic tomographic projection data sets using cross-correlation. *Physics in Medicine and Biology*, 43(6):1645–1657, 1998. URL http://stacks.iop.org/0031-9155/43/1645.

[2] E. E. Fitchard, J. S. Aldridge, P. J. Reckwerdt, G. H. Olivera, T. R. Mackie, and A. Iosevic. Six parameter patient registration directly from projection data. *Nuclear Instruments and Methods in Physics Research*, 421(1-2):342–351, January 1999. doi: 10.1016/S0168-9002(98)01236-4.

[3] K. Lance Gould, Tinsu Pan, Catalin Loghin, Nils P. Johnson, Ashrith Guha, and Stefano Sdringola. Frequent Diagnostic Errors in Cardiac PET/CT Due to Misregistration of CT Attenuation and Emission PET Images: A Definitive Analysis of Causes, Consequences, and Corrections. *J Nucl Med*, 48(7):1112–1121, 2007. doi: 10.2967/jnumed.107.039792. URL http://jnm.snmjournals.org/cgi/content/abstract/48/7/1112.

[4] A. Reilhac, C. Lartizien, N. Costes, S. Sans, C. Comtat, R.N. Gunn, and A.C. Evans. Pet-sorteo: a monte carlo-based simulator with high count rate capabilities. *Nuclear Science, IEEE Transactions on*, 51(1):46 – 52, feb. 2004. ISSN 0018-9499. doi: 10.1109/TNS.2003.823011.