Data-Driven Motion Correction in Single Photon Emission Computed Tomography of the Brain

A thesis submitted in fulfilment of the requirements for the award of the degree

Masters by Research

from

UNIVERSITY OF WOLLONGONG

by

Andre Z. Kyme, BSc (Hons.)

Medical Radiation Physics

2004

CERTIFICATE OF AUTHORSHIP / ORIGINALITY

I, Andre Kyme, declare that this thesis, submitted in fulfilment of the requirements for the award of Masters by Research, in the Department of Medical Radiation Physics, University of Wollongong, is wholly my own work unless otherwise referenced or acknowledged. The document has not been submitted for qualifications at any other academic institution.

Andre Z. Kyme

30th January 2004

ACKNOWLEDGEMENTS

I wish to acknowledge the support, technical and otherwise, received from others during the course of this thesis. Many thanks to my supervisor Brian Hutton for his direction, eagerness to discuss and think through problems, for his insights, guidance and encouragement. I am very thankful to my colleagues Rochelle Hatton and David Skerrett for the countless times they have assisted in helping me think through various aspects of the work and also for solving many of my IT related struggles. Thanks go to Leighton Barnden and Setayesh Behin-Ain for their willingness to help troubleshoot and for acquiring all of the real experimental data used in this work. Finally, thanks go to Roger Fulton for his provision of software and his insights and expertise on motion correction in SPECT, to Michael Braun for many helpful discussions of a mathematical sort, and to ANSTO for their generous support which enabled me to participate in a thoroughly stimulating overseas conference during the course of my research.

A. K., Sydney, January 2004.

CONTENTS

FIGURES

FIGURE

TABLES

TABLE

LIST OF ABBREVIATIONS

2D	Two-Dimensional
3D	Three-Dimensional
COM	Centre-of-Mass
CT	Computed Tomography
DOF	Degree-of-Freedom
DSA	Digital Subtraction Angiography
EM	Expectation Maximisation
fMRI	Functional Magnetic Resonance Imaging
FOV	Field-of-View
FWHM	Full Width at Half Maximum
LOR	Line of Response
MI	Mutual Information
ML	Maximum Likelihood
MRE	Mean Registration Error
MRI	Magnetic Resonance Imaging
MSD	Mean of Squared Differences
MSDR	Mean Squared Difference Improvement Ratio
OS	Ordered-Subsets
PET	Positron Emission Tomography
RAMLA	Row-Action Maximum Likelihood Algorithm
RBI	Re-scaled Block-Iterative
rCT	Reprojected Computed Tomography
SPECT	Single Photon Emission Computed Tomography
SSD	Sum of Squared Differences

ABSTRACT

Introduction: Patient head motion is a well-recognised problem in single photon emission computed tomography (SPECT) of the brain. Motion occurring between or during the acquisition of projections can lead to reconstruction artifacts that compromise accurate patient diagnosis. Although some form of restraint tends to be used in practice, motion incidence and magnitude is still high enough to warrant frequent repeat studies or the application of motion correction. The motivation for this work was the outstanding need for a high-performance motion correction strategy for brain studies. Such a strategy should accurately correct general rigid-body motion. The optimal strategy would also be non-invasive, have a high degree of automation, and be fast, convenient (requiring little or no calibration, patient cooperation, and extra gadgetry), and robust with respect to noise. We describe and implement a fully 3D, non-invasive, data-driven approach that is suitable for use with clinical data and is potentially automatic. The approach is based on a comparison of measured and estimated projection data. Acquired projections are segregated into groups corresponding to discrete locations held by the brain during scanning and the largest group is reconstructed. The position and orientation of this reconstruction is optimised for each remaining group by comparing the measured projections with those generated from the transformed reconstruction. After each optimisation, the current reconstruction estimate is updated with the relevant projections using the ordered-subsets expectation maximisation (OSEM) algorithm.

Methods: Three sets of experiments were carried out on different types of data to validate the motion correction procedure and investigate practical aspects of implementing the approach clinically. In the initial set of experiments, seven noisy motion-corrupted projection sets simulating 2-4 head positions were generated from the digital Hoffman brain phantom. The angular location and extent of movement and the magnitude of rotation and translation with respect to each axis was varied for each set. Motion correction was applied to these data using various regimes: with/without attenuation included in the optimisation; with/without a second iteration. Extracted motion parameters were compared with the applied movements. The error between the extracted and applied parameters was quantified in terms of the mean registration error (MRE), an average displacement of the vertices of a box surrounding the brain. Overall improvement from motion correction was quantified in terms of a mean squared difference improvement ratio (MSDR). Corrected, uncorrected, and motion-free slices were also compared visually. For the second group of experiments, three physical Hoffman phantom studies containing single or double movements were obtained. The Polaris motion tracker was used to provide an independent measurement of motion. Motion parameters were extracted using our approach and compared with those measured by the Polaris. An investigation of cost function behaviour was also carried out by mapping the cost function in the neighbourhood of the Polaris solution. The third group of experiments constituted a preliminary clinical validation. Three volunteers underwent a motion-free scan followed by a scan in which they performed one head movement. A fourth volunteer underwent two scans, holding a single (but

different) brain location in each. Again the Polaris was used to measure the motion independent of our technique. Data from the fourth volunteer was used to simulate two single-movement studies, facilitating a rigorous quantification of the improvement obtained from motion correction. Optimisations were performed with and without reduced projections, scatter correction, thresholding of background counts, compensation to avoid biasing from truncated data, and pre-smoothing of the acquired data.

Results: In the digital phantom experiments, estimated rotations and translations were mostly within 2° and 1mm of the applied values. The MRE was less than 1 pixel in most cases. Accurate motion estimates could be obtained at over twice the speed by leaving attenuation out of the optimisation stage. Visually, there was a clear reduction in motion-induced artifacts after correction. Most MSDR values were well in excess of 2, and the MSDR tended to increase with increasing corruption. A second iteration of correction did not provide sufficient improvement to warrant the additional time cost. In the physical phantom experiments there was good agreement between the extracted and Polaris measurements for the x and y-rotation and z-translation parameters. A systematic discrepancy existed for the remaining parameters. The discrepancy was reduced for the third dataset (two movements); in this case the corrected study closely resembled that obtained using the Polaris values. Analysis of the cost function indicated that the MSD was fairly insensitive to large rotations whilst being sensitive to typical translations. Discrepancies appeared to be the result of object symmetry. In all of the volunteer studies, sets of motion parameters were obtained that closely followed the trend of the Polaris. In general, however, there was a systematic discrepancy from the actual Polaris values. Scatter correction had little effect on accuracy. Using reduced projections (greater proportion of the image occupied by brain) tended to provide estimates as good or better than using larger projections. Pre-smoothing generally lead to less accurate estimates. For large movements, tracking the plane of truncation was necessary to obtain sensible estimates. Thresholding was important in removing background counts and confining the solution to a sensible portion of the cost function. For all volunteers there was a clear improvement in image symmetry and contrast after using our approach. In certain cases, correction was better than that obtained from the Polaris. Of particular concern is the method used for attaching the head target. Poor attachment can lead to decoupling of target and head movement. For the two semi-simulated studies, the MSD improved by approximately 4 and 2 respectively, whereas the Polaris provided no improvement.

Conclusion: We have demonstrated that complex brain movements in simulated and real data can be accurately estimated and corrected using this data-driven approach.

CHAPTER 1

INTRODUCTION

1.1 Single Photon Emission Computed Tomography

Single photon emission computed tomography (SPECT) is a nuclear medicine imaging modality used for assessing physiological function. A radiopharmaceutical is injected into the patient and a series of two-dimensional (2D) images of the radionuclide distribution are collected using a rotating gamma camera. These data are then reconstructed to give a stack of transverse slices representing the estimated three-dimensional (3D) distribution of radioactivity in the organ or region of interest.

SPECT resolution and sensitivity are far worse than for the structural modalities like computed tomography (CT) and magnetic resonance imaging (MRI). The external collimator, required to produce useful images, is the primary cause of poor detection efficiency, leading to the low-count (high noise) images characteristic of SPECT. Efficiency can be improved by increasing the administered dose and/or the scan time, though patient comfort and safety impose a limit on both of these. Photon attenuation in the patient tissues, and scatter within the tissues and detector crystal, are additional factors affecting image quality and quantification. Compensating for these and other physical and geometric factors is non-trivial. Compared with planar gamma camera imaging, SPECT provides greater image contrast and better delineation of overlapping structures, as well as reproducible and quantifiable activity distributions. Despite its limitations, SPECT remains a widely used and vital tool for assessing and diagnosing the functional state of organs and tissues, particularly in brain and heart. A thorough coverage of SPECT physics and quantification can be found in Fahey et al (1996) and Rosenthal et al (1995).

1.2 Patient Motion

Patient movement during an acquisition is a well-known cause of distortion and artifacts in reconstructed SPECT data. This has been shown to compromise accurate patient diagnosis. The character, extent, and incidence of artifacts have a complex dependence on many factors, including the size, direction, timing and duration of motion. Motion incidence is expected to be high in brain studies due to the large number of patients presenting with epilepsy and dementia. Moreover, when patients are physically limited and/or non-compliant, the task of acquiring useful, motion-free data is made more difficult.

Head restraint is used routinely in brain SPECT acquisitions to minimise voluntary and involuntary patient motion. Despite this, motion can be significant enough to thwart attempts to reconstruct diagnostically accurate images. Unless a repeat study free of motion is performed, robust motion correction is required. A brief review of relevant literature on patient motion and various detection and correction strategies is provided in the next chapter.

An optimal motion correction strategy for brain SPECT would certainly compensate for six degree-of-freedom (DOF) rigid-body motion. The strategy should also be

- ★ Fully automated
- ★ Fast
- ★ Non-invasive
- ★ Robust with respect to noise
- ★ Convenient, i.e. requiring minimal additional equipment, calibration, and patient involvement

The motivation for this work was the outstanding need for a strategy that satisfies most of these desirable features. We describe and implement a fully 3D, non-invasive, data-driven approach that is suitable for use with clinical data and which is potentially automatic.

1.3 Principle of Data-Driven Motion Correction

Fundamentally, motion correction is the task of obtaining consistent projection data. When motion occurs between discrete acquired projections, consistency is lost and errors will be propagated to the reconstructed estimate. Within certain limits, the reprojections generated from this estimate are distinguishable from their measured counterparts and can yield information about the corruption. This in turn may be used to reverse the effects of motion. A number of previous authors have demonstrated that motion correction is feasible in SPECT using a comparison of the acquired projections with projections generated from the reconstruction[1] (eg. Lee and Barber 1998, Arata et al 1995). Armed with knowledge of the detector orientation relative to the object at each projection angle, provided projections can be incorporated into the reconstruction according to this orientation, then a means for reversing motion is obtained. This principle is illustrated schematically in Figure 1.1.

[1] Relating measured and estimated projection data has a much broader scope and application than SPECT, eg. Penczek et al 1994.

1.4 Introduction to the Motion Correction Approach

The principles of the approach have been described in detail in Hutton et al (2002). Moreover the feasibility of detecting and correcting general 3D patient brain movement using this method has been established (Hutton et al 2002, Kyme et al 2001). A brief summary is provided here.

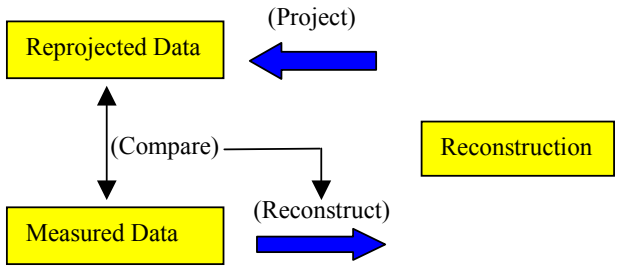

Figure 1.1. Principle of motion correction using a comparison of the measured and estimated (reprojected) projections. Motion information from this comparison enables the measured projections to be incorporated into the reconstruction consistent with the orientation they were acquired at.

A dual-90° detector gamma camera is used to acquire a full set of projection data consisting of orthogonal pairs of projections. These data are then reconstructed and the resulting estimate is reprojected. Comparison of the measured projection pairs with the corresponding reprojected pairs using a mean square difference (MSD) similarity measure enables groups of projections acquired for the same brain orientation to be identified. For each group, the projection pair exhibiting the greatest dissimilarity (largest MSD) is selected as the pair to determine the orientation for the group. The downhill-Simplex optimisation routine is used to try various rigid-body (3 rotations, 3 translations) transformations of the reconstruction. After each 3D transformation tested by the Simplex routine, the reconstruction is reprojected and the relevant projection pair compared with its measured counterpart. Convergence to a solution is determined by whether the similarity measure is within a user-defined tolerance or if a specified number of iterations are reached. Once the orientation has been estimated, the projection pairs associated with this orientation must be incorporated into the reconstruction. The ordered-subsets expectation maximisation (OSEM) algorithm is used for this. OSEM allows starting reconstructed estimates to be updated with additional projection information. Using the reconstruction resulting from the Simplex optimisation stage as the starting image, all projection pairs belonging to the same group as the optimised pair are made to update the reconstruction. The procedure repeats until all projection pairs from all groups have been incorporated consistently into the reconstruction. A flow chart of the approach is shown in Figure 1.2.

The main results from previous work are summarised below:

- Comparison of the measured and estimated (reprojected) projections enabled clear identification of flat peaks and troughs corresponding to different locations of the brain during acquisition.
- Increasing noise in the projection data caused an increasing MSD offset in the curves without compromising the ability to distinguish peaks and troughs. This was true at a level of noise typical of a clinical study and also at much higher levels.
- Optimising the orientation of the reconstruction enabled six DOF simulated motion to be estimated accurately. Accuracy varied for different DOF at any particular noise level, and worsened for all parameters as noise increased.
- The feasibility of correcting discrete 3D motion using this technique was demonstrated.

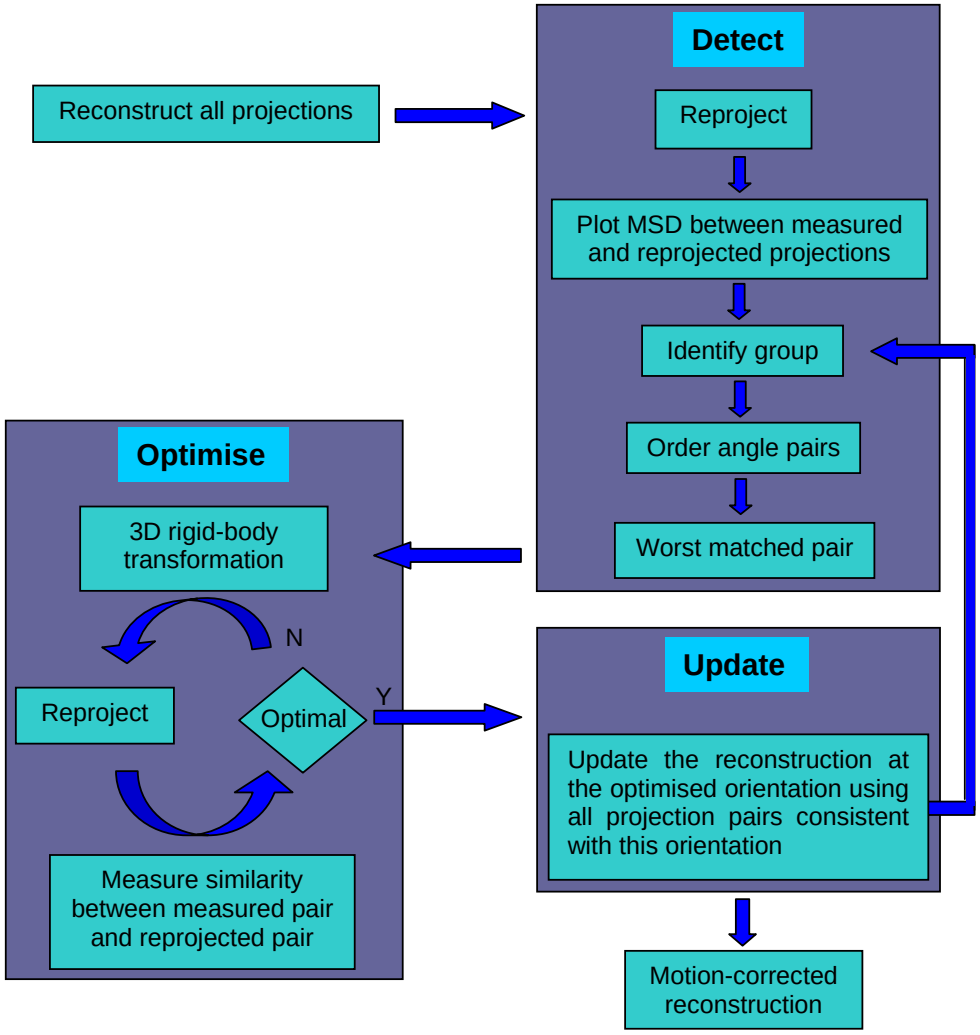

Figure 1.2. Flow diagram of the data-driven motion-correction approach. Additional details can be found in the text and in Hutton et al (2002) and Kyme et al (2001).

1.5 Objectives of the Present Work

The present work builds on previous methodology and results. We seek to validate the approach for simulated and real data as well as provide a detailed framework for practical implementation of the approach. Our objectives can be summarised as:

(i). Validate the approach for correction of complex brain motion in simulated SPECT data

(ii). Demonstrate feasibility of the approach for correcting complex brain motion in real data (phantom and human subjects)

(iii). Identify practical considerations when implementing the approach with real data

(iv). Suggest a protocol for data-driven motion correction in clinical brain SPECT

1.6 Thesis Structure

The present work has been structured as follows:

- Chapter 2 provides background material on image reconstruction, registration and processing relevant to the motion correction technique. The second half of the chapter provides a review of the literature on motion in medical imaging and methods for tackling this problem.
- Chapter 3 provides a formal introduction to the algorithm used in the simulation and real data validation stages of the work.
- Chapter 4 details the methods and results of the digital phantom validation experiments.
- Chapter 5 details the methods and results of the physical phantom validation experiments. A description of the gold standard used for the real data experiments (Chapters 5 and 6) is also provided.
- Chapter 6 details the methods and results of the clinical validation experiments. Important practical considerations in implementing the approach with real data are also described.
- Chapter 7 provides a general discussion and summary, describes avenues for future investigation of this problem, and describes ways in which the technique could be further tested and enhanced. We also provide general conclusions of the work.

1.7 Publications and Published Abstracts

The following is a list of publications and published abstracts arising from this work.

(1) Kyme AZ, Hutton BF, Hatton RL, Skerrett DW, and Barnden LR. Practical aspects of a data-driven motion correction approach for brain SPECT. *IEEE Trans. Med. Imag*., 22:722-729, 2003.

(2) Hutton BF, Kyme AZ, Lau YH, Skerrett DW, and Fulton RR. A hybrid 3D reconstruction/registration algorithm for correction of head motion in emission tomography. *Trans. Nucl. Sci.*, 49:188-194, 2002.

(3) Kyme AZ, Hutton BF, Hatton RL, Skerrett DW, and Barnden LR. In Yves Bizais, editor, *Proceedings of the VIIth International Conference on Fully 3D Reconstruction in Radiology and Nuclear Medicine*, Saint-Malo, France, July 2003.

(4) Kyme AZ, Hutton BF, Hatton RL, Skerrett DW, and Barnden LR. Considerations in applying data-driven motion correction in real SPECT data. *J. Nucl. Med.*, 44 (5 suppl.), page 64P, 2003. (Abstract.)

(5) Kyme AZ, Hutton BF, Hatton RL, and Barnden LR. Optimizing data-driven motion correction in brain SPECT: partial reconstruction and attenuation correction. *J. Nucl. Med.*, 43 (5 suppl.), page 222P, 2002. (Abstract.)

(6) Kyme AZ, Hutton BF, Hatton RL, and Skerrett DW, Barnden L. Practical aspects of a data-driven motion correction approach for brain SPECT. *ANZ Nucl. Med*., December:143-147, 2002.

(7) Kyme AZ, Hutton BF, Hatton RL, and Skerrett DW. Evaluation of factors that influence the performance of an automatic motion correction algorithm for brain SPECT. *Nucl. Med. Commun*., 22:923, 2001. (Abstract.)

CHAPTER 2

BACKGROUND AND REVIEW OF LITERATURE

PART A: BACKGROUND

2A.1 Image Reconstruction

A brief description of the reconstruction problem and methods is provided here. Lalush (2003), Defrise (2002), and Bruyant (2002a) provide a comprehensive coverage of the theory and practice of analytical and iterative reconstruction techniques.

SPECT image reconstruction is the task of solving for an emitting distribution given projections of the distribution. The problem has traditionally been performed using 2D back-projection where data from all projection bins corresponding to a particular transverse slice are projected back through the imaging volume. Appropriate filtering (convolution) of the projection data removes back-projection artifacts. Provided there is sufficient angular sampling, no noise, and no physical effects such as attenuation, filtered back-projection provides a theoretically exact reconstruction. Analytic 2D reconstructions are also performed in the Fourier domain, taking advantage of the fact that the Fourier transform of a line in the projection space corresponds to an oblique line through the Fourier transform of the object (Fourier projection theorem). The object distribution is obtained from the transformed projections by applying the inverse Fourier transform. The method is prone to artifacts induced in the Fourier domain interpolation. In general, all analytic techniques are limited because only simplistic modelling of the photon physics is possible.

Many good iterative methods have now been developed to solve for tomographic reconstructions. Since these methods are numerical, the physics of photon emission and detection may be accounted for at any level of complexity. All iterative techniques are based on a feedback model aimed at achieving consistency between estimated and measured data. At the outset, the object (distribution) is initialised to some arbitrary value. This is projected and the estimated projections compared with measured data to derive a set of error projections. Error projections are in turn transformed to the image space and used to update the estimate. The process continues until the errors are within some specified tolerance. A general schematic for iterative reconstruction is shown in Figure 2.1.

Mathematically, projection of the distribution is modelled by a matrix H, the elements h_{ji} of which describe the probability that voxel i in the object contributes to projection bin j. All of the physics modelling emission, passage through a medium, and detection of photons may be incorporated into this matrix.

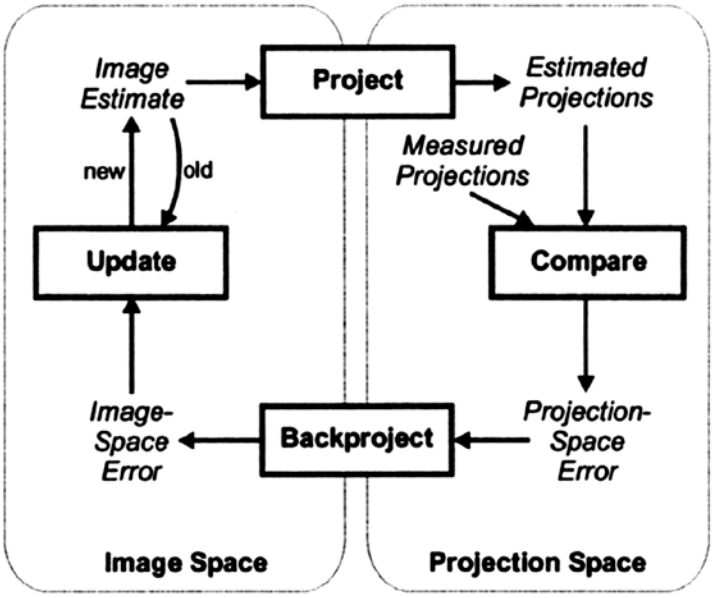

Figure 2.1. Generalised flow model of an iterative reconstruction algorithm. (Lalush DS. Iterative image reconstruction. In Wernick and Aarsvold, editors, *Emission Tomography: The fundamentals of PET and SPECT*. Elsevier (to appear).

Noise and other inadequacies in the data (such as motion) mean that clinical data are almost never 'consistent' – i.e. there is no object that, when projected, would give rise to the measured data. Statistical reconstruction algorithms account for the fact that there is variation expected in the acquired data. The problem then becomes one of finding the most likely solution from an array of possible solutions. Variation in the data, due, for example, to the random nature of nuclear decay, gets described by a statistical criterion. The iterative algorithm seeks the solution specified by this criterion.

Early on, the gold standard in iterative algorithms was maximum likelihood expectation maximisation (ML-EM) which combines the ML Poisson statistical criterion with the EM algorithm (Shepp and Vardi 1982). Iterative algorithms such as ML-EM provide good image quality and noise characteristics (eg. Shepp et al 1984, Chornoboy et al 1990, Soares et al 2000). Although the likelihood of a correct solution increases with increasing iteration number, noise also increases, thereby requiring that the process be terminated after some specified number of iterations. The solution subsequently may be smoothed to a visually and diagnostically acceptable level.

Ordered-subsets EM (OSEM) (Hudson and Larkin 1994) is a block-iterative variant of ML-EM that provides significant computational advantages. Here projection data are grouped into subsets of a desired size. Each update of the object is based on information from a single subset rather than the complete set of projections. This provides an update of similar quality to one, complete iteration of ML-EM, but

proportionately faster by a factor approximately equal to the number of subsets. An underlying assumption of OSEM is that each subset is equally likely to contain counts[1]. This is known formally as the subset balance condition. As a result, subsets are normally chosen to have maximum angular variation within and between them. However, as Lalush (2003) points out, reconstructions do not appear overly sensitive to ordering.

OSEM has no proof of convergence, even in the noise-free case as ML-EM does. It can be shown that OSEM converges to multiple solutions, and successive iterations oscillate between these in a process known as limit cycling. Despite the theoretical limitations, from a practical point of view OSEM has proved to be a popular and stable algorithm to work with in SPECT, possibly because the limit cycle solutions are very similar in general. It has been shown that OSEM is a special case of a more general block-iterative formulation that has a stronger theoretical basis (Soares et al 2000). In this formulation, convergence to a unique solution is proven in the noise-free case. One implementation, the rescaled block-iterative (RBI) algorithm, is not constrained by subset balance and therefore does not have any specific requirements on the projections within each subset (Byrne 1996, Byrne 1997, Byrne 1998). The row-action maximum likelihood algorithm (RAMLA) (Browne and Di Pierro 1996) is also very similar to OSEM. A relaxation parameter controls the impact of successive iterations and also permits single projection subsets to be used.

It has been shown that OSEM reconstruction facilitates 3D motion correction (Fulton et al 1994, Hutton et al 1997). If projections corresponding to similar locations of the object are grouped into subsets, an initial estimate may be reoriented according to motion knowledge such that each projection subset is used consistent with the orientation it was acquired at. This methodology is a core component of our motion correction approach and will be discussed in more detail in Chapter 3.

2A.2 Image Registration

Image registration refers to the spatial or temporal alignment of images. It is particularly relevant to medical imaging since patients routinely undergo imaging on multiple modalities and combining information from these scans is diagnostically advantageous (eg. Pelizzari et al 1989). Registration of images from one or more modalities can improve identification of disease, quantification of anatomical and physiological parameters, treatment planning and follow-up (eg. Picard and Thompson 1997, Radau et al 2000, Meltzer et al 1990). Much literature exists on the theory and practice of intra and inter-modality registration techniques. Van den Elsen et al (1993), Maintz and Viergever (1998a), and Hill et al (2001) provide a thorough coverage of the subject.

[1] Mathematically this is equivalent to requiring $\sum h_{ji}$ be constant over all subsets.

Registration is the problem of finding the mapping from one image coordinate space to another (Barber 1992). The computational demand of this task is determined by the mathematical complexity of the mapping function. For some applications such as intra-subject brain registration, a rigid-body model may be sufficient (eg. van den Elsen 2002). For many other applications (eg. blood vessel visualisation and inter-subject registrations), an affine model is more correct. There are numerous examples in the literature of non-rigid spatial warping methods to twist a group of control points on two separate images into coincidence (eg. Pickens and Price 1987, Barber 1992, Dann et al 1989).

Given sufficient reference landmarks the mapping can be solved algebraically or a solution can be searched for iteratively. Generally, registration procedures make use of a searching algorithm (optimisation strategy) to minimise or maximise a cost function describing the similarity between images. Press et al (1992) provides a summary of the principles of some common optimisation strategies and Capek et al (2001), Meijering et al (1999), Hill et al (2001), and Penney et al (1998) have described the choice of and performance of common similarity measures.

Comparison of grey values is a simple and intuitive method for intra-modality registration (eg. Eberl et al 1996). Measures used to quantify spatial agreement include minimising the sum of squared differences (eg. Slomka et al 1995), maximising the cross-correlation (eg. Eisner et al 1987), and minimising the standard deviation of the ratio of the images (eg. Woods et al 1992). Direct grey value comparison techniques need to be applied differently for inter-modality registration problems since the images have a non-linear grey value relationship. Intensity re-mapping (eg. van den Elsen et al 1994) can be applied first so that the intensity distributions resemble each other more closely.

Inter-modality registration can also be achieved using surface (eg. Turkington et al 1995) or marker (eg. Malison et al 1993) matching. However, such methods are limited by the consistency of the data (Levin et al 1984). An alternative is to use a global similarity measure, of which mutual information (MI) (Wells et al 1996, Maes et al 1999) is a popular choice. MI operates under the principle that improved registration leads to lower variance in the joint pixel-intensity histogram of the two images. In other words, as registration improves, there will be an increased probability of correctly predicting the grey value of a pixel of one image given the value in its assumed counterpart in the second image. A multi-variate MI registration approach where additional information is included from a third image can improve robustness in locating global extrema (Boes and Meyer 1999).

Segmentation is the manual or automated extraction of specific structures in an image. Numerous intra and inter-modality approaches rely on segmentation of one or more of the images to achieve registration (eg. Collins et al 1995 , Ardekani et al 1995).

Many motion correction approaches in the literature are based on registration of volume images from a time sequence (eg. Maas et al 1997, Picard and Thompson 1997, Pappata et al 2002) or alignment of planar images in the projection space prior to reconstruction (eg. Groch et al 1985, Levin et al 1984, Geckle et al 1988, Leslie et al 1997, Pellot-Barakat et al 1998). Essentially the motion correction problem is equivalent to the registration problem since motion leads to temporal or spatial misalignment between images.

A registration technique resembling our own approach to motion correction in SPECT has been described by (Surova-Trojanova et al 2000). Here planar emission images were registered to reprojected CT (rCT) images based on a comparison in projection space. The orientation of the CT reconstruction was optimised so as to maximise similarity between the rCT image and planar transmission image. Following this step, the rCT image and planar emission image were in registration by virtue of the fact that the emission and transmission data were acquired simultaneously. An extension of this method to register SPECT and CT volumes using two planar views has been described (de Vries and Moore 2001).

2A.3 Similarity Measures (Cost Functions)

A similarity measure, or cost function, is a criterion of registration quality. It is used to quantify the degree of similarity between two or more sets of data. Specifying an appropriate cost function for a particular registration problem is critical since this determines the multi-dimensional shape of the solution space, as well as the speed and reliability with which the solution may be found by an optimisation procedure. The optimal choice of cost function has been shown to depend on the modalities involved (eg. Penney et al 1998). Processing techniques such as smoothing alters the cost function shape and can make the measure more favourable (eg. Roche et al 1999). Some of the more common similarity measures used in medical image registrations are described here.

2A.3.1 Cross-Correlation

The cross-correlation function value is a summation of the product of corresponding pixels in each image, over an arbitrary number of dimensions. Variable offset parameters for each dimension are optimised to maximise the function and thereby obtain the shift in each dimension necessary to register the two images. The cross-correlation function has a high threshold due to noise and is unsuitable for detecting rotations and large-amplitude motion spread out in time (Eisner et al 1987). A description of the measure can be found in Buzug et al (1998) and Meijering et al (1999), and examples of its use in Eisner et al (1987), O'Connor et al (1998) and Penney et al (1998).

2A.3.2 Sum of Squared Differences (SSD) and Mean of Squared Differences (MSD)

As the name suggests, the SSD function is calculated by summing all values of the squared difference image. Dividing this value by the number of elements in the image gives the MSD function (also known as the cross-structure function). Like the cross-correlation function, these measures are based on actual grey values. A description of the measure can be found in Buzug et al (1998) and examples of its use in Bidaut and Vallee (2001) and Hoh et al (1993).

2A.3.3 Sum of Absolute Differences

This is calculated by summing all values of the absolute difference image and is a minimum at registration. (Eberl et al 1996) found it to be the most reliable and accurate similarity measure in registering SPECT and PET images when compared with the stochastic sign change (see deterministic sign change below), sum of products, and standard deviation of ratios measures. Other examples of using this measure can be found in Slomka et al (1995) and Hoh et al (1993).

2A.3.4 Deterministic Sign Change

A periodic pattern is imposed on one image before it is subtracted from the second image. Registration assumes the periodic pattern is maximally preserved in the difference image. Progress to a global solution can be hindered by local extrema. Gradient-based optimisation procedures cannot be used with this measure since it is discrete. A description of the deterministic sign change technique can be found in Buzug et al (1998) and examples of its use in Penney et al (1998) and Slomka et al (1995).

2A.3.5 Histogram-based methods

These methods operate on the principle that the difference of two registered images will have low-contrast. Common to them is an initial subtraction of images followed by formation of the normalised grey value histogram. A number of functions, such as energy and entropy of the histogram of differences, have well-defined extrema in certain applications. Probably the most common histogram measure is the mutual information (MI) (Wells et al 1996, Maes et al 1999). MI operates under the principle that improved registration leads to lower variance in the joint pixel-intensity histogram of the two images. It is used for global image registration and is particularly useful when a non-linear pixel relationship exists between images. MI produces well-defined extrema with excellent smoothness and steepness in CT-MRI registrations (Capek et al 2001). A multi-variate MI registration approach which includes information from a third image has demonstrated improved robustness in locating global extrema (Boes and Meyer 1999). A further description of histogram-based measures can be found in Buzug et al (1998), Meijering et al (1999) and Maintz et al (1998b), and examples of their use in Penney et al (1998) and Roche et al (1999).

2A.4 Optimisation Strategies

An optimisation strategy is a search procedure to locate the optimal state (extremum) specified by the criterion of registration quality (cost function). The search usually takes place in a multi-dimensional space. Optimisation procedures vary in robustness with regard to local extrema and convergence to a global solution. Some of the commonly used optimisation algorithms are described here. More detailed descriptions of these algorithms may be found in Press et al (1992), Ebert (1997a), Ebert (1997b) and Maes et al (1999).

2A.4.1 Downhill-Simplex Routine

The downhill-Simplex algorithm (Nelder and Mead 1965, Press et al 1992) is a non-exhaustive search procedure based on contortions of a geometric structure (simplex) that traverses the multi-dimensional space. Successive iterations of the routine are characterised by expansions, contractions, reflections, and scaling of the structure. Maes et al (1999) showed that the Simplex routine reliably located the global minimum and out-performed Powell's method and the steepest gradient descent method in CT-MR image registrations. Other examples of its use can be found in Slomka et al (1995) and Hoh et al (1993).

2A.4.2 Powell's Method

The cost function is minimised for each dimension independently. The solution in one dimension becomes the origin for the following search in the next dimension. It has been found to converge more rapidly than the downhill-Simplex and gradient methods for high resolution studies, but slower in low-resolution studies. An example of its use can be found in Faber et al (1991).

2A.4.3 Steepest Gradient Descent

This is similar to Powell's method but, at each minimum, optimisation occurs in the direction of steepest descent. In general it is not the optimal algorithm choice because of slow convergence.

2A.4.4 Stochastic Methods

Stochastic methods such as simulated annealing are another branch of optimisation algorithms. They are less 'systematic' and rely on random perturbations to explore the solution space. The changes are accepted if the cost function proceeds closer to the optimal value. If local minima are expected to be problematic, a stochastic approach is preferable. A drawback of stochastic methods is slow convergence.

2A.5 Image Processing

2A.5.1 Volume Projection

Many techniques for volume projection are described in the literature. These are broadly classified into pixel-driven and ray-driven methods depending on whether the process is considered as a filling of projection bins or an emission from volume bins respectively. Approaches differ in accuracy and speed according to the complexity of sampling and interpolation used. Computationally efficient projection of image data for tomographic image reconstruction and simulation can be achieved by appropriate rotation of the data followed by summation along the primary matrix axes (eg. DiBella et al 1996, Lalush 2003). This approach has been used for reprojecting image volumes in the present work.

2A.5.2 Interpolation

Standard image processing often requires interpolation of sampled data. For example, non-orthogonal rotation requires interpolation of either the original image or the rotated image. If the interpolation scheme is inappropriate, image counts may not be preserved, uniformity can be degraded, positional errors can be introduced, and excessive or angle-dependent blurring can result (Wallis and Miller 1997). This in turn can degrade optimisation procedures, registrations (Hill et al 2001, Meijering et al 1999), and quantitative analyses (eg. Kuhle 1992).

The ideal kernel for interpolation is the (infinite) sinc function. For practical, efficient processing, this is often approximated using either piecewise polynomials or windowing. The simplest piecewise polynomials give rise to the nearest neighbour and linear interpolation schemes, the latter being by far the most widely used approach in general applications. Meijering et al (1999) showed that linear interpolation is the best option when two grid points are used to define the kernel. A cubic kernel is the best approach when four grid points are used. When additional grid points are used, windowing the sinc provides more accurate interpolation than polynomial approximations, though the choice of window function is critical. A rectangular window (truncated sinc) for example gives poor results. The choice of kernel is complicated by a dependence on the imaging modalities involved and the 3D plane selected. Wallis and Miller (1997) described a rotator that redistributes source image counts according to Gaussian blobs. Performance with respect to count preservation, uniformity, positional errors, computational speed, and angle-dependent blurring was superior to other algorithms tested (including bilinear interpolation). In the present work we used the bilinear scheme for 3D transformations because it provides the best speed and accuracy for a simple kernel and cubic-voxel basis. Bilinear interpolation is also applied in our OSEM reconstruction software.

PART B: REVIEW OF LITERATURE

2B.1 Motion in Medical Imaging Modalities

Many authors have sought to describe and characterise the sources and impact of patient motion in various medical imaging modalities. A brief coverage of the literature is provided here with particular attention paid to the nuclear medicine modalities SPECT and PET.

2B.1.1 Types of Motion

Patient motion can be classified broadly into involuntary and voluntary motion, and subdivided further into rigid and non-rigid (elastic) motion. Heartbeat motion, respiratory motion (and associated movement of organs with the diaphragm) (eg. Klein et al 2002, Bruyant et al 2002b, Nehmeh et al 2002, Blodgett 2002), cardiac "creep" (translation of the heart up the chest after exercise due to changing diaphragm position) (Friedman et al 1989, Eisner et al 1988, Tsui et al 2000), and peristalsis (eg. abdominal) are all examples of involuntary, non-rigid motions that do occur during medical imaging. The periodicity and temporal characteristics of these motions vary considerably.

Other examples of involuntary motion that may or may not be expressed as rigid motion are vibrational motion such as cardiac or respiratory-induced head motion (Freire et al 2002, Maas et al 1997, Calhoun et al 1998), sneezing, coughing, pain reflexes, tremor and other movement disorders, and slow rotational and translational drifting motions (Fulton 2000, Pandos et al 2002, Bloomfield et al 2003).

Voluntary movement can result from general restlessness, adjustment to relieve pain or discomfort, leg-crossing, response to stimulus in the room, and compliance with activation paradigms where the physiological effect of performing certain tasks is being tested (Lopresti et al 1999, Linney and Gregson 2001, Fulton 2000, Pandos et al 2002, Freire et al 2002, Maas et al 1997, Calhoun et al 1998, Zeffiro 1996). Placing arms behind the head in cardiac SPECT is a common source of rapid, discrete heart motion (Geckle et al 1988). Germano et al (1994) and Pandos et al (2002) have observed a 'settling' period early in the scan and subsequent 'discomfort' period late in the scan, with less motion mid-scan.

Cumulative head motion increases with increasing scan time irrespective of restraint (Green et al 1994, Fulton 2000). In any given time period there are multiple statistically significant mean positions. The cumulative increase therefore results from differences in the frequency and amplitude of non-random movements rather than an increase in the frequency and amplitude of random movement about a single mean (Lopresti et al 1999). Random motion is spatially symmetric about an average for each DOF and is generally low amplitude and high frequency. Non-random movement tends to be high amplitude and low frequency, leads to a change in mean position, and results from speaking, leg-crossing, sneezing, coughing and pain. Slow, continuous movement can last minutes and often follows large, non-random movements (Lopresti et al 1999). These results are in agreement with Bloomfield et al (2003) and Fulton

(2000) who detected sharp, rapid movements (up to 4° rotations and 15mm translations) interspersed by periods of stability or slow drift (1.5° and 2-3mm) when monitoring normal volunteers and epilepsy patients. Pandos et al (2002) found a similar pattern amongst patients (dementia, psychiatric, fibromyalgia, and chronic fatigue syndrome) and normal volunteers. Thus commonplace head motion is characterised by slow shifts in position in all DOF, as well as random, discrete, rapid movements of relatively large amplitude.

2B.1.2 Effects of Motion

It is well known that patient motion can cause artifacts in medical imaging modalities, particularly tomographic modalities (eg. Yang et al 1982, Eisner et al 1987, Silver et al 1994). Cooper and McCandless (1995) defined motion-induced artifacts as any streaking, smudging, defects or hot spots unexplained by physiology or pathology that are improved by motion correction.

Many authors have sought to characterise artifacts produced by various types of patient movement (eg. Silver et al 1994). One group defined the primary groupings for SPECT artifacts as streaking at high contrast boundaries occurring for all types and combinations of motion, boundary discontinuities induced by translational motion, and distortions of shape and size in weak contrast objects (Yang et al 1982). In general, artifacts can be manifest in many ways, some examples being blurring and loss of resolution, hot and cold spots masking true physiology, general distortion, reduced contrast, mis-positioning/splitting of activity, and deformation (Depuey and Garcia 1989, Klein et al 2002, Germano et al 1993, Ivanovic et al 2001).

The presence of motion-induced artifacts in any modality can affect quantitation. Examples are false positive results for SPECT radiotracer uptake in brain structures (Pappata et al 2002), erroneous lung lesion characterisation using PET (Nehmeh et al 2002), and altered lung density values from CT (Blodgett 2002). In digital subtraction angiography (DSA)[1], tiny patient motions can lead to false blood vessel visualisation and distortion of true vessels (Meijering et al 1999). Small motion is also problematic in functional magnetic resonance imaging (fMRI)[2], leading to spurious signal intensity changes unrelated to the specific task performed (Zeffiro 1996, Freire et al 2002, Maas et al 1997, Calhoun et al 1998).

Random motion appears to induce smaller artifacts than non-random motion (O'Connor et al 1998), and non-returning motion more serious artifacts than periodic motion (O'Connor et al 1998, Matsumoto et al 2001). Movements affecting a single projection have little effect on the final reconstruction, though for any given amplitude, artifacts become more severe as the number of projections affected increases (Silver

[1] DSA involves subtraction of a baseline X-ray image from a contrast-enhanced X-ray image in order to visualise blood vessels. A good coverage of motion in this modality can be found in Meijering et al (1999).

[2] fMRI is used for imaging task-induced changes in cerebral blood flow and typically involves a time series acquisition. Motion of the head caused by cardiac and respiratory vibration, or performance of the task itself, regularly leads to inaccurate activation signals due to mis-registration of the time series images.

et al 1994). Gross patient motion is considered by some authors (eg. Germano et al 1994) to be the cause of most, and typically the worst, motion-induced artifacts. It has been demonstrated that artifacts induced by complex motion involving more than one DOF exhibit the same features as those induced by the individual components of the motion, suggesting the impact is approximately additive (Yang et al 1982, Botvinick et al 1993). However, there are no simple rules governing artifact production. In general, the location, extent, and abundance of motion-induced artifacts has a complex dependency on the time of occurrence in the acquisition, direction, amplitude and duration of motion (Botvinick et al 1993, Prigent et al 1993, Friedman et al 1989). Moreover, alignment of transmission data for attenuation correction (Bailey 1998), the object size and location in the field of view (FOV) (Li et al 1995b), the detector resolution (Green et al 1994), the direction of camera rotation, scan time, and clinical pathology (eg. Pandos et al 2002) are all influential factors. Modalities become increasingly susceptible to motion-induced artifacts the better the scanner resolution. For example, high resolution CT (sub-mm resolution) and state-of-the-art PET (2-3 mm resolution) make small patient movements problematic. In nuclear medicine studies, occurrence of motion artifacts is exacerbated by long scan times, typically 20 mins to 2 hours (Ivanovic et al 2001, Silver et al 1994). As a result, motion is more likely in single-detector scans where the scan time must be increased compared to multi-detector acquisitions so that sufficient counts are obtained (Matsumoto et al 2001). Moreover, various authors (eg. O'Connor et al 1998, Matsumoto et al 2001, Hillier and Wallis 1999) have made the observation that multi-detector acquisitions are more sensitive to motion (and hence motion-induced artifacts) since multiple views are acquired simultaneously. By contrast, multi-rotation schemes employing alternating rotation of detectors has been shown to be less sensitive to motion (eg. Britten et al 1998, Germano et al 1994). Complicating the matter further is the fact that artifacts show dependency on the properties of restraining devices (Green et al 1994, Menke et al 1996).

The threshold of artifact production is an important consideration for day-to-day scanning. However, the complex dependency of artifact production on the nature of motion and other variables makes dogmatic specification of such a threshold difficult. Practically, it is helpful to have rough guidelines, and numerous authors have suggested one-pixel motion (Prigent et al 1993, Botvinick et al 1993, Matsumoto et al 2001), or more conservatively, half-pixel motion (Eisner et al 1987, Eisner et al 1988, Friedman et al 1989). Motion that can be identified visually was the threshold suggested by Cooper et al (1993), though Yang et al (1982) reported significant artifacts based on a more conservative limit. The threshold of artifact production is also a vital consideration from the perspective of design and performance of correction algorithms. A range of algorithms with low sensitivity (~ 4-5 pixels) (eg. Passalaqua and Narayanaswamy 1995) to high sensitivity (~ 0.5 pixels) (eg. Germano et al 1993) have been published.

Authors report motion occurrence in cardiac studies for between 10-40% of cases (eg. Eisner et al 1988, Cooper and McCandless 1995, Botvinick et al 1993, Prigent et al 1993, Germano et al 1994). The number of mis-diagnoses reported from scans identified as having motion is also quite variable (4-40%) (eg.

Prigent et al 1993, Botvinick et al 1993, Cooper and McCandless 1995). O'Connor et al (1998) estimated 10-20% of cardiac SPECT studies either required application of motion correction or a repeat scan because of significant motion-induced artifacts. Such figures can be misleading in isolation and should be interpreted with respect to the type and size of motion. Head rotations exceeding 1° are readily observed in Parkinsonian patients according to Menke et al (1996). Pandos et al (2002) monitored the head motion of 15 dementia, 8 psychiatric, 10 fibromyalgia and 18 chronic fatigue syndrome patients, as well as 23 normal volunteers. They found 71% of subjects moved > 1mm or 1°, and 10% moved > 4mm or 4°. Moreover, 65% of the subjects moved at least three times during the 30-minute scan.

In certain cases it has been demonstrated that applying motion correction can produce or enhance motion artifacts (eg. Cunneen et al 2001), however this appears to be due to limitations in the algorithms (Matsumoto et al 2001).

2B.2 Prevention of Motion in Medical Imaging

It is said that prevention is the best form of medicine and many authors hold to a preventative philosophy with regard to patient motion. In fact, authors across the modalities generally agree that irrespective of whether motion correction is applied to data or not, some form of motion prevention during the acquisition should be mandatory practice.

Immobilisation techniques such as stereotactic frames (eg. Bergstrom et al 1981, Enzmann and Freimarck 1984) have been used to secure the head. More recent approaches include mouldable polyurethane head holders or thermoplastic masks (Mawlawi et al 1999, Green et al 1994, Ruttiman et al 1995). Simpler measures such as Velcro® straps appear to be commonplace in many nuclear medicine departments. Such measures reduce, but do not eliminate, motion (Fulton 2000). For example Mawlawi et al (1999) quoted average head movements of approximately 2mm and 4mm using polyurethane holders and thermoplastic masks respectively, and Green et al (1994) observed movements of 3-10mm for subjects wearing the thermoplastic mask.

Motion prevention devices requiring patient cooperation such as bite plates and handgrips (eg. Feldkamp 1994) have been used to restrict motion, but these may not be appropriate for less compliant patients. Patient positioning and comfort is a simple but important motion prevention measure, particularly for long scans. To limit internal and gross patient motion in cardiac SPECT, Kiat et al (1992) employed prone positioning. Devices such as the Patient Support Device® (Cooper and McCandless 1995) provide additional support to parts of the body such as the back and arms to promote comfort.

Artifacts induced by respiratory motion have been minimised using appropriate breath-holding protocols (eg. Blodgett 2002). Gating, although not strictly a preventative technique, allows data acquired during different intervals of the respiratory cycle to be segregated and binned into separate images. Cycle information is measured via optical monitoring of the chest surface, spirometric techniques, chest impedance techniques, air flow and temperature techniques, pressure transducer techniques, and data-driven techniques (eg. Klein et al 1998, Cho et al 1999, Klein et al 2001). Cardiac gating involves the same principle applied to the heart cycle and makes use of the ECG signal as a gating trigger (Germano et al 2001).

Other preventative measures described in the literature include reduction of scan time (eg. Richie et al 1992, Germano et al 1994, Passalaqua and Narayanaswamy 1995), administration of drugs such as glucagon to minimise peristaltic motion in the abdomen (Harrington 1983, Rabe et al 1982), anaesthesia to minimise gross patient motion (particularly useful for paediatric and non-compliant patients), and use of devices such as abdominal compression bands to displace overlying organs and tissue from the area of interest (eg. Boxt 1983).

To summarise, preventative measures alone do not, and in certain cases cannot, eliminate motion. They may also be impractical and cumbersome for routine clinical use. Immobilisation techniques can be invasive, cause a considerable amount of discomfort, and even be counter-productive (Cooper and McCandless 1995, Bloomfield et al 2003). Other forms of restraint are often insufficient to prevent motion-induced artifacts, and in some cases lead to device-dependent features (Green et al 1994). Moreover, restraint does not limit motion to the level necessitated by state-of-the-art scanner resolution. Promotion of patient comfort is important and necessary but incidence of motion can still be high (Cooper and McCandless 1995).

2B.3 Solving the Motion Problem

There are two critical stages in any method for solving the motion problem in medical imaging. The first stage is motion detection, which may or may not include estimation of the motion. The second stage is motion correction where information from the detection stage is applied in some way to compensate for motion and give rise to a more consistent set of data for the task required. Detection and correction will be treated separately, though it should be noted that most authors describe detection and correction methodologies together when presenting a motion correction approach. The reason for treating the stages separately is because, in many cases, techniques from each stage can be ‘mixed and matched’.

2B.4 Motion Detection

As mentioned above, stage one of any motion correction approach is to detect motion. All motion detection techniques can be classified as either external or internal. External detection techniques are based on independent measurement of the patient motion, i.e. motion information is separate from the actual image data of interest. Internal, or data-based, techniques use the measured data itself to detect motion.

2B.4.1 External Motion Detection

Electromagnetic motion tracking systems have been used to measure head motion in PET (eg. Green et al 1994, Goldstein et al 1997). A reference field induces currents in orthogonal coils fixed to the patient. The size and direction of currents can be converted to positional information via calibration data. Eddy currents induced by nearby conductors (eg. detector gantry) degrade the stability of electromagnetic systems (Goldstein et al 1997).

Mechanical tracking arms have been implemented for motion detection in nuclear medicine (eg. Fulton 2000). The arm is secured to the object of interest and sensors at the device joints generate electrical signals based on the degree of rotation enabling object movements to be computed.

Optical techniques dominate the literature. For example Picard and Thompson (1997) used video cameras to detect the position of light emitting diodes on the patient. A similar system developed by Goldstein et al (1997) for emission tomography involved tracking the triangulation of three incandescent lights fixed to the head. This device had similar performance to an electromagnetic tracking system. Web cameras were used in SPECT studies by Gennert et al (2002b) to enable detection and computation of the patient surface as it changed with motion.

Infrared devices have the advantage of being insensitive to ambient light levels. Menke et al (1996) achieved high performance (50 Hz sampling, 0.2mm accuracy) 3D head motion tracking in PET using two CCD video cameras to collect reflections from infrared targets attached to a device secured to the patient's teeth. The system in use by many research groups presently is the infrared Polaris motion tracker developed by Northern Digital Inc[1]. Four infrared reflectors on a target secured to the patient head enable six DOF (3 position, 3 orientation) data to be sampled. The Polaris was applied in PET by Lopresti et al (1999) who sampled at 1 Hz and obtained 0.5 mm and 0.5° accuracy for translations and rotations respectively, suitably inside the recommendations suggested by Green et al (1994). An accuracy assessment by Watabe et al (2002b) indicated < 2mm and < 0.5° errors for translation and rotation respectively. They also demonstrated that the error propagated to the whole motion-corrected brain was <

[1] Northern Digital Inc., Waterloo, Ontario, Canada.

2%. Numerous other authors have recently applied the Polaris to brain SPECT (eg. Pandos et al 2002) and PET (Fulton 2000, Fulton 2002, Woo et al 2002, Watabe et al 2002a, Bloomfield et al 2003).

If quantitative information from an external device is to be applied in a motion correction scheme, the mapping from device coordinates to imaging (eg. gamma camera) coordinates must be computed. Gennert et al (2002a) developed a calibration phantom for this task. Other examples of obtaining the mapping for particular devices can be found in Fulton (2000) and Goldstein et al (1997).

2B.4.2 Internal (Data-Based) Motion Detection

As a general quality control measure during SPECT imaging, an operator typically views the cinematic sequence of projection images (cine). This also serves as a qualitative means of assessing the presence and degree of patient motion (eg. Depuey and Garcia 1989, Matsumoto et al 2001, Germano et al 1994), being additionally sensitive for multi-detector acquisitions (O'Connor et al 1998). The approach is limited by being operator-dependent.

A sinogram is a plot of source distribution along the detector FOV as a function of angular position of the detector. For a general distribution this appears as a complex pattern of superimposed sine curves. Discontinuities in the sinogram caused by patient motion can be detected visually (Eisner et al 1987) provided the motion is large enough. A less subjective approach using the sinogram was described by Lu and Mackie (2002) who incorporated high activity/contrast nodal markers in the object. The motion-corrupted sinogram trace of nodal regions appears as prominent but distorted sine waves in sinogram space. This trace can be used to estimate the expected trace, from which motion parameters are derivable. They demonstrated feasibility of detecting in-plane respiratory motion in CT and emission CT applications using this method. Linney and Gregson (2001) described a motion detection approach in brain CT using opposing views acquired at different times to estimate the time derivative of the sinogram. From this, a sinogram of motion was derived. However, only limited extraction of the motion component was demonstrated, even for simplistic in-plane movements.

Motion is also identifiable via the frequency spectrum of acquired data. Such an approach is not common in SPECT because of the long scan times. One example is Groch et al (1985) who identified regular, non-random peaks in the frequency domain resulting from cardiac, breathing, and exercising movements in multi-gated blood-pool studies. Frequency domain techniques are far more common in fMRI because of the much smaller temporal scale. For example the 'navigator' echo technique (eg. Anderson and Gore 1994, Ulug et al 1995) enables the measurement of phase-shift information related to patient movement. Another example is phase-only matched filtering which has been used to extract rotational offsets between two images (eg. Maas et al 1997, Calhoun et al 1998).

Cross-Correlation Techniques: Cross-correlation in the projection space is the basis of many approaches to motion detection (eg. Eisner et al 1987, Lee and Barber 1998, Arata et al 1995, Ivanovic et al 1998, Pellot-Barakat et al 1998, Passalaqua and Narayanaswamy 1995, Britten et al 1998, Leslie et al 1997). For example, Eisner used the cross-correlation function to detect in-plane translational motion in cardiac SPECT. Motion was seen as abrupt spikes in the frame-to-frame correlation. In the dual-scan approach of Passalaqua, cross-correlation of projections from an initial fast scan and a subsequent slow scan was used to determine inter-scan translational motion. Ivanovic and Pellot-Barakat obtained multiple projections at each angle using a full 360° orbit protocol with triple-detector gamma camera. Projections were then cross-correlated with two others at the same angle to produce correlation plateaus. The method was extended to hybrid PET/SPECT systems by cross-correlation of partial sinograms (Pellot-Barakat et al 2001) but this technique was limited to the detection of fast, discrete motion and had low sensitivity (translations > 1cm, rotations > 10°). A multi-rotation scheme over 180° followed by cross-correlation of corresponding projections was used by Britten to measure 1D axial translation of the heart. Cross-correlation of a cutaneous fiducial marker was shown to enable more accurate tracking of cardiac motion than cross-correlation of cardiac activity (Leslie et al 1997). In general, frame-to-frame effects such as object asymmetry, and gantry motion inconsistencies can hamper precise quantification using cross-correlation measures. A more sophisticated correlation approach was described by Dhanantwari et al (2001) for x-ray CT cardiac imaging. A spatial overlap correlator scheme was applied to data from multiple rotations of the scanner to extract the phase-amplitude stamp of organ motion as a function of time. The method is limited to the detection of 'cyclical' (not necessarily periodic) motion and is therefore not generally applicable in brain SPECT where patient motion has no predictable trace. Moreover, a high temporal sampling rate is required which is not feasible in SPECT due to low sensitivity.

Centre-of-Mass (COM) Techniques: A COM technique to detect head movement during cerebral blood flow studies of the brain was performed by Reichmann et al (2002). Here two complete (360°) acquisitions using a triple-detector gamma camera were collected fifteen minutes apart. COM matrices were derived for each scan and quantitatively compared to infer whether significant motion had occurred. The method was sensitive to motion amplitudes ~2mm. It has been well reported that respiratory motion causes organ movement during chest and torso emission imaging. Gating of data according to the respiratory cycle is a more sophisticated alternative to breath-holding regimes to correct for such motion. Recent papers have demonstrated that effective respiratory gating is possible using the data itself. Computation of the axial COM in dynamically acquired SPECT (eg. Bruyant et al 2002b) or in list mode PET (eg. Klein et al 2001, Klein et al 2002) provides a robust measure of heart position. The amount of data available for each COM calculation can be small, this being the main limitation of a data-driven approach. The "diverging squares" method (Geckle et al 1988) and 2D-fit method (Cooper et al 1993) are two other COM techniques described in the literature. In the former, a region is positioned over the heart in each image of a cardiac sequence so as to maximise counts, and motion is estimated from the offset

between regions in successive frames. Cooper and McCandless (1995) found the "diverging squares" method tracked motion superiorly to cross-correlation, but also resulted in false-positive detection. 2D-fitting successfully followed sinusoidal and linear, random and non-random movements for both single and dual-90° detectors (down to 0.7mm amplitude). Movements of this size were found to be below the threshold of cross-correlation (Eisner et al 1987).

Marker Techniques: Various authors (eg. Friedman et al 1989, Botvinick et al 1993) have used external radioactive markers to detect gross heart motion. Groch et al (1985) tracked the cardiac centroid in sequential projections using an external ^{241}Am marker. Low counting statistics in the ^{241}Am window and inability to detect or correct motion perpendicular to the imaging plane limited this technique. Furthermore, upward "creep" of the heart post-exercise caused the chest marker to become uncorrelated with internal heart movement. Germano et al (1993) also corrected gross patient motion in myocardial SPECT studies by tracking a point source placed on the sternum. The motion histogram for the source centroid over the sequence of projections indicated peaks at the significant positions held during the study. Ivanovic et al (2001) used ^{22}Na point sources attached to phantoms to enable robust, automated registration of multiple, fast scans. However, in real studies, use of external radioactive markers in well-defined anatomical locations such as the xiphoid or segments of the spine may not enable consistent tracking of rigid-body heart movement and does not enable tracking of internal motion (Geckle et al 1988). Moreover, internal heart motion can mask rigid-body motion (Friedman et al 1989, Botvinick et al 1993).

Reprojection techniques: Motion in SPECT can be detected and estimated by comparing the measured projections with projections generated from the reconstruction (Lee and Barber 1998, Arata et al 1995, Matsumoto et al 2001, Hutton et al 2002, Kyme et al 2001, Kyme et al 2003). Apart from Hutton and Kyme, the techniques described in the literature are restricted to detection of translational motion in the projection plane.

2B.5 Motion Correction

The second stage in solving the motion problem involves applying information from the detection stage to perform correction. For tomographic modalities, motion correction techniques may be classified broadly as projection-based or volume-based. This distinction is drawn based on whether motion correction is achieved primarily through processing of raw projections or reconstructed data respectively.

2B.5.1 Projection-Based Correction

In projection-based methods, motion correction is achieved by obtaining geometric consistency in the projection data prior to reconstructing. The most common way of doing this is to retrospectively align the projection data in accordance with motion estimates that have been derived from an external device, the

data itself, or visually by the user. These methods seek to utilise all of the acquired data. The alternative is to leave out data that is not geometrically consistent, an approach referred to as ‘motion purging’.

Eisner et al (1987) and Passalaqua and Narayanaswamy (1995) performed shifting of projections based on the cross-correlation between successive projections in a scan. Germano et al (1993) also used this approach but tried to enhance the accuracy using a point source on the sternum. Axial and transaxial translations were corrected with an accuracy of 0.5mm. In multi-gated blood-pool studies, Groch et al (1985) determined patient motion by tracking the coordinates of an ^{241}Am point source attached to the patient, and applied these data as correction offsets to ^{99m}Tc data to correct for rigid-body heart motion of exercising patients. The “diverging squares” method was used by Geckle et al (1988) to obtain translational offsets necessary to align motion-corrupted projections. They demonstrated a reduction in streak artifacts, hot spots and defects caused by 2.5cm in-plane translations, though the presence of background counts did cause images to become worse in some cases. O’Connor et al (1998) showed that simple manual shifting of projection data was effective in correcting small, continuous and bouncing sinusoidal motions, and linear random and non-random translations in a cardiac SPECT phantom.

Passalaqua and Narayanaswamy (1995) shifted projections based on the cross-correlation between projections in a fast scan (assumed to be motion free) and a subsequent slow scan. After alignment, the corrected slow-scan data were reconstructed and summed with the fast-scan data to produce a motion-corrected study. Compared to cross-correlation between successive projections from a single scan, the dual-scan correction approach produced better visual and quantitative results. Bruyant et al (2002b) used tracking of the axial COM in dynamically acquired SPECT cardiac data to determine the translational shifts necessary to correct for respiration-induced cardiac motion. These shifts were applied to align the projection data prior to reconstruction.

A number of authors have tried to align measured projections based on a correlation between these data and projections generated from the reconstruction (eg. Huang and Yu 1992, Lee and Barber 1998, Arata et al 1995, Matsumoto et al 2001). Translational shifts in the projection plane (Arata et al 1995, Matsumoto et al 2001) and axially (Lee and Barber 1998) were demonstrated initially. Arata applied various filters to the raw data and reprojected data prior to the correlation. Lee and Barber recognised that applying the method iteratively may improve the motion estimates. Surova-Trojanova et al (2000), Hutton et al (2002), and Kyme et al (2003) have demonstrated the applicability of the approach to general 3D motion by incorporating an optimisation stage.

Analytic and iterative reconstruction algorithms have been modified to incorporate motion information and account for motion directly. For example Li et al (1995a) described a filtered back-projection algorithm for fanbeam SPECT to compensate for in-plane motion. The same group described an approximate 3D extension to this technique that included axial tilt correction (Li et al 1995b). Analytic

reconstruction techniques only enable handling of simplistic motions. In order to account for fully general motion, 3D geometric consistency between projections must be achieved, i.e. each projection must be incorporated into the reconstruction consistent with the 3D location in space at which it was acquired. This has been achieved using iterative techniques such as OSEM (eg. Fulton et al 1994, Fulton et al 1999). Here, patient motion was conceptualised as groups of projections being acquired on the surface of different cylinders, translated and/or rotated with respect to each other. By applying the inverse movement to the current reconstruction estimate, projections on any particular "cylinder" were used to update the estimate correctly. The present work is based on this approach to motion correction.

Traditionally SPECT and PET data are binned into 2D projections (frames) with each one containing photons detected over a particular area for a period of seconds to minutes. During this time the patient is free to move and thus induce inconsistencies in the data. Subsequently, many motion correction techniques in the literature rely on an assumption that the patient moves only between frames (eg. Weinzapfel and Hutchins 2001, Ivanovic et al 2001). High frame acquisition rates are necessary for this assumption to be reasonable (Fulton et al 2002). More recently, list mode data acquisition has been used in PET to circumvent this assumption. However, it is feasible to correct for motion within a frame by considering the object voxels as having a time varying position, and then incorporating a time-averaged probability matrix[1] in the iterative reconstruction algorithm. With this adaptation, abrupt movements during a frame were shown to be correctable (Fulton et al 1993).

In PET, obtaining motion-free raw data by geometrically transforming lines of response (LORs) in accordance with motion knowledge was proposed initially by Daube-Witherspoon et al (1990). Numerous authors have since adopted the approach. For example Menke et al (1996) implemented the geometric shifting of LORs and demonstrated improved speed using a look-up table of LOR correction factors. More recently authors have demonstrated feasibility of this method in list mode PET where each coincidence event is corrected based on measurements by the Polaris motion tracker (eg., Watabe et al 2002a, Woo et al 2002, Bloomfield et al 2003). A hardware concept for achieving this in real time has been proposed by Jones (2001). Error propagation experiments on clinical data indicated that brain images could be corrected for motion with < 2% error using this method (Watabe et al 2002b). Bloomfield provides a thorough description of the technique and includes normalisation corrections for detector efficiency and geometry, crystal interference and dead time (also Qi and Huesman 2002). Corrected sinograms were rebinned to 2D before being reconstructed with 2D FBP. Other authors have focussed on issues related to the reconstruction of the motion-corrected sinograms, in particular incorporating motion information into EM reconstruction algorithms (eg. Thielemans and Mustafovic 2003, Rahmim et al 2003). Current work is focussed on adequately accounting for re-positioned LORs falling outside the FOV and instances where the object leaves the FOV for some part of the scan.

[1] The transition matrix modelling the emission and detection process is described in §*2A.1*.

Lu and Mackie (2002) described a technique for emission CT applications in which the sinogram trace of nodal density/contrast regions was used to estimate an expected trace. The sinogram was interpolated and scaled so as to coincide with the expected trace and produce a motion-corrected sinogram. Reasonable correction was obtained for data corrupted by simulated respiratory motion. The technique is limited to motions with a simple mathematical description.

Although DSA is not a tomographic modality, motion correction is projection-based, and is regularly required to compensate for tiny voluntary and involuntary motions occurring between the two images (eg. Buzug et al 1998). For example, Levin et al (1984) used simple manual shifting to align images and Pickens and Price (1987) used a series of fiducial markers on the patient to drive a polynomial warping of the contrast image to the baseline image.

In X-ray CT imaging, Dhanantwari et al (2001) described a data-driven motion correction approach based on knowledge of the phase-amplitude stamp of the motion. From this, coherent (motion-free) sinograms were compiled at selected phases of interest. The technique is appropriate for cyclical motion where organ positions are repeated over time without necessarily being periodic.

Methods based on the exclusion of corrupt data can be used to correct general 3D motion since they do not depend on accurate estimates of motion. Such methods however do require robust detection of motion. The temporal image fractionation protocol for SPECT (Germano et al 1994) is based on rotating a triple-detector camera repeatedly through a 120° partial orbit in alternating directions to acquire multiple projections at each angle. The patient has a greater opportunity to remain still since up to two out of the three projections at each angle may be purged. Alternation of the rotation reduces susceptibility to directional movements that cause artifacts (eg. upward “creep” of the heart). The method was shown to correct for axial and lateral non-returning motion and simulated upward “creep”, as well as improve perfusion artifacts in clinical cardiac studies. Cross-correlation between corresponding projections from multiple scans was used by Ivanovic et al (1998) and Pellot-Barakat et al (1998) to determine the maximum amount of consistent data at each projection angle. This was then assimilated to produce a final motion-free set in which the inconsistent data was left out. They showed that completely motion-free sets could be reconstituted provided the patient was motionless for at least $^1/_3$ of the study. Applying the method to sets corrupted with axial translations (0.5cm) and anticlockwise rotations (0.5°), they were able to generate reconstructions extremely similar to the motion-free reference.

2B.5.2 Volume-Based Correction

As the name suggests, volume-based methods operate in the object domain. The most common approach involves registration of multiple volume images (each with negligible/minimal motion) using motion estimates derived from an external device (eg. motion tracker), from an automated registration program, or visually.

Green et al (1994) proposed the idea of registering and summing reconstructed PET data obtained from a number of fast, short acquisitions rather than a single, long study. Another approach by Picard and Thompson (1997) involves segregating 2D PET data into a new frame each time head movement above a given threshold is detected by a motion monitoring system. The separate frames are reconstructed and aligned to a common reference based on the motion measurements, then summed to produce the final image. Fulton et al (2002) demonstrated correction of head motion in a volunteer by applying this method using six DOF motion data from the Polaris to align the frames. The approach was also extended to 3D PET acquisitions with appropriate compensation for attenuation and scatter (Fulton et al 2001). Moreover, a variation on this technique involving list mode data acquisition has been shown to correct for motion that would normally occur within a frame. Here, list mode data are binned into separate frames in response to the detection of significant movements. Each frame is reconstructed separately and aligned according to the Polaris measurements.

Summing multiple dynamic frames of projection data can reduce motion artifacts through averaging provided the majority of projections represent a single brain location (Ivanovic et al 1998, Pellot-Barakat et al 1998). Moreover, summing dynamic frames acquired using an alternating direction of rotation can also reduce motion artifacts in the reconstruction by averaging uni-directional and drifting motions (Germano et al 1994).

Menke et al (1996) described a volume-based method involving local convolution of corrupt reconstructed slices using kernels derived from motion knowledge. Pellot-Barakat et al (2001) described a technique for hybrid PET/SPECT systems where partial sinograms of events at each gantry angle get cross-correlated across multiple temporal sets. Pre-motion and post-motion data are reconstructed separately, then registered and summed for motion correction. Ivanovic et al (2001) used a more simplistic approach involving ^{22}Na point source markers to drive the registration of reconstructions from multiple fast scans.

In fMRI, signal changes due to activation paradigms are generally as small as 2-5%. Therefore it is critical to compensate for motion between reconstructed temporal frames. Correcting for the motion is complicated by inhomogeneity in the external static and rotating magnetic fields, inhomogeneity in the magnetisation of the head, and correlation between the stimulus and motion (eg. Freire et al 2002, Calhoun et al 1998). The latter complication has also been demonstrated in PET (eg. Pappata et al 2002). fMRI time series motion correction is typically implemented by rigid registration of the multiple reconstructed images (eg. Woods et al 1992, Woods et al 1993). A Fourier-space approach has been described by Maas et al (1997). Here, rotational mis-registration between images in the time series is determined using phase-matched filtering. After reversal of the rotation, translational mis-registration is corrected by multiplication with a complex exponential. This technique was optimised with respect to interpolation and speed by Calhoun et al (1998). The “navigator echo” technique (eg. Lee et al 1996) is

another method for fMRI motion correction. An initial echo is used to determine the correct plane for the imaging echo. A similar idea was used in diffusion-weighted fMRI (Anderson and Gore 1994, Ulug et al 1995).

2B.6 Summary

In summary, patient motion is a well-recognised problem in medical imaging and has been particularly well documented in the nuclear medicine modalities SPECT and PET. The lengthy scan times in nuclear medicine has lead to a plethora of methods being developed to rectify the problem. These include both prevention and correction techniques. However, despite some form of restraint generally being used in practice, robust motion correction solutions remain elusive.

CHAPTER 3

GENERAL METHODOLOGY AND IMPLEMENTATION

The principles underlying our novel approach to 3D motion correction and the feasibility of implementation have been described briefly in Chapter 1 and detailed in previous work by this group (Hutton et al 2002, Kyme et al 2001). In this chapter we present formally the comprehensive algorithm used for the phantom and real data validation experiments detailed in Chapters 4-6.

3.1 General Methodology

3.1.1 Definitions

We wish to estimate the radionuclide distribution $g(r)$, where r denotes position in the 3D-object space. We define the acquired (measured) projections, in projection space s, according to

$$P(s) = \{p_i(s) : i \in I\}, \tag{1}$$

where p_i is the i-th simultaneously-acquired group of projections, i is the set $\{0, 1, \ldots, i_{max}\}$, and i_{max} is given by

$$i_{\max} = \frac{total\ \ projections}{number\ \ of\ \ camera\ \ detectors} - 1. \tag{2}$$

Each p_i contains the same number of projections, equal to the number of detectors on the gamma camera.

Let $\mathbf{F}_i$ be the operation of forward projecting (reprojecting) at the angles corresponding to the i-th group. We also define the process

$$g_s^* = \mathbf{R}[Q, g_s], \tag{3}$$

where $g_s^*(r)$ is the updated reconstruction obtained when an iterative reconstruction algorithm, $\mathbf{R}$, is used to reconstruct an arbitrary set of projections, $Q(s)$, beginning with an initial image estimate, $g_s(r)$. (Note that the position variables, r and s, have been left out for simplicity.)

3.1.2 Identification of Misaligned Projections

We obtain the first estimate, $\hat{g}_u$, by reconstructing the full set of measured projections, P, using a flat (grey) image, U, as the starting reconstruction, viz.

$$\hat{g}_u = \mathbf{R}\left[P, U\right]. \tag{4}$$

The subscript denotes that this reconstruction is uncorrected and therefore may contain motion artifacts.

To identify groups of projections corresponding to discrete locations of the brain, we compute the square of the norm of the difference between the forward projections and measured projections. The similarity measure for two discrete (projection or object) functions, x and y, is given by

$$\mathrm{C}(x, y) = \frac{\left\|x - y\right\|^2}{N_x}, \tag{5}$$

where N_x is the number of non-zero elements in x. If for each i, the value Γ given by

$$\Gamma = \mathrm{C}\left(p_i, \mathbf{F}_i\left[\hat{g}_u\right]\right) \tag{6}$$

is calculated, sets of p_i can be identified for which the position/orientation of the brain was fixed. In other words, we identify $P_m \subseteq P$ such that

$$P_m = \left\{p_i : i \in I_m\right\}. \tag{7}$$

Here I_m is a subset of I containing the indices of all acquired groups p_i that were collected with the brain at, or 'close to', location m.

3.1.3 Estimation of Motion

To estimate motion parameters for each change in brain location, consider using a transformation operator **T** to apply a rigid-body transformation (3 rotations, 3 translations) in the object space. The aim is to choose the transformation so as to minimise $\mathrm{C}(P_m, P'_m)$, where P'_m here is given by

$$P'_m = \left\{\mathbf{F}_i\left[\hat{g}_u\left(\mathbf{T}\right)\right] : i \in I_m\right\}. \tag{8}$$

P'_m is thus the set of forward projections generated from (transformed) $\hat{g}_u$ at all angles identified as belonging to movement m. Denoting the optimised transformation from location *m-1* to location m as $\mathbf{T}_m$, the set of $\mathbf{T}_m$'s ($m \geq 1$) are sought using a direct-search optimisation algorithm.

3.1.4 Motion Correction

We form the partial reconstruction using projections acquired at the *m = 0* brain location

$$\hat{g}_c^{(0)} = \mathbf{R}\left[P_0, U\right]. \tag{9}$$

The superscript on $\hat{g}_c$ indicates that this reconstruction is taken to be the initial motion-corrected estimate. If we define a cumulative transformation operator according to

$$\begin{aligned} \mathbf{S}_0 &= Id \ , \ m = 0 \\ \mathbf{S}_m &= \mathbf{T}_m . \mathbf{T}_{m-1} ..... \mathbf{T}_1 \ , \ m > 0 \ , \end{aligned} \tag{10}$$

then updated motion-corrected estimates (m>0) are given by the recurrence relation

$$\hat{g}_c^{(m)}(\mathbf{S}_m) = \mathbf{R}\left[P_m, \hat{g}_c^{(m-1)}(\mathbf{T}_m . \mathbf{S}_{m-1})\right]. \tag{11}$$

When $m = M = max\{m\}$, the motion-corrected result is obtained. This can be returned to the initial frame of reference by applying the inverse cumulative transformation, $\mathbf{S}_M^{-1}$. The general methodology is illustrated in Figure 3.1.

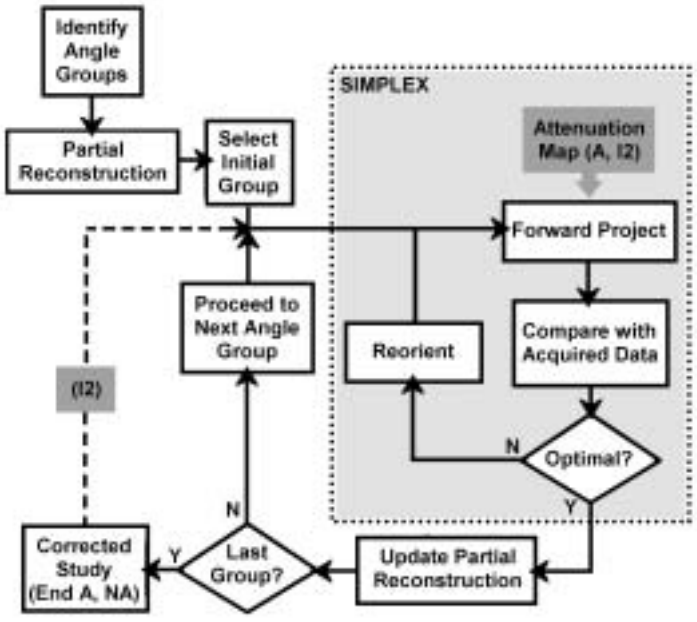

Figure 3.1. Flow model showing the basic methodology of data-driven motion correction. Various modes of correction (*A*, *NA*, *I2*) described in chapter 4 are also illustrated here.

3.2 Implementation

3.2.1 Detector Geometry

The methodology in §*3.1* is applicable to any multi-detector gamma camera. Components of rotational or translational head motion that are perpendicular to the face of a detector cannot be resolved using a comparison of the projection images obtained from that detector [Milanfar 1999]. This means that to detect all rigid-body DOF, a single detector is inadequate – as is any pair of diametrically opposed detectors. For this technique, natural choices of detector geometry would be an orthogonal configuration[1] in the dual-detector case, and an equilateral triangle configuration (detectors spaced at intervals of 120°) in the triple-detector case. We simulated dual-90° detector acquisitions for the digital phantom experiments (Chapter 4), and obtained triple-120° detector acquisitions for the physical phantom experiments (Chapter 5) and volunteer experiments (Chapter 6).

3.2.2 Reconstruction

The operator **R** was defined above as any iterative, tomographic reconstruction algorithm capable of updating some specified start image with an arbitrary set of projections. We chose the OSEM algorithm (Hudson and Larkin 1994). The projection sets, P_m, used in each motion-correction update were divided into OSEM subsets. For the digital phantom experiments in Chapter 4, a subset size of 2 was used when $^1/_3$ of the data or less were available, otherwise a subset size of 4 was used. For the physical phantom and clinical experiments in Chapters 5 and 6, simultaneously-acquired projections were grouped together using a subset size of 3.

It has been suggested (eg. Hudson and Larkin 1994) that OSEM reconstructions preferably should be based on subsets of well-dispersed angles so that the information added per sub-iteration is maximised. In the present work, we applied this principle of a balanced addition of the available projection data by manually creating OSEM files with high angular dispersion within and between subsets. Also, where possible, $\{P_m\}$ were chosen so that consecutive m's did not correspond to spatially adjacent angle groups (see Figure 4.5). The potential limitations of OSEM with respect to subset structuring and possible alternatives that may prove more robust are discussed further in Chapter 7.

3.2.3 Optimisation

The similarity measure was minimised using the downhill-Simplex algorithm (Nelder and Mead 1965, Press et al 1992). This optimisation was terminated if successive function evaluations differed by less than a specified tolerance or if a specified maximum number of function evaluations were computed. We used a tolerance of 10^{-5} and set the maximum number of function evaluations to between 250 and 350. The downhill-Simplex algorithm was chosen for its simplicity of implementation, its common use in

[1] Robinson and Milanfar (2001) sought to estimate 2D motion fields by estimating the components of projected motion for at least two independent directions. Using orthogonal views they demonstrated excellent estimation of 2D translational motion.

registration problems, and its good performance in reasonably well behaved parameter spaces (Maes et al 1999).

3.2.4 Detection of Motion Groups

The first stage of the algorithm is to reconstruct $\hat{g}_u$ using all the measured data and then to compute the similarity function from (6) for each allowed value of i. Plots of these data are used to indicate differentiated groups of projections. Even if clearly differentiated groups are visible, this plot alone is not sufficient to define the movement groups P_m conclusively since completely different motions could easily result in the same MSD. The limitation is true of all motion detection methods that are 'blind' to the actual motion. Such methods rely on the presence of inconsistencies within the data, and although they enable projections to be classified as 'aligned' or 'misaligned', without additional information, they provide no correlation between a particular degree of inconsistency and specific movement. In our case, we seek to obtain additional information regarding the start and end points of head movements by transforming $\hat{g}_u$ prior to forward projection. The transformation in this case is an arbitrary set of rigid-body parameters. We used a random number generator to obtain transformations within specified bounds. Computing $\Gamma = C(p_i, \mathbf{F}_i[\hat{g}_u])$ as a function of i after applying each of these transformations to $\hat{g}_u$ enables a more conclusive picture of the motion-relationships between angles to be established. An example of this technique is shown in Figure 4.3.

In the present work, identification of motion groups from the multiple $\Gamma = C(p_i, \mathbf{F}_i[\hat{g}_u])$ versus i curves was done visually. However, since the variation of $\Gamma = C(p_i, \mathbf{F}_i[\hat{g}_u])$ within a motion group tends to be much smaller than between groups (Hutton et al 2002, Kyme et al 2001), an automatic peak-finding and correlation method to ascertain the motion groups from the multiple plots is feasible and will be one of the avenues of future development of this technique.

3.2.5 Collimator Geometry

It was stated above that motion components perpendicular to the face of a detector are 'invisible' to that detector. Technically this is only true for parallel-hole collimation. In the case of fanbeam collimation, the slanted holes enable motion components perpendicular to the detector face to be measured via a distance-dependent magnification effect. Fanbeam collimator geometry was investigated to determine the extent to which this effect enhances motion detectability. Simple translational motion away from a single detector was simulated using the Hoffman brain phantom[1] (Hoffman et al 1990). The magnitude of simulated motion was increased from 1 to 6 pixels (pixel size = 4.4mm). The resulting motion-corrupted reconstructions were reprojected at the relevant angle and the projection compared with the measured projection for this angle. A repeat of the experiment was performed for parallel-hole collimation.

[1] Note that a full description of how data were simulated from the Hoffman brain phantom is provided in Chapter 4 as part of the digital phantom validation experiments.

The results are plotted in Figure 3.2. Similarity between the reprojected and measured projections is recorded on the vertical axis. Movement perpendicular to the detector was hardly resolved using parallel-hole collimation. By contrast, this movement was detected using fanbeam collimation, and the detectability increased with motion magnitude. However, MSD values were considerably smaller than what has been observed previously for 1-3 pixel motion using orthogonal parallel-hole detectors (Hutton et al 2002), suggesting the additional detection benefit using fanbeam collimation is small. In the present work we have chosen to use a multi-detector parallel-hole regimen in all experiments.

3.2.6 Partial Reconstruction

Instead of optimising the orientation of $\hat{g}_u$ (uncorrected reconstruction) to estimate the motion parameters for each movement (§*3.1.3*), a partial reconstruction can be used. The partial reconstruction is generated from a consistent, but incomplete, set of projection data. First we form the initial motion-corrected estimate as in (9) but choose P_0 as the largest of the P_m's. Then the $\mathbf{T}_m$'s are obtained by minimising the similarity function $\Gamma = C(P_m, P'_m)$, where P'_m in this case is given by

$$P'_m = \left\{ \mathbf{F}_i \left[\hat{g}_c^{(m-1)} \left(\mathbf{S}_{m-1} \right) \right], \, i \in I_m \right\}. \qquad (12)$$

P'_m is therefore the set of forward projections generated from the (m-1)th estimate at all angles identified as belonging to movement m. Note here that $m > 0$ and updates proceed as before.

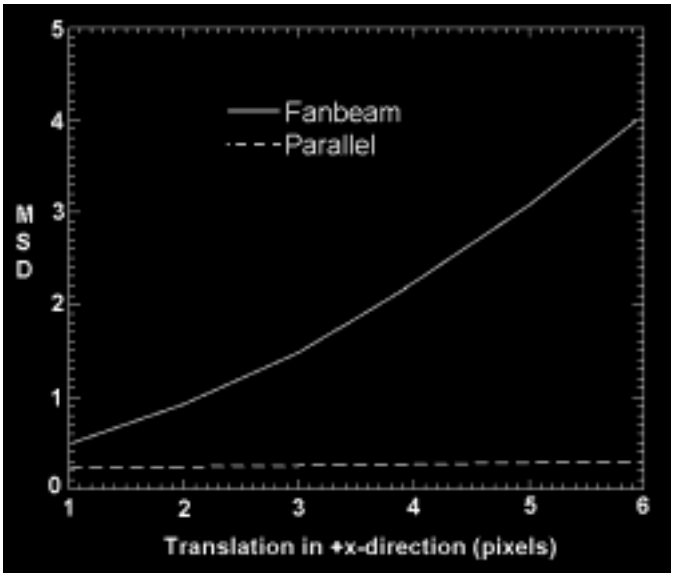

Figure 3.2. Mean square difference between acquired and generated projections for simulated translational motion normal to a single detector. The motion was practically 'invisible' for parallel-hole collimation at all magnitudes. Fanbeam geometry enabled detection, this improving with translation magnitude. (Pixel size = 4.4mm.)

3.2.7 Basis for Using a Partial Reconstruction

Here we provide justification for §*3.2.6* – that is, use of a partial reconstruction in order to get improved motion parameter estimates. Two experiments were conducted to determine the impact of partial reconstruction on the detection and correction of motion respectively.

In the first experiment, two complete sets of projections, *P1* and *P2*, were generated by analytically projecting the digital Hoffman phantom at two different orientations. 'Hybrid' projection sets, H_i, containing an increasing proportion of *P2* projections were formed. To determine whether forward projections generated from the fully or partially reconstructed data enabled better differentiation between *P1* and *P2*, a distinguishability index was calculated as

$$d = \mathrm{C}\left(P1, \mathbf{FR}\left[H_i^*, U\right]\right) - \mathrm{C}\left(P2, \mathbf{FR}\left[H_i^*, U\right]\right), \qquad (13)$$

where, for the full reconstruction case H_i^* is equivalent to H_i, and for partial reconstruction case H_i^* is the subset of H_i belonging to *P1*. Operator **F** (without subscript) denotes forward projection at all angles. The experiment is illustrated in Figure 3.3.

Results are shown in Figure 3.4. This is a plot of *d* versus the number of *P1* projections contained in the full and partial reconstructions. For the full reconstruction case, the effectiveness of using the forward projections to differentiate *P1* and *P2* data rapidly declined to a minimum as the proportion of *P2* data in the hybrid set, *H,* increased to 50%. By comparison, forward projections generated from the partial reconstruction enabled excellent differentiation even when $< {}^1/_4$ of the *P1* data were used to reconstruct. On the assumption that improved motion parameter estimation results when there is improved differentiability of data, these results suggest that use of a partial reconstruction should provide more robust motion correction.

In the second experiment, two noisy motion-corrupted datasets were simulated from the digital Hoffman brain phantom[1]. The two datasets correspond to sets 1 and 7 respectively in Figure 4.2. Set 1 contained a single motion (two brain positions), and set 2 contained three motions (four brain positions). The actual movements simulated are provided in Table 4.1. The movements in each set were estimated by optimising the orientation of (a) the full reconstruction (reconstructed from the full set of projections) and (b) a partial reconstruction (reconstructed from a subset of consistent projections). These movement estimates were compared with the applied values.

[1] Note that a full description of how data were simulated from the Hoffman brain phantom is provided in Chapter 4 as part of the digital phantom validation experiments.

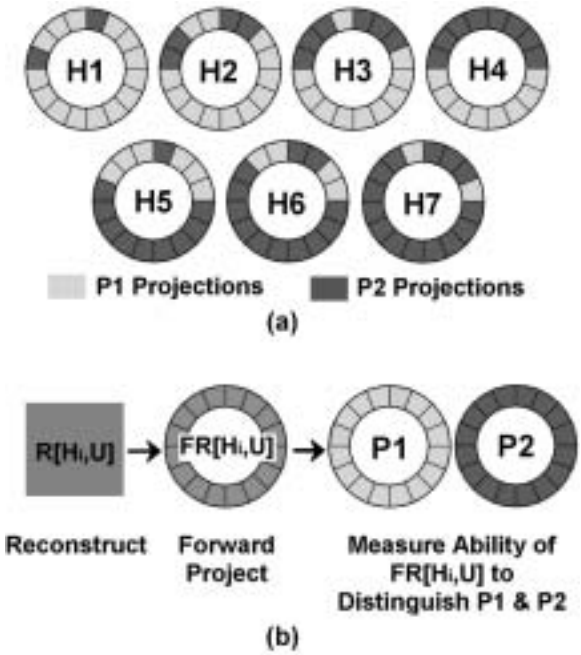

Figure 3.3. (a) The hybrid projection sets formed from *P1* by increasing the proportion of *P2* projections. (b) Each hybrid set was reconstructed and the resulting reconstruction forward-projected. Using (13) we estimated the ability of the full reconstruction to distinguish between *P1* and *P2* data. Performance of the partial reconstruction to differentiate *P1* and *P2* data was estimated in much the same way but in this case H_i was made up of *P1* data exclusively.

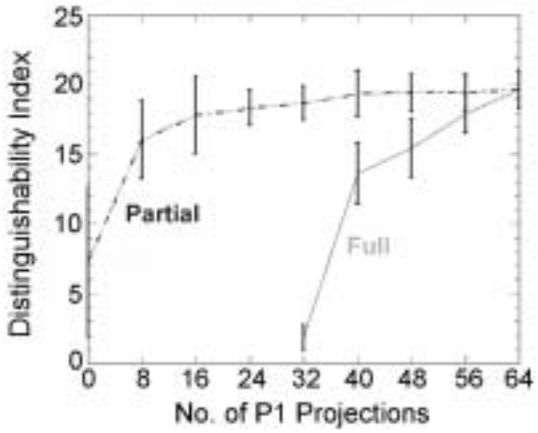

Figure 3.4. Results showing how effectively the full and partial reconstructions enabled the sets *P1* and *P2* to be distinguished (see Figure 3.3). Error bars indicate the standard deviation of MSD/projection values used to derive the distinguishability index.

Table 3.1 shows the movement parameters extracted using the full and partial reconstruction regimes for the two datasets. The performance of the two regimes was comparable for the lesser corruption (dataset 1). Note the large deviations in rotational parameters for dataset 2 using the full reconstruction and the general improvement in accuracy when a partial reconstruction was used.

In the present work a partial reconstruction approach was used in all motion correction experiments in order to enhance our ability to detect and correct for motion.

Table 3.1. Movement parameters extracted for the two datasets using a partial and full reconstruction. For dataset 2 the movements were estimated in the order: movement 2 (16 projections), movement 1 (16 projections), movement 3 (16 projections). Rotations are in degrees and translations in pixels.

	Movement 1						**Movement 2**						**Movement 3**					
Dataset 1	**X°**	**Y°**	**Z°**	**X**	**Y**	**Z**												
Applied	-8.0	-3.0	5.0	-1.0	0.5	-2.0												
Full	-8.3	-1.2	4.3	-0.8	0.4	-2.0												
Partial	-6.2	-2.2	4.9	-1.0	0.8	-1.9												
Dataset 2	**X°**	**Y°**	**Z°**	**X**	**Y**	**Z**	**X°**	**Y°**	**Z°**	**X**	**Y**	**Z**	**X°**	**Y°**	**Z°**	**X**	**Y**	**Z**
Applied	0.0	4.0	4.0	0.0	1.5	1.0	2.0	7.0	7.0	0.5	2.0	1.5	4.0	10.0	7.0	1.0	2.5	2.0
Full	1.4	3.7	3.3	0.1	1.5	1.1	2.4	4.3	3.4	0.4	2.0	1.4	2.5	4.0	2.9	0.9	2.2	1.7
Partial	0.1	5.1	7.2	0.0	1.5	1.0	1.5	6.5	5.3	0.4	1.8	1.4	4.1	11.2	8.9	0.9	2.5	1.9

CHAPTER 4

DIGITAL PHANTOM VALIDATION EXPERIMENTS

4.1 Data Simulation

Projection data for the simulations were obtained from a 128 x 128 x 80 voxel version of the digital Hoffman brain phantom (Hoffman et al 1990). Voxels were 2.2 mm in each dimension. This phantom was projected using in-house analytical projection software. We modelled high-resolution parallel-hole collimation with depth-dependent blur and included uniform tissue attenuation. Scatter was not modelled. To generate the attenuation map, the slices of the phantom were converted to binary, filled to give uniform closed regions, Gaussian filtered to enlarge the extent of the brain, and lastly converted back to binary grey levels. Thus the attenuation slices followed the contour of the corresponding brain slices but were slightly larger to account for skull attenuation. The attenuation coefficient modelled was $\mu = 0.12$ cm^{-1}. Using the projector we produced 64 emission projections of dimension 64 x 40 pixels. Pixels were square with length 4.4 mm. To model Poission noise in the projection data, each pixel was replaced by a random Poisson deviate using the actual grey value as the distribution mean. The Poisson-distributed values were generated after scaling the data so that the maximum total counts per projection were 50k. This value was judged to give a level of noise equivalent to a typical clinical brain study.

To simulate projection sets corrupted by patient motion, various rigid-body transformations were applied to the digital Hoffman brain phantom. After each transformation it was projected as above. Projections from the resulting sets were then combined to simulate a motion-corrupted acquisition. Since we were modelling dual-90° detector geometry, projections were always transferred as orthogonal pairs. Motion was simulated to be discrete and as having occurred between projections. Motion occurring during the acquisition of a projection pair was not simulated[1]. The method is illustrated in Figure 4.1.

4.2 Datasets

Seven studies containing from one to three discrete head movements (two to four discrete head positions respectively) were simulated. The motions included 3D head tilting, twisting and sliding, i.e. feasible sharp and discrete movements that may result from coughing, discomfort, tiredness, and other involuntary actions of a real patient (Green et al 1994, Lopresti et al 1999, Fulton PhD, Barnden personal com). The

[1] Treatment of motion occurring during a projection can be found in Hutton et al (2002).

seven simulated sets were constructed to have considerable variation in the angular location and extent of movement and the magnitude of motion in each DOF. A diagrammatic representation of the motion-corrupted datasets is shown in Figure 4.2 and a summary of the simulated motions is given in Table 4.1.

4.3 Motion Detection

Projection groups were identified as described in §*3.2.4*. An example illustrating the application of this method to dataset 4 is provided in Figure 4.3.

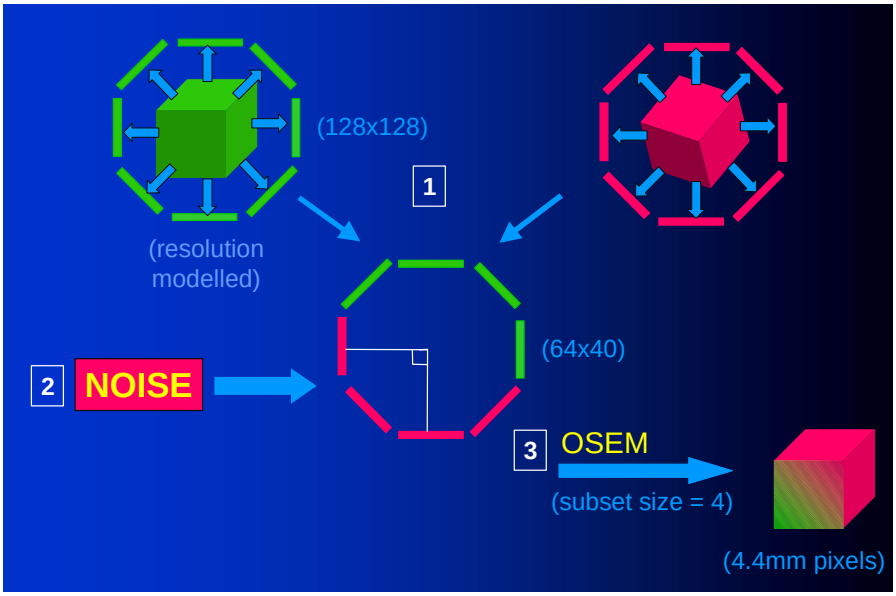

Figure 4.1. Steps involved in simulating motion-corrupted datasets.

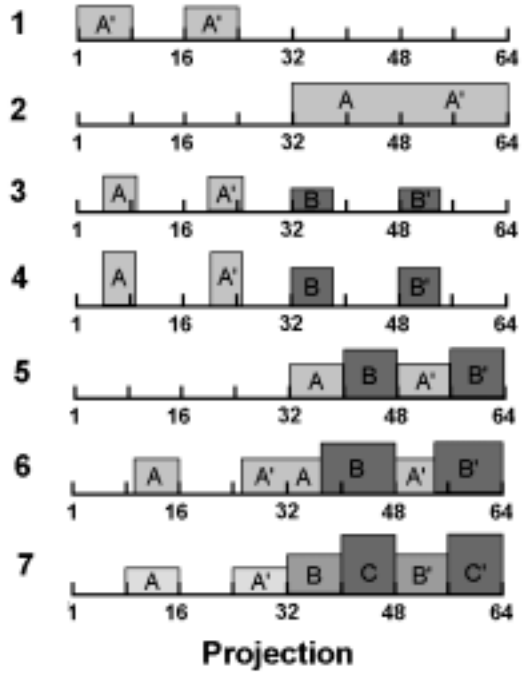

Figure. 4.2. Schematic showing the location of movements (A, B, C) incorporated into the 7 digital phantom datasets. Table 4.1 shows the parameters corresponding to each motion. A', B' and C' represent angles at 90° to A, B and C respectively.

Table 4.1. Simulated rotations (°) and translations (pixels) and the number of angle pairs affected for the 7 digital phantom datasets

Data	Label	Pairs	Applied Movements					
			X°	Y°	Z°	X	Y	Z
1	A	8	-8	-3	5	-1.0	0.5	-2.0
2	A	16	-8	-3	5	-1.0	0.5	-2.0
3	A	6	3	0	6	0.4	-0.5	-1.0
	B	6	-3	3	-5	-0.3	0.0	2.0
4	A	6	6	0	12	1.5	-1.5	-3.0
	B	6	-7	8	-10	-1.5	0.0	4.0
5	A	8	-2	5	12	-0.8	1.2	-0.2
	B	8	-3	0	6	-0.5	-0.5	-1.2
6	A	11	-2	5	12	-0.8	1.2	-0.2
	B	11	-3	0	6	-0.5	-0.5	-1.2
7	A	8	0	4	4	0.0	1.5	1.0
	B	8	2	7	7	0.5	2.0	1.5
	C	8	4	10	7	1.0	2.5	2.0

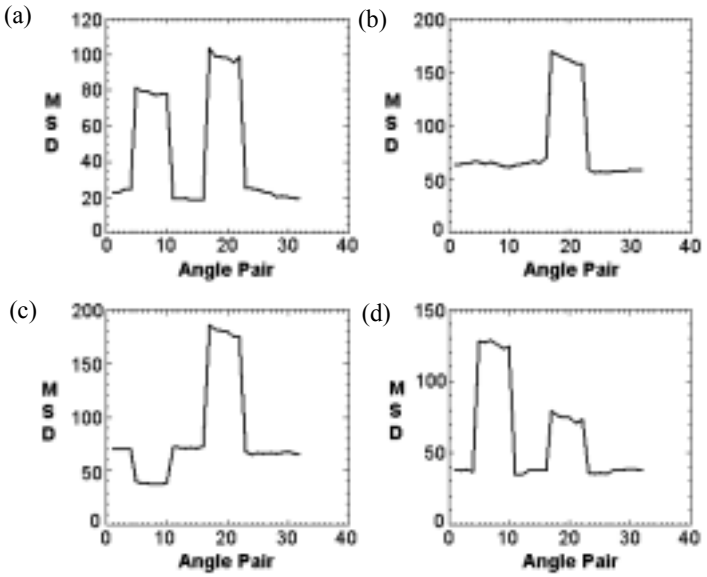

Figure 4.3. Similarity (MSD) between acquired and forward projections for the indexed (*i*) angle pairs of a simulated phantom study (dataset 4). With no transformation of the reconstruction (a) was obtained. Transforming the reconstruction 3 different ways before calculating the similarity gave (b), (c), and (d). Note that at certain orientations of the brain (e.g., (b)), some movements may not be detected. These data suggest 3 distinct angle groups (ie. 3 brain positions adopted during the acquisition): pairs {5-10}, pairs {17-22}, and pairs {1-4, 11-16, 23-32}.

4.4 Motion Correction

As described in Chapter 3, motion correction involves identification of motion groups followed by a series of Simplex-driven optimisations to estimate the brain orientation for each identified group. After each optimisation the current estimate is updated with the projections from the relevant group. Once all data are included, the resulting reconstruction is regarded as motion-corrected. Four motion correction regimes based on this basic methodology were compared for each of the 7 datasets.

Regime A ("attenuation"): Attenuation was included in the reprojection process. Recall that reprojection is used when identifying the projection groups (detection stage) and when estimating the motion for each group (optimisation stage). To ensure that the attenuation map was registered to the partial reconstruction, they were always transformed synchronously[1]. For any particular orientation of the partial reconstruction, a 2D attenuated projection was generated by:

[1] See Fulton et al (2002) for another example of this approach.

(i) Rotation of the partial reconstruction and attenuation map axially by the negative of the desired projection angle (measured from the vertical)
(ii) Formation of the cumulative sum of the attenuation map in the vertical direction
(iii) Exponentiation of the cumulative attenuation map
(iv) Voxel-by-voxel multiplication of the partial reconstruction by the exponentiated array
(v) Summation vertically through the columns

No resolution (blurring) effects were modelled in the reprojection process. Once the motion of each group had been estimated, the partial reconstruction was updated with the relevant projections. The synchronously transformed attenuation map was used to correct the current estimate for attenuation at each update.

Regime NA ("no attenuation"): It has been suggested previously that leaving attenuation out of the motion identification and estimation stages of the algorithm should not adversely affect the accuracy of motion estimation (Hutton et al 2002). We used the *NA* regime to test this hypothesis. Although attenuation was unaccounted for to obtain each motion estimate and updated reconstruction, once all estimates were obtained, the initial partial reconstruction was updated from the beginning, this time with attenuation correction. Thus the final motion-corrected studies presented here were all corrected for attenuation irrespective of whether or not attenuation was included in the motion identification and estimation stages.

Regime I2 ("second iteration"): This regime constitutes a second iteration of regime *A*. The regime *A* result was substituted for the initial partial reconstruction. Motion parameters for each group were re-estimated by optimising the orientation of this reconstruction. After each estimate of motion, no updates were performed. Once estimates for all motion groups were obtained, the initial partial reconstruction was successively updated (as for regime *A*) to produce a motion-corrected study.

Regime A^+ ("ideal"): For this regime we used the known movements as the motion estimates. This represented the best correction achievable from the current implementation.

The four correction regimes are illustrated in Figure 4.4 and Figure 3.1. For each regime the updates were carried out as described in §*3.2.2*. Projections to be included at each update were structured into subsets of size 2 or 4. A projection and its orthogonal partner (acquired simultaneously) always shared the same subset. Achieving angular dispersion of projections within and between subsets was done manually. For datasets containing multiple brain orientations, successive Simplex optimisations did not necessarily correspond to adjacent motion groups. By choosing the projection groups in an alternating fashion (Figure 4.5), the angular information provided to OSEM was further dispersed.

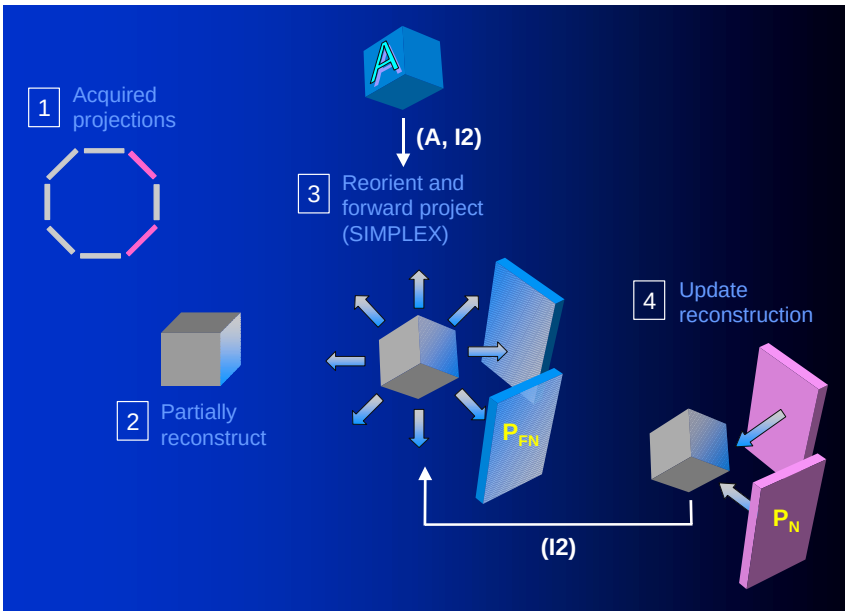

Figure 4.4. The four regimes used to estimate motion parameters and derive motion-corrected studies. All regimes contained the basic steps of (1) acquisition, (2) partial reconstruction, (3) optimisation, and (4) updating. Regime *A* and *I2* involved input of an attenuation map in step (3). This map was reoriented according to the emission volume and used to generate attenuated reprojections. Once the corrected study was obtained using regime *A*, a second iteration (*I2*) could be performed using this reconstruction in place of a partial reconstruction.

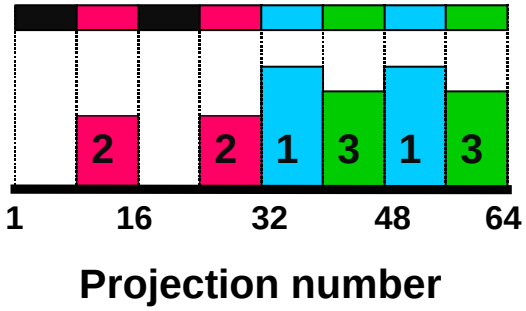

Figure 4.5. When multiple movement groups are present (represented here by the different colours), the order of optimisation and updating becomes more important. This is because the reconstruction is biased toward information contained in the particular subset used in an update. By choosing successive motion groups carefully, we can avoid adding spatially adjacent angles on consecutive updates and thereby achieve a more balanced incorporation of the projection data. Adding angles this way should give rise to a truer estimate at any particular stage, and therefore better motion estimates. In this example the black projections represent those used for the initial partial reconstruction. The order then proceeds as blue followed by red followed by green.

4.5 Analysis

Visual and quantitative assessments were performed after motion correction to measure the performance of the algorithm and the improvement in motion-induced perfusion artifacts.

For each digital phantom simulation, difference images were formed between the motion-free reference ($\hat{g}_r$) and each motion-corrupted and corrected study in order to assess residual errors. Extracted rigid-body motion parameters (from Simplex) for each movement were compared with the applied parameters. The mean registration error (MRE) was used to provide a summary figure for the quantitative accuracy of an optimisation. MRE is the average linear distance between the vertices of a bounding box enclosing the brain in the true location and the extracted location (Figure 4.6).

To quantify the overall improvement in image accuracy derived from motion correction we used a mean square difference ratio (MSDR). This was calculated as

$$MSDR = \frac{C\left(\hat{g}_r(\mathbf{s}_M), \hat{g}_u(\mathbf{s}_M)\right)}{C\left(\hat{g}_r(\mathbf{s}_M), \hat{g}_c^{(M)}(\mathbf{s}_M)\right)}. \qquad (12)$$

For this calculation, the motion-corrupted ($\hat{g}_u$) and motion-free ($\hat{g}_r$) reconstructions were transformed to the orientation of the motion-corrected reconstruction ($\hat{g}_c$). 3D convolution[1] with a Gaussian kernel (FWHM = 9 mm) was applied to the studies before measuring the MSDR so that measured differences were due primarily to corruption rather than noise. The MSDR was calculated over 19 central brain slices.

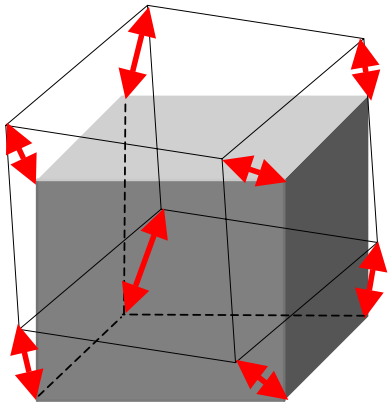

Figure 4.6. Calculation of MRE. The shaded box represents the target object and the wire frame the rigidly-registered object. Registration error was assessed by computing the average distance (pixels) between each target vertex and the corresponding registered vertex (red arrows).

[1] Convolution was first applied in-plane using a 2D kernel and then across-plane using a 1D kernel.

4.6 Results and Discussion

The accuracy of extracted motion parameters for each of the applied movements is summarised in Table 4.2 in terms of the MRE. The majority of values were considerably less than 1 pixel. Generally the MRE obtained after a second iteration of motion correction (*I2*) was an improvement on the single iteration (*A*) value as expected. There was no clear indication that including attenuation (*A*) in the motion estimation process produced better motion estimates on a consistent basis. In Figure 4.7a, the deviation (°) of the extracted x, y and z-axis rotations from the applied values is plotted for the 13 movements. Most points were clustered within 2° of the applied movement, with tight groups near 0-deviation except in the case of z-axis rotation. For translations (Figure 4.7b), most data points were within 1mm of the applied movement. The general impression of these data was that a second iteration gave slightly more accurate motion estimates, and for a single iteration, attenuation did not have a significant influence on accuracy. Clearly all of the rotational and translational DOF were transparent to the algorithm. Maximum deviation from the applied values (across all optimisations) was 3.6° (x-rotation), 4.2° (y-rotation), 5.6° (z-rotation), 1.8mm (x), 2.2mm (y), and 0.9mm (z). The mean absolute deviation, calculated as |applied – extracted|, is shown in Table 4.3 for each DOF and each correction regime. These were similar using a single iteration with and without attenuation (*A* and *NA*) but were smaller in all DOF using a second iteration (*I2*). The z-axis rotation was the most error-prone rotation parameter and the z-axis translation the most accurate translation parameter, irrespective of the correction regime used.

Upper, middle, and lower brain slices are shown for datasets 5 and 7 in Figure 4.8a and 4.8b respectively. Rows (top to bottom) correspond to motion-free (R), motion-corrupted (U), and motion-corrected (A, NA, I2) slices. Severe edge defects and distortions in perfusion were evident in the motion-corrupted slices of both datasets. After motion correction using each regime, slices resembled their motion-free counterparts much more closely. Only minor differences were apparent between the correction regimes.

The same vertical sequence of slices is shown in Figure 4.9 with the difference image between each slice and its motion-free counterpart shown alongside. The last row (A^+) corresponds to the best correction achievable, i.e. that obtained using the known parameters. For dataset 5 (Figure. 4.9a), prominent differences around the perimeter of the motion-corrupted slice were reduced by a similar extent using all correction regimes. For dataset 6 (Figure. 4.9b), distortion in the motion-corrupted slice was predominantly anterior and posterior. The bulk of these differences were removed by motion correction. A second iteration (*I2*) removed the postero-lateral difference still present after correction using regime *A*. Importantly, the same defect was removed using a single iteration without attenuation (*NA*). Clear residual differences existed after motion correction. However, the close resemblance between the difference images indicates motion-induced distortions were reduced close to the potential of the technique. Possible causes of the residual errors are noise, and interpolation incurred between OSEM updates.

TABLE 4.2. Mean registration error (MRE) for the estimate of each movement in the 7 simulated datasets using three correction regimes. Values are in pixels. A second value appears in parentheses in some cases where the movement was detected as two separate groups.

		Movement Label		
Data	Regime	A	B	C
1	A	0.72		
	NA	0.68		
	I2	0.21		
2	A	2.08		
	NA	1.05		
	I2	1.25		
3	A	0.19	1.78	
	NA	0.69	2.39	
	I2	0.20	1.28	
4	A	0.80	1.08	
	NA	1.08	1.50	
	I2	0.43	0.83	
5	A	0.38	0.32 (0.22)	
	NA	0.75	0.23 (0.81)	
	I2	0.46	0.22	
6	A	0.83 (0.66)	0.66	
	NA	0.76 (1.03)	0.23	
	I2	0.59	0.71	
7	A	1.39	0.76	0.85
	NA	0.74	0.72	0.70
	I2	0.90	0.78	0.24

TABLE 4.3. Mean absolute deviation of extracted rotations (°) and translations (pixels) from applied parameters, shown for each DOF for each correction regime

Regime	X°	Y°	Z°	X	Y	Z
A	0.63	0.79	1.5	0.57	0.75	0.31
NA	0.73	0.56	1.61	0.79	0.70	0.48
I2	0.20	0.45	1.22	0.40	0.35	0.18

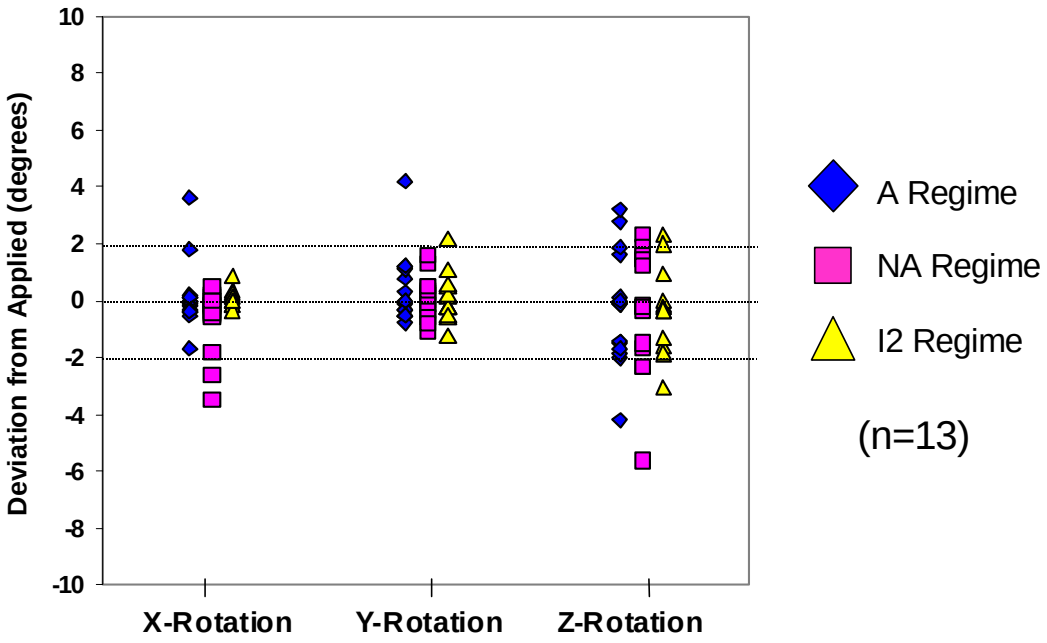

Figure 4.7a. Deviation (°) of extracted x, y, and z-axis rotations from the applied motion parameters. There were a total of 13 movements across the 7 datasets. Values are plotted here according to the correction regime used.

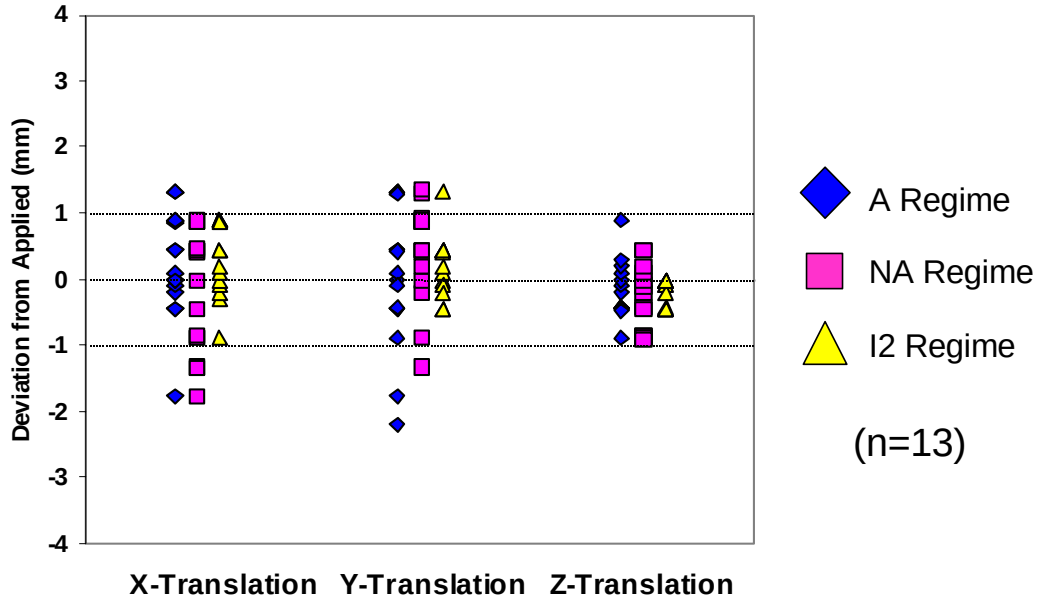

Figure 4.7b. Deviation (pixels) of extracted x, y, and z-axis translations from the applied motion parameters. There were a total of 13 movements across the 7 datasets. Values are plotted here according to the correction regime used.

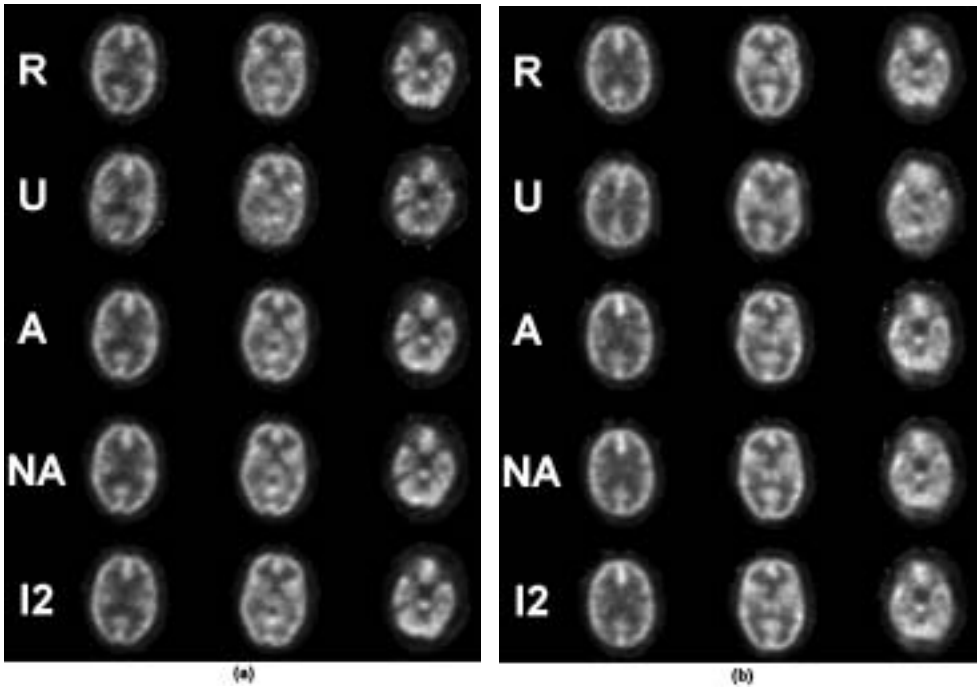

Figure. 4.8. Upper, middle, and lower (left to right) brain slices for (a) dataset 5 and (b) dataset 7. The 3 slices are shown for the motion-free (R), motion-corrupted (U), and motion-corrected (A, NA, I2) reconstructions.

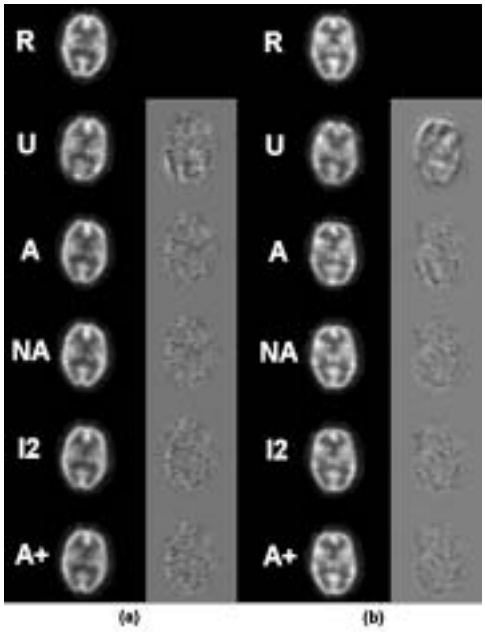

Figure. 4.9. (a) Dataset 5 and (b) dataset 7. A middle brain slice (left column) is shown for the motion-free (R), motion-corrupted (U), motion-corrected (A, NA, I2), and ideal-corrected (A+) reconstructions. Shown alongside each slice is the difference image formed by subtracting the corresponding motion-free slice. Difference images were scaled to the same range.

The MSDR values calculated for each correction regime are shown in Table 4.4. All corrections resulted in an improvement, generally by a factor of 2 or more. Comparison of the values obtained for datasets 1 and 2, 3 and 4, and 5 and 6 indicates that improvement increased with the magnitude and extent of corruption. Applying a second iteration of motion correction (*I2*) bettered the single iteration result in all cases, though this additional benefit was minor (dataset 7 the exception). This suggests most errors get corrected in the first pass. Moreover, ignoring attenuation during optimisation (*NA*) gave higher MSDR values than the *I2* regime for 3 of the datasets (2, 6, 7), and equal or marginally smaller MSDR values for the remainder.

MSDR results are expressed as a fraction of the ideal values (obtained for A^+ reconstructions) in Figure 4.10. For the *NA* and *I2* regimes this fraction was > 80% in all cases. Regime *A* became less effective as the magnitude and extent of corruption increased, being least effective for dataset 7.

TABLE 4.4. MSDR values after motion correction using the variuos regimes.

Regime	Dataset						
	1	2	3	4	5	6	7
A	2.1	3.8	1.6	2.9	3.3	4.5	5.9
NA	2.4	4.3	1.6	2.8	3.1	5.4	9.4
I2	2.4	4.2	1.6	3.0	3.4	4.8	9.3
A+	2.5	4.5	1.7	3.1	3.8	6.1	11.4

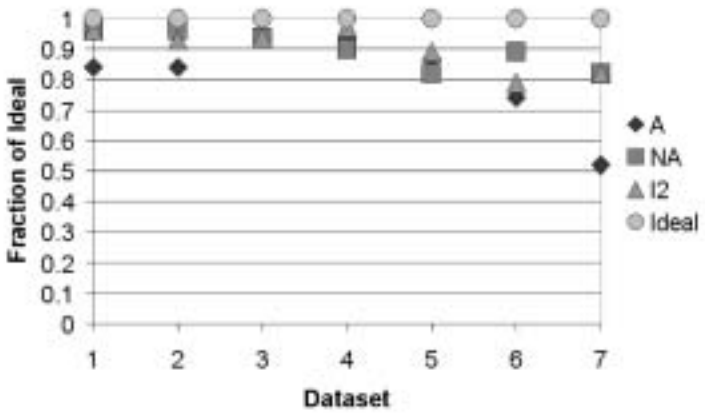

Figure. 4.10. MSDR values for each correction regime, expressed as a fraction of the MSDR for the ideal (A^+) correction.

CHAPTER 5

PHYSICAL PHANTOM VALIDATION EXPERIMENTS

After the preliminary validation experiments using simulated data, the next phase of validation was to test the method on real data, beginning with acquisitions of the physical Hoffman brain phantom. Unlike the simulations outlined in Chapter 4 in which movements were known, an independent reference (gold-standard) motion measurement was required for the real data experiments in order to be able to validate the data-driven estimates. In this chapter we provide a brief outline of the theory of rigid-body matrix transformations and then go on to describe the operation and implementation of the Polaris motion tracker used for independent motion measurement. The second half details the methods and results of the physical Hoffman brain phantom experiments performed.

5.1 Theory

A 3D vector is typically represented as the sum of scaled orthogonal unit vectors, viz.

$$\mathbf{t} = x\mathbf{x} + y\mathbf{y} + z\mathbf{z}. \tag{13}$$

Here, **x**, **y** and **z** are the unit vectors and x, y and z the corresponding scale factors. An alternative is to represent the vector using matrix notation:

$$\mathbf{t} = \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}. \tag{14}$$

The latter representation is convenient because certain geometric transformations, including translation, rotation, and scaling, can be described using a matrix. Computing the transformed vector $\mathbf{t}'$ is then equivalent to performing the matrix multiplication

$$\mathbf{t}' = \mathbf{H}\mathbf{t}\,, \tag{15}$$

where **H** is the matrix describing the transformation. In the present work, the transformations of interest are translation and rotation.

Translation by the vector $a\mathbf{x} + b\mathbf{y} + c\mathbf{z}$ is equivalent to multiplying $\mathbf{t}$ from the left by:

$$\mathbf{T}(a,b,c) = \begin{bmatrix} 1 & 0 & 0 & a \\ 0 & 1 & 0 & b \\ 0 & 0 & 1 & c \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{16}$$

Similarly, rotations about each of the x, y, and z-axes are described respectively by:

$$R_x(\psi) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\psi) & -\sin(\psi) & 0 \\ 0 & \sin(\psi) & \cos(\psi) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{17}$$

$$R_y(\theta) = \begin{bmatrix} \cos(\theta) & 0 & \sin(\theta) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\theta) & 0 & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{18}$$

$$R_z(\phi) = \begin{bmatrix} \cos(\phi) & -\sin(\phi) & 0 & 0 \\ \sin(\phi) & \cos(\phi) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{19}$$

Positive angles substituted into these rotation matrices will describe anti-clockwise rotations about the relevant axis when viewed from the positive direction toward the origin. Multiple translations or rotations applied in sequence have a description as a single transformation matrix. This matrix is calculated by successively multiplying each component matrix from the left. For example, the matrix $\mathbf{M}$ describing a general translation followed by rotation ψ about the x-axis then rotation θ about the y-axis then rotation ϕ about the z-axis is given by:

$$\mathbf{M} = \mathbf{R}_z(\phi)\mathbf{R}_y(\theta)\mathbf{R}_x(\psi)\mathbf{T}_{xyz}\ . \tag{20}$$

The multiplication of component matrices is not commutative. Therefore to generate the matrix for a desired transformation it is necessary to define an order of operations. Likewise, decomposing a transformation matrix into the component translational and rotational matrices requires a particular order of operation be specified.

Since a 4x4 matrix can describe the 3D transformation of a vector, it also describes the switch in coordinates in going from one frame of reference to a second frame that is rotated and translated with respect to the first. This property is applied in §*5.2.3*.

Quaternion mathematics (Hamilton 1969) provides an alternative means of describing 3D rotations. A quaternion is a complex number with three imaginary parts and represents a single rotation about an arbitrary axis that is not, in general, one of the primary axes. Unit quaternions are written:

$$q = q_0 + iw_x + jw_y + kw_z \ , \tag{21}$$

where q_0 is the real part, w_x, w_y and w_z are the imaginary parts, and

$$\begin{array}{lll} i.i = -1 & j.j = -1 & k.k = -1 \\ i.j = k & j.k = i & k.i = j \\ j.i = -k & k.j = -i & i.k = -j \end{array} \ , \tag{22}$$

$$q_0 = \cos\left(\frac{\theta}{2}\right) \quad \text{and} \quad \sqrt{w_x^2 + + w_y^2 + w_z^2} = \sin\left(\frac{\theta}{2}\right). \tag{23}$$

Using a single quaternion to describe any combination of rotations is compact and avoids phenomena such as gimbal lock. The conversion from a unit quaternion to standard rotation matrix form is given by

$$\begin{bmatrix} \left(q_0^2 + q_x^2 - q_y^2 - q_z^2\right) & 2\left(q_x q_y - q_0 q_z\right) & 2\left(q_x q_z + q_0 q_y\right) & 0 \\ 2\left(q_y q_x + q_0 q_z\right) & \left(q_0^2 - q_x^2 + q_y^2 - q_z^2\right) & 2\left(q_y q_z - q_0 q_x\right) & 0 \\ 2\left(q_z q_x - q_0 q_y\right) & 2\left(q_z q_y + q_0 q_x\right) & q_0^2 - q_x^2 - q_y^2 + q_z^2 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{24}$$

5.2 Motion Measurement System (Polaris)

5.2.1 Basic Operation

The Polaris is an optical motion-tracking device manufactured by Northern Digital Inc[1]. A rigid target is firmly secured to the object of interest and monitored at regular intervals by the Polaris in order to determine changes in position and orientation of the object. Two modes of operation are available. In passive mode, the Polaris exposes the field of view with infrared radiation and detects signals reflected

[1] Northern Digitial Inc., Waterloo, Ontario, Canada

back from passive reflective markers on the target. In active mode the target is made up of infrared light emitting diodes so that both the Polaris and the target emit infrared radiation. A physical connection between the Polaris and light emitting diodes is required in active mode. The Polaris was operated in passive mode for all of our experiments.

Up to three separate targets may be monitored at any one time. The Polaris samples the reflective markers on these targets at up to 20Hz in passive mode. A sample consists of a position and orientation measurement for each target being monitored. Position data are returned as three-element vectors, (p_x, p_y, p_z), and orientation data as unit quaternions, q.

5.2.2 Target

Three to six passive reflective markers are arranged on a rigid substrate. Inter-marker separation distances must be constructed in accordance with rules provided by the manufacturer. An accurate description of the target is then stored in an on-line target file. The object target consisted of a 3mm thick, grey polyethylene square substrate with four spherical infrared reflective markers attached as shown in Figure 5.1a. For the physical phantom experiments, the object target was attached to the phantom securely using adhesive tape. The reference target consisted of four commercial adhesive-backed 10mm diameter reflective discs bonded to black paper on an aluminium substrate (Figure 5.1b). This target was permanently fixed to the gamma camera gantry. The arrangement of reflector positions on the reference target was different from the object target since the former had no size restriction. Four markers were used on each target to introduce redundancy into the location measurement and thereby improve accuracy.

5.2.3 Coordinate Transformation

Since SPECT data are collected with respect to the camera coordinate system and motion data with respect to the Polaris system, it is necessary to determine the transformation between these systems in order to apply the measurements usefully for motion correction. This is achieved by attaching a tiny point source of activity to the centre of the reflective marker classified as the target origin. The target is moved to various locations in the gamma camera (and Polaris) field of view, and two orthogonal SPECT views are acquired at each location for long enough to give quality planar images of the point source. The centroid of the point source is then determined from the two views. Once the location of the source is known in the gamma camera and Polaris frames, the transformation matrix $\mathbf{T}^c$ relating these two frames can be determined according to the method detailed in Fulton (2000), *pp*. 102-105. Staff at the institution from which our data were supplied used this method to calibrate the Polaris. We were provided with the calculated $\mathbf{T}^c$ matrix.

The transformation matrix $\mathbf{T}^c$ is only valid for the particular gantry-Polaris positioning at the time the calibration was performed. Clearly this is inflexible for routine use, i.e. where the Polaris is likely to be

re-positioned and/or the gantry adjusted. To enable flexible re-positioning without the need for re-calculation of $\mathbf{T}^c$, a reference target was constructed and secured to an immovable part of the SPECT gantry. If $\mathbf{T}_r$ denotes the position and orientation of the reference target in Polaris coordinates when $\mathbf{T}^c$ was computed, and $\mathbf{T}_{r'}$ denotes the position and orientation after a relative movement between the Polaris and gantry, then the change in position of the reference target in Polaris coordinates is given by:

$$\mathbf{T}_{\Delta r} = \mathbf{T}_{r'}\mathbf{T}_r^{-1}. \qquad (25)$$

Without loss of generality we can consider any relative movement between the Polaris and gantry to be caused by movement of the Polaris. The change for the Polaris is then the inverse of (25), i.e. $(\mathbf{T}_{\Delta r})^{-1}$. The modified transformation matrix $\mathbf{T}^{c'}$ is then given by:

$$\mathbf{T}^{c'} = \mathbf{T}^c\left(\mathbf{T}_{\Delta r}\right)^{-1}. \qquad (26)$$

The reference target set-up can be seen in Figure 5.1 and the principle of using a reference target for immunity to Polaris and/or gantry adjustment is illustrated in Figure 5.2.

5.2.4 Interface and Raw Data Processing

The requested brain phantom scans were acquired on a Philips Irix triple-detector gamma camera[1] with the three detectors spaced at 120° from each other. The reconstructed pixel size was 3.5mm/pixel (no zoom). In each acquisition, 120 2D projections (40 per detector) were collected over 360°. Polaris motion data were acquired simultaneously[2].

Northern Digital Inc. provide a PC interface for the Polaris hardware. Custom software[3] was used to interface the Polaris with the SUN platform. This software was supplied to Queen Elizabeth Hospital for their own implementation of the Polaris using SUN workstations. Based on the raw data output from the Polaris, the software produces a text file listing the position and orientation data for each sample and from all targets monitored by the Polaris in a given run. For *n* monitored targets the file contains *n* lines per sample. The most important values on each line are a target identifier and six floating point values describing the x, y, and z-axis rotational components of object orientation and the x, y, and z-axis translational components of object position. In our case there were two targets: a reference target and object target. An example segment of one of the text files generated is shown in Figure 5.3. Note that all motion components are tabulated in Polaris coordinates. For each acquisition the Polaris recorded a total of 80 measurements, i.e. two samples per SPECT projection.

[1] Philips Medical Systems, Cleveland, Ohio, USA.
[2] Image and motion data were acquired by staff at the Department of Nuclear Medicine, Queen Elizabeth Hospital, Adelaide.
[3] C language utility programs written by Roger Fulton, Department of Nuclear Medicine, Royal Prince Alfred Hospital, Sydney.

Figure 5.1a. The head target. Four infrared-reflective spheres were attached to a 3mm square slab of polyethylene to form this target.

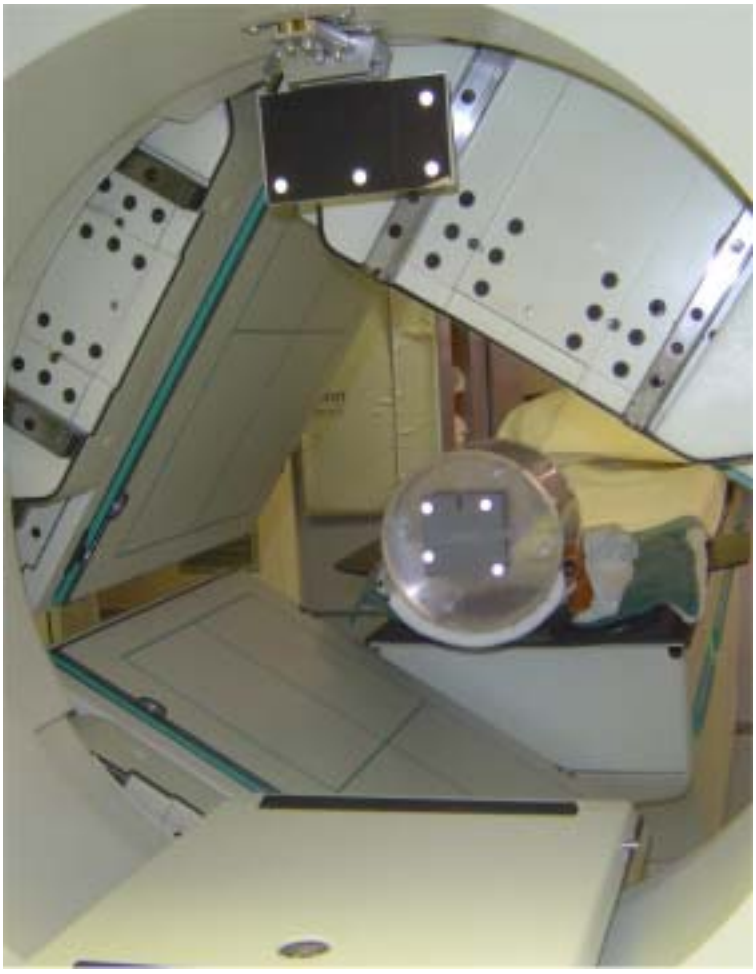

Figure 5.1b. View from the Polaris. The larger reference target fixed to the gantry and the smaller head target attached to the perspex case of the brain phantom are both in the Polaris field-of-view. The three detectors used for imaging the phantom are also seen in this photograph.

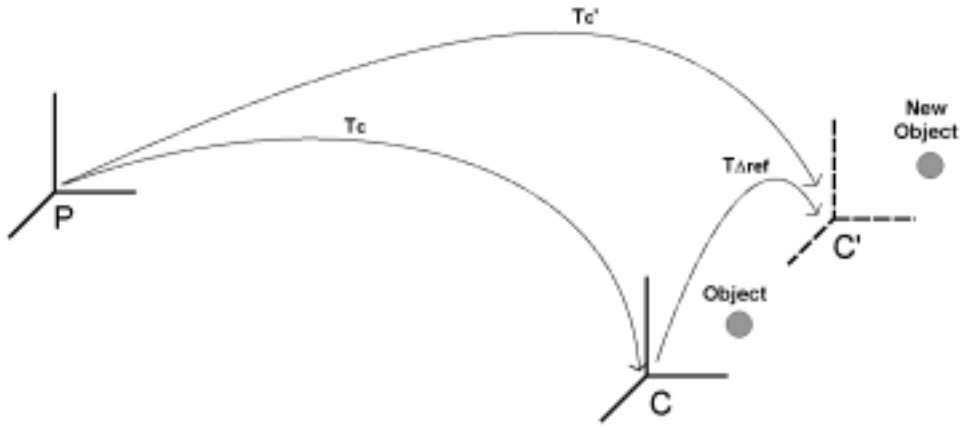

Figure 5.2. The principle of using a reference target to achieve immunity from relative motion between the gamma camera gantry and the Polaris motion tracker. Here *P* refers to the Polaris coordinate system, *C* to the original camera system, and *C'* to a translated and/or rotated camera frame. Without tracking of the relative motion between the gantry and Polaris, the Polaris would record the location of the "new object" with respect to the original *C* frame, since the matrix $\mathbf{T}^{c}$ was calibrated to transform between these two systems. Thus such relative movement will appear as head movement. However, knowledge of the shift/rotation undergone by the reference target enables us to compute the location of the "new object" with respect to the shifted/rotated frame *C'*. Using the modified matrix $\mathbf{T}^{c'}$, true head movements are distinguished.

```
 1 10:02:37 A    1.1   -2.0   -89.5  -163.7    35.3 -1461.0
 1 10:02:37 B    6.3   -1.1  -155.3   201.5    -4.1 -1938.5
 2 10:02:56 A    1.2   -1.9   -89.5  -163.7    35.3 -1461.1
 2 10:02:56 B    6.2   -1.3  -155.3   201.5    -4.1 -1938.4
 3 10:03:15 A    1.2   -1.8   -89.5  -163.7    35.3 -1461.1
 3 10:03:15 B    6.2   -1.3  -155.3   201.5    -4.1 -1938.5
      ⋮                       ⋮                      ⋮
61 10:21:38 A    1.2   -1.8   -89.5  -163.6    35.3 -1461.2
61 10:21:38 B    5.2   -1.3  -154.7   203.9    -4.6 -1938.2
62 10:21:57 A    1.1   -2.0   -89.5  -163.6    35.3 -1461.1
62 10:21:57 B  -12.6    8.6   161.0   288.6     0.6 -1966.0
63 10:22:16 A    1.2   -1.8   -89.5  -163.6    35.3 -1461.2
63 10:22:16 B  -12.3    8.4   160.8   287.9     0.6 -1966.0
```

Figure 5.3. A section of the text output produced by the Polaris motion tracker in response to sampling of the physical Hoffman phantom with fiducial marker attached. The first digit represents the sample number. For each sample number, the time of measurement is recorded and the position of the reference target (A) and phantom target (B) are recorded in yaw, pitch, roll, x, y, z format (the six floating point numbers on the right). It can be seen here that the reference target remained fixed relative to the Polaris, but the phantom (target B) underwent a change in location at sample 61-62. Altogether there were 80 samples recorded.

Communication between the Polaris and gamma camera workstation can be set up in order to facilitate synchronisation between Polaris sampling and the start and stop of a SPECT acquisition. However, since the movements in our phantom and patient studies were of significant duration and were modelled as discrete, this accuracy was not required, and a simpler approach of initiating the two systems independently at approximately the same time was employed.

We were supplied with an IDL[1] program[2] to process the Polaris data from a given run and enable the user to select Polaris motion samples between which the change in object location should be computed (Figure 5.5). The program decomposes this change from a standard transformation matrix to pitch, yaw, roll (the orientation components) and x, y, z (the position components) format, in gamma camera coordinates. The operation of the processing program is summarised as follows:

1. $\mathbf{T}^c$ (transformation matrix describing Polaris to camera coordinate conversion) read in from file
2. $\mathbf{T}_r$ (transformation matrix describing the reference target location in Polaris coordinates at the time when $\mathbf{T}^c$ was derived) read in from file
3. Reference target data from the Polaris measurement record (text) is read in to six (pitch, yaw, roll, x, y, z) 80-element (number of Polaris samples) vectors
4. Object target data from the Polaris measurement record is read in to six (pitch, yaw, roll, x, y, z) 80-element (number of Polaris samples) vectors
5. User selects two Polaris samples (graphically)
6. Transformation matrix computed for each selected sample $\Rightarrow \mathbf{T}_{s1}$ and $\mathbf{T}_{s2}$
7. Transformation matrix representing the change computed as $\mathbf{T}_{s2}(\mathbf{T}_{s1})^{-1}$
8. Result is decomposed into six pure rigid-body component motions

5.3 Coordinate System Validation

It was pointed out in §*5.1* that when motions are considered as a sequence of pure rotational and translational movements about/along the primary axes, specifying component angles and shifts (eg. in pitch, yaw, roll, x, y, z format) alone is ambiguous; the order of application must also be specified since matrix multiplication is non-commutative.

During motion estimation in the data-driven approach, the Simplex optimisation routine samples sets of six rigid-body parameters based on the current cost function value. Two modifications were required in order to achieve consistency with the processed Polaris motion estimates. The first was a reordering of

[1] Interactive Data Language, RSI Systems, Boulder, Colorado, USA.
[2] Original program written by Leighton Barnden, Department of Nuclear Medicine, Queen Elizabeth Hospital, Adelaide. This was modified for use at our own institution.

the projection data. Gamma camera settings forced the first projection collected to be posterior to the supine phantom. Direct reconstruction of these acquired data produced inverted transverse slices. We reordered the projections so that the first projection was anterior. Without this modification the data-driven routine would extract motion parameters of the same magnitude (ideally) as the Polaris, but the signs would be wrong. We modelled 3D transformations by concatenating a sequence of pure-movement matrices. The sequence was x-rotation, y-rotation, z-rotation, and lastly (3D) translation. Firstly we confirmed that the processed Polaris motion estimates reflected this order of operation. With this confirmed, the parameters were directly comparable provided the definition of axes was identical in both cases. We confirmed the orientation of axes was common, but that the definition of positive and negative was not. Therefore the second modification was to reverse three of the signs in the processed Polaris motion measurements to achieve overall axis consistency between the two measurement systems.

After applying these modifications, a simple validation experiment was carried out to demonstrate that the motion measurements from the Polaris were sensible and comparable with our own system. The phantom was filled with approximately 250 MBq of ^{99m}Tc thoroughly mixed in water, and then positioned supine on the patient bed in the "unmoved" position. A catheter was wrapped obliquely around the outer cylindrical casing of the phantom and secured with tape. It was then filled via syringe with approximately 50 MBq of activity in water. The phantom and fiducial marker were imaged in the unmoved position, then reoriented and shifted to a new, arbitrary location in the field of view and imaged a second time. Acquisitions were performed with the parallel-hole low-energy high-resolution collimators mounted and contained a total of 120 projections (40/detector), each collected into a 128^2 pixel matrix for 30 seconds. The reconstructed pixel size was 3.5mm (zoom = 1). Some sample projection images from the two acquisitions are shown in Figure 5.4. We used an independent 3D-motion registration package to register the first study (unmoved position) to the second study (moved position). The registration package is based on a Simplex-driven maximisation of the normalised mutual information cost function. We modified the registration package to incorporate the same 3D transformation routine as used in our data-driven approach. Results of this validation experiment are shown in Table 5.1. Clearly the signs were consistent and good agreement in the magnitudes suggests the Polaris was providing sensible, comparable output.

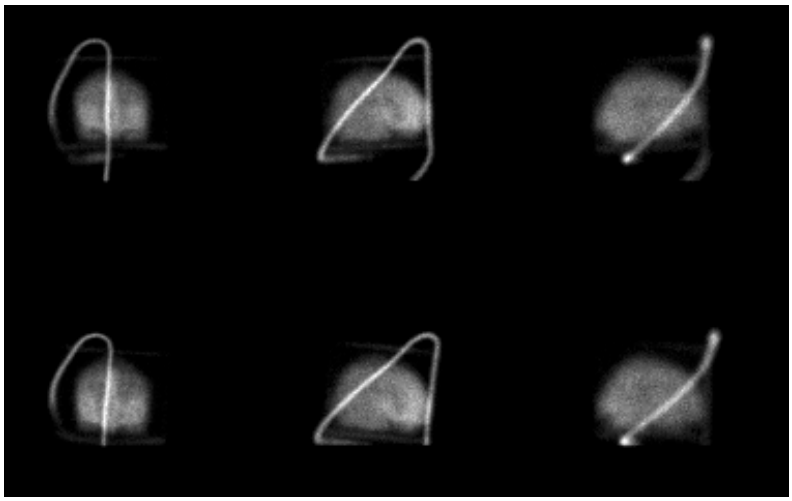

Figure 5.4. Projection images (1, 20 and 40 moving left to right) of the physical Hoffman phantom used in the coordinate system validation experiment. The top and bottom rows correspond to the unmoved and moved positions respectively. The hot fiducial marker (catheter) is clearly seen wrapped around the surface of the phantom.

Table 5.1. Motion parameters for movement of the physical Hoffman phantom with fiducial marker attached. Parameters were obtained using both the Polaris and an independent registration.

	X°	**Y°**	**Z°**	**X (pix.)**	**Y (pix.)**	**Z (pix.)**
Polaris	-7.1	-4.0	2.0	1.1	-3.6	12.7
Independent Registration	-5.5	-3.7	-2.0	1.6	-3.6	12.3

5.4 Acquisitions

Three physical Hoffman phantom studies were used in our experiments. Staff at a separate institution acquired these data upon our request[1]. The phantom was filled with approximately 200 MBq of ^{99m}Tc thoroughly mixed with water. Care was taken to remove large air bubbles. A total of 120 projection images of the phantom were collected over 360° using a 128^2 pixel matrix. This corresponded to 40 projections per detector. Each projection was collected for 30 seconds. The reconstructed pixel size was 3.5mm (zoom = 1).

[1] Department of Nuclear Medicine, Queen Elizabeth Hospital, Adelaide.

Movements of the phantom were applied manually by an operator. This was done (rapidly) between two neighbouring projections so as to minimise motion corruption within projections. Two of the studies acquired had a single 3D movement lasting for $< {}^{1}/_{4}$ of the acquisition and the third had two 3D movements each lasting for ${}^{1}/_{3}$ of the acquisition. Independent measurement of the applied motion was obtained using the Polaris motion tracking system. An example motion trace recorded by the Polaris is shown in Figure 5.5.

5.5 Motion Detection and Correction

Initially projection data were reconstructed without attenuation using 16 OSEM subsets and 8 projections per subset. Angular dispersion of these subsets was maximised. Based on the method outlined previously, motion groups were detected and a partial reconstruction formed from the largest group of consistent projections. A directed 6-parameter downhill-Simplex search was performed to match the orientation of this reconstruction to each motion group. In the case of dataset 3, the partial reconstruction was updated with projections from the first optimised group before the second optimisation was performed. No account was taken of attenuation during motion detection and correction.

5.6 Analysis

Extracted motion parameters were compared with those measured by the Polaris. Motion-corrected reconstructions obtained using our motion estimates were assessed visually and compared with the corresponding motion-corrected reconstruction derived using the Polaris parameters. To characterise the behaviour of the cost function in the neighbourhood of a search solution, relevant parameters were manually driven about the solution values whilst other parameters were fixed at their solution values.

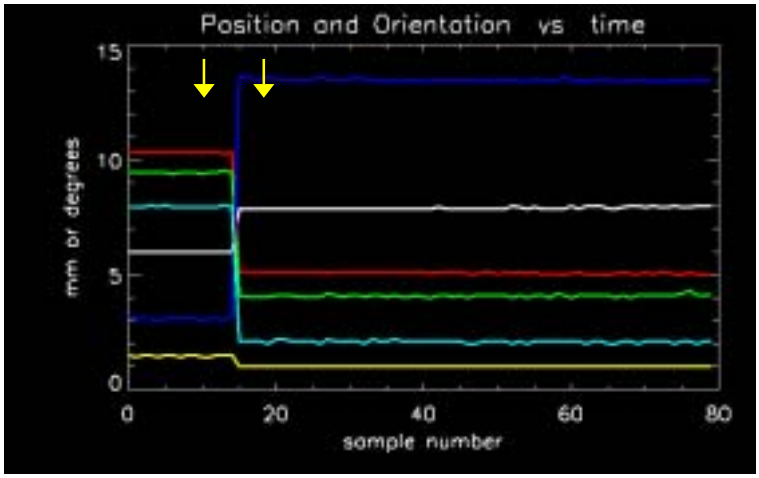

Figure 5.5. Polaris motion trace for the first physical Hoffman brain phantom dataset. The colour coding is: cyan, x-rotation; blue, y-rotation; orange, z-rotation; white, x-translation; red, y-translation; green, z-translation. This scan contained a single movement of the phantom. The Polaris data is displayed here without reordering. Note the user-defined samples (marked with yellow arrows) between which the location change is computed.

5.7 Results and Discussion

The motion parameters recorded by the Polaris are listed in Table 5.2. Shown alongside these are the corresponding parameters extracted using the data-driven approach. For datasets 1 and 2 (each containing a single corrupting movement), there was good agreement for the x and y-rotation and z-translation parameters. However, a reasonable discrepancy existed for the remaining DOF (x and y-translation and z-rotation parameters).

A surface plot of the cost function is shown in Figure 5.6 for dataset 1. The similarity measure was computed as the x-translational and z-rotational parameters were manually driven through the range –1 to 8 pixels, and –4° to 5° respectively. Note that all other parameters were fixed at the Polaris values. The topology of this map indicates there was ambiguity identifying the parameters corresponding to the solution (minimum value of C) due to a broad, slowly-varying cost function. Based on this and other cost function maps (not shown) it was observed that each rigid-body parameter had a varying influence on the cost-function. The results suggest that the MSD is fairly insensitive to large rotations whilst being sensitive to typical translations (Figure 5.6). The significance of this for motion correction is that there may be difficulty in locating a minimum in the direction of low sensitivity unless the downhill-Simplex algorithm can be made to search directly along such paths in the multi-dimensional space. Furthermore, combinations of the rigid-body parameters are correlated to a different extent. Correlation means that error in one parameter will propagate to error in other parameters. Z-translation is the only parameter that tends to act fairly independently from parameters paired with it.

Table 5.2. Data-driven and Polaris measured rotations (°) and translations (pixels) for the first (left) and second (right) physical phantom datasets. Both datasets contained single corrupting movements.

DOF	Extracted	Polaris	DOF	Extracted	Polaris
X°	-9.8	-9.1	X°	-0.6	-0.1
Y°	8.6	7.7	Y°	0.0	0.1
Z°	1.5	-0.7	Z°	-2.2	-0.2
X	1.1	2.8	X	3.9	2.4
Y	0.4	2.5	Y	3.2	4.1
Z	2.7	3.0	Z	-4.0	-3.7

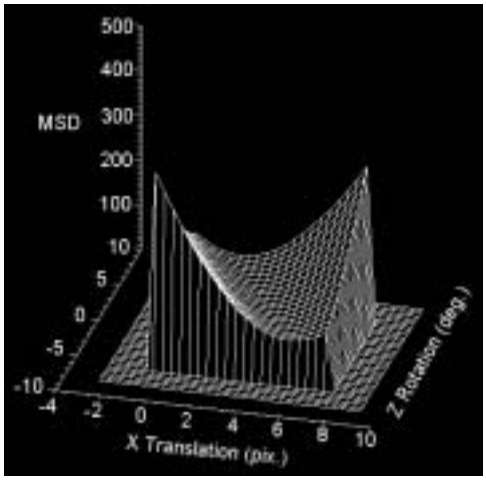

Figure. 5.6. Surface plot of the cost function near the solution extracted for the first physical phantom dataset. This shows a "valley" traversing the x-translation-z-rotation parameter space. The shallowness made locating the minimum difficult and resulted in a discrepancy between the extracted and Polaris-measured values.

Much better overall agreement between the extracted parameters and Polaris parameters existed for the two movements in dataset 3 (Table 5.3). Reconstructed slices before correction, and after correction using our extracted parameters and the Polaris parameters, are shown in Figure 5.7. Motion artifacts, including severe anterior distortion, are clearly evident in the uncorrected slice. Motion correction resulted in a significant improvement. Some asymmetry in the uptake is noticeable in the motion-corrected slices. A possible reason for this is that the unmoved phantom axis was slightly misaligned with the tomograph axis, i.e. these are not quite true transverse slices. It may also be a consequence of subset structuring, though this is yet to be investigated. It should be noted that the similarity of the corrected slices suggests that data-driven motion correction performed close to its potential in this case.

Table 5.3. Data-driven and Polaris measured rotations (°) and translations (pixels) for the third physical phantom dataset. This dataset contained two corrupting movements.

	Movement 1		Movement 2	
DOF	Extracted	Polaris	Extracted	Polaris
X°	-1.2	-1.3	-0.7	0.1
Y°	1.4	1.9	8.1	8.4
Z°	-26.8	-25.9	-1.7	0.2
X	-0.5	-0.1	-3.9	-5.7
Y	-0.2	-0.1	-0.5	0.1
Z	-0.3	-0.3	-1.4	-1.4

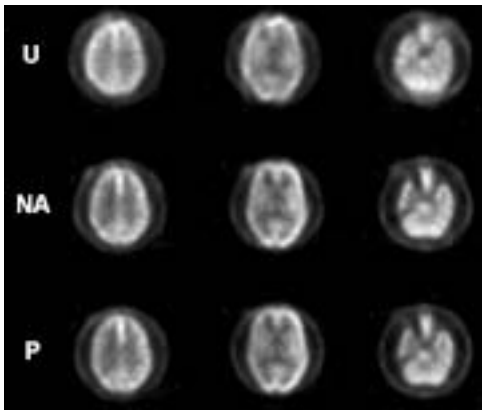

Figure 5.7. Upper, middle and lower brain slices (left to right) shown for physical phantom dataset 3 (3 head positions). The uncorrected study (U) had significant artifacts. Using the *NA* regime (NA) and Polaris measurements (P) for correction resulted in very similar slices, each with significantly reduced distortion.

We postulate that discrepancy between the extracted and Polaris motion parameter measurements for certain DOF is a consequence of object symmetry. To demonstrate this we simulated the set of projections acquired from a uniform, off-centre cylinder of activity. A single, discrete axial rotation of the cylinder was incorporated into the acquisition according to the method of §*4.1*. The cost function neighbourhood of the true solution as a function of z-rotation and y-translation is shown in Figure 5.8. Clearly there is no unique solution (distinct minimum). An off-axis, axial rotation of a cylindrically symmetric uniform activity distribution cannot be uniquely estimated using the projection images. Cylindrical symmetry leads to ambiguity between axial rotations and x/y translations, both of which can give rise to the same projection image. This principle is illustrated in Figure 5.9a. The design of the physical phantom permits activity to spread around the central circular insert and also above and below

the insert. Thus there are significant activity contributions from regions with cylindrical symmetry. These appear in the projection data as a hot rectangle of activity surrounding the brain (Figure 5.9b). It is likely that this biases the parameter estimation by confusing axial rotation with translational movement in the plane of projection.

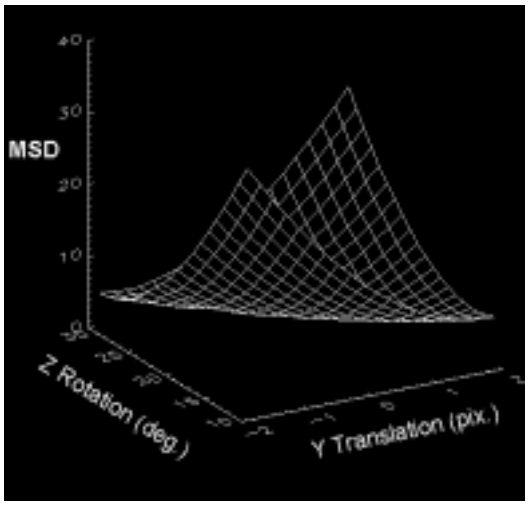

Figure 5.8. Cost function shape near the solution (pure axial rotation) obtained for an off-axis, uniform cylindrical phantom. This is shown as a function of the z-rotation and y-translation parameters.

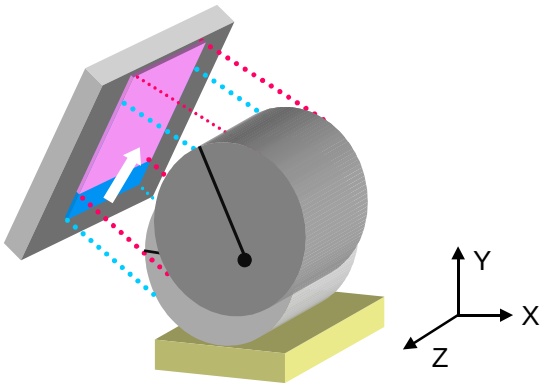

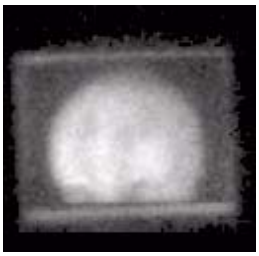

(a) (b)

Figure 5.9. (a) When projections are being used to determine 3D movement, an off-axis axial rotation of a uniform cylinder is indistinguishable from a simple translation in the plane of projection; (b) a projection of the physical Hoffman phantom showing the 'hot rectangle' around the periphery caused by activity between the brain and outer perspex case and also in compartments above and below the brain.

CHAPTER 6

CLINICAL VALIDATION

6.1 Introduction

The final section of the present work describes methods and results obtained from data-driven motion correction applied to human volunteer data. This constitutes a preliminary clinical feasibility study that it is hoped will be a solid basis for further development of the method in future work. The aims of this section were to

- Investigate the influence of various parameters on the data-driven detection and correction of motion in human subjects
- Highlight important practical considerations of applying our method to human subject data
- Propose a novel data-driven working procedure for obtaining improved clinical outcome in human studies affected by motion

6.2 Methods

6.2.1 Acquisition of Clinical Data

We received controlled brain SPECT studies of three normal volunteers. Staff at a separate institution[1] acquired these data in conjunction with their own development of a research database of brain SPECT normals.

Volunteers were injected with 300-400 MBq of ^{99m}Tc bound to HMPAO[2]. This dose was within NHMRC[3] specifications for volunteer subjects. This dose is less than the standard dose of ≥ 600 MBq administered to patients undergoing a brain SPECT procedure. The Polaris motion tracker was set up to collect an independent account of patient motion over the course of the studies. It was operated in passive mode at a sampling rate of 2/projection (~0.05 Hz). The infrared reflecting target was the same as that used in the physical phantom studies described in §*5.2.2*. A standard welding mask with the visor removed was fitted to the patient and used as a base for attaching the target. Three semi-rigid plastic straps from the mask wrapped across the forehead, across the crown of the head, and under the back of the head respectively. The target was secured firmly to the modified mask at the crown of the head as shown in Figure 6.1.

[1] Department of Nuclear Medicine, Queen Elizabeth Hospital, Adelaide.
[2] Hexa-methyl phenyl amine oxime
[3] National Health and Medical Research Committee

Figure 6.1. Modified welding mask used to attach the head target (grey slab with spherical reflectors) to the patient head. There are three semi-rigid plastic straps: one passing across the forehead, one across the back of the head, and one across the crown of the head. The target is secured to the top of the crown strap.

A Philips Irix triple-detector gamma camera mounted with parallel-hole low-energy high-resolution collimators was used for the acquisitions. For each study, 120 projections were collected over 360° (40 projections/detector, 38s/projection) giving a total acquisition time of approximately 25 minutes. The projection matrix size was 128^2 pixels. A photopeak image (centre 140 keV, width 20%) and scatter image (centre 118 keV, width 10%) were collected in each acquisition.

Each volunteer consented to undergo two scans. They were instructed to remain as still as possible for the duration of the first scan. This constituted our reference (motion-free) study. After a short break the second study was performed in which the volunteer was instructed to remain still until prompted by staff to alter their head position and orientation. This prompt was timed so that movement was substantially completed between projections. The resulting acquisition constituted our movement (motion-corrupted) study. Note that only single-movement[1] studies were collected and that in each case the two head locations involved were both well represented in terms of the total acquisition time.

[1] Although a single step-wise movement was requested, other non-voluntary movements could have occurred. Therefore the Polaris trace was checked to ensure that this was a valid assumption. In all cases volunteers remained in a roughly constant location either side of the movement (Figure 6.10). Minor drifting was ignored; such movement, however, was minimal and unlikely to be the cause of significant artifacts visible over and above noise.

6.2.2 Preparatory Data Processing

Three preparatory processing steps were performed to generate input data for the motion detection and estimation stages. Firstly, all projection data were reordered as described in §*5.3* so that the first projection was anterior. This ensured the reconstructed slices would be correctly oriented. Secondly, the reordered movement study for each volunteer was corrected for scatter. The image collected using a 10% window centred at 118 keV served as a reasonable approximation of the scatter distribution in the corresponding photopeak image. We estimated scatter fraction using a simple linear drop-off approximation. The scale factor to multiply the scatter image was computed as

$$SF = \frac{1}{2}\left(\frac{W_{Photo}}{W_{Scatter}}\right), \qquad (27)$$

where W_{Photo} and $W_{scatter}$ are the % widths of the photopeak and scatter windows respectively. The principle is illustrated in Figure 6.2. In our case *SF* was unity, allowing us to use the scatter images unmodified. Before subtracting the scatter image from the corresponding photopeak image, it was filtered using a 2D Butterworth kernel (order = 5, cut-off = 0.3 x Nyquist) in order to reduce noise. After subtraction, any negative counts were zeroed. In our optimisations, both scatter-corrected and non scatter-corrected data were used to estimate motion.

The projected brain occupied less than half of the total available projection area in the 128^2 images. Therefore it was reasonable to extract, from each projection, a reduced projection containing the whole brain but free of much of the zero-space in the original image. Based on the cinematic viewing of projections, a reduced matrix of 80^2 pixels was chosen in order to include the whole brain in both the unmoved and moved locations across all projections. In the transaxial dimension the matrix was extracted symmetrically to avoid introducing artificial centre-of-rotation errors. Axially, asymmetric extractions are permissible provided the same region is extracted in all projections. In our case the 80 axial rows were chosen to exclude much of the zero-space above the top of the head as well as a buffer[1] of zero-space at the base of the image. The matrix extraction step was performed on the reordered movement scan for each volunteer. It enabled us to determine the effect that increasing the proportion of data-of-interest would have on motion estimation[2].

[1] This buffer was a zero-padded region created by the acquisition software to enable square images to be generated from non-square FOV detectors.

[2] Note that an added advantage of reducing matrix size is shorter processing times. Since the rate-limiting step of our Simplex optimisation procedure was the 3D transformation of the reconstruction, reducing the matrix size from 128^3 to 80^3 reduced optimisation time by approximately a factor of 4.

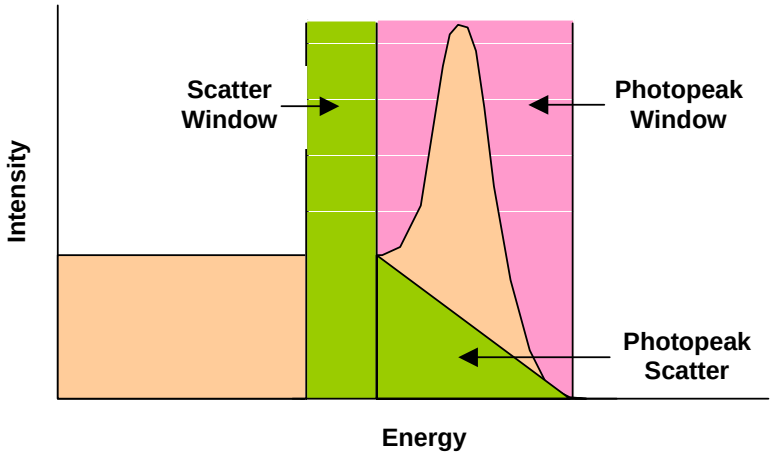

Figure 6.2. Simple scatter subtraction. A window just downstream of the ^{99m}Tc photopeak is collected to provide an estimate of the distribution of scatter in the photopeak. The scatter profile in the photopeak is assumed to drop away linearly and the scatter image is scaled by a factor of half the ratio of window widths to account for this.

6.2.3 Semi-Simulated Studies

In the three volunteer studies staff did not attempt to ensure that the starting location of the head in the movement acquisition coincided with the location held during the reference acquisition. Indeed the break between acquisitions, during which the head target was removed and volunteers left the bed to return some minutes later, meant this could not be assumed. Although the reference target provides immunity from relative movement between the gantry and Polaris (even between scans), separate scans become disconnected if there is re-positioning of the volunteer relative to the bed, or of the head target relative to the head[1]. To compare one of our motion-corrected reconstructions with its corresponding reference the two needed to be aligned. Since the misalignment between the two studies was not directly measured, an independent 3D volume registration post-correction was required to align them[2].

We wished to quantify, using a single figure, the improvement obtained from motion correction, and then to compare this with the performance of motion correction based on the Polaris measurements. One robust method for doing this involves constructing new studies in which the reference and motion-corrected data share a common brain location. This circumvents the need for independent alignment of the studies and the associated registration errors. The semi-simulation procedure can be summarised as follows:

[1] Note that relative movement between the head target and head is problematic even within the same scan since head movement is de-coupled from the target. This scenario is not unlikely and is discussed in greater detail later in this chapter.
[2] The 3D independent registration is described later in this chapter.

(i) Reference and movement data (each consisting of photopeak and scatter images) were reordered.

(ii) Movement data (photopeak and scatter) were corrected for radionuclide decay. The decay factor was calculated based on the time of acquisition recorded in the data header. This correction normalised counts in the movement data to counts in the reference data.

(iii) Reference and (decay-corrected) movement data were scatter-corrected (see §*6.2.2*).

(iv) Projections from the movement scan were then transplanted into the reference scan to simulate a discrete motion. Provided the location of the brain in the movement scan differed from the location in the reference scan, any such transplanted projection constitutes a 'movement'.

There are a number of advantages of generating semi-simulated studies. Firstly, we have a complete reference study and the known movement to align it with the motion-corrected study. This movement is simply the inverse of the Simplex solution. Secondly, there is no reliance on the accuracy of the Polaris measurements. And thirdly, there is no reliance on an independent 3D registration (with its associated errors) to enable a direct quantitative comparison.

Therefore, in addition to the initial three volunteer studies, we requested a fourth volunteer study in order to implement the semi-simulation technique. Similarly to the earlier volunteers, the fourth volunteer underwent two brain scans. However, unlike the earlier scans, this volunteer held a single head position for the full duration of each scan and changed position between scans. Photopeak and scatter images were collected as before. From these data we derived two motion-corrupted semi-simulated data sets. These represented single-movement[1] scans where the second position was held for $^1/_2$ and $^1/_8$ of the total scan respectively.

Using the modified welding mask, operators reported observing a 'catching' of the posterior plastic strap (around the back of the head) with the patient bed as the volunteers moved. To minimise this scenario and the associated risk of the head and target movement becoming decoupled, the semi-rigid strap was removed from the mask and replaced by a soft fabric strap that fastened across the back of the head with Velcro® (Figure 6.3). Also, unlike the previous three volunteers, we requested that the fourth subject not have the head target removed between scans. This prevented relative movement between the target and head between scans. It also enabled us to transplant part of the Polaris output (eg. Figure 5.3) from the first scan directly into the output from the second scan, and use this hybridised text file to compute the movement that the Polaris would have measured had the movement occurred within a single scan. This is only permissible if the fixation of the target to the head is identical in the two studies. Moreover, we requested that the two scans be performed back-to-back to ensure that patient position relative to the bed was unchanged between the end of the first scan and start of the second.

[1] See footnote 1 on p. 66.

Figure 6.3. Head target attachment mechanism for the fourth volunteer. The semi-rigid plastic strap that formerly passed across the back of the head was cut away and replaced with a soft strap fastened using Velcro®. This modification was aimed at minimising 'catching' of the device on the patient bed and thereby minimising the decoupling of target and head motion.

6.2.4 Motion Detection

Motion groups were identified according to the procedure described in §*3.1.2* and §*3.2.4*.

6.2.5 Simplex Optimisations (Motion Estimation)

Downhill-Simplex optimisations (directed parameter searches) were used to estimate motion in the human subject studies. These searches were continued for a minimum of 300 iterations and a maximum sufficient to obtain relative stability in the estimates for at least 50 iterations. In practice this iteration value turned out to be between 300 and 450. It should be noted that stability in the parameters was often reached well before the minimum of 300 iterations. The simplex *scale* arguments were fixed for each particular volunteer to be approximately 10° for rotational DOF and approximately 5 pixels for translational DOF[1]. The *tolerance* argument was set to a value of 10^{-5} for all optimisations for all volunteers. Although the preceding arguments were not necessarily the same in each run, they were not the variables of interest. Rather, they were set to produce sensible, unchanging output. The variables of interest, i.e. those systematically changed from optimisation to optimisation, are described below.

[1] Initial *scale* specification is not critical in most cases as the simplex can expand and contract in the multi-dimensional space.

Projection Size

Optimisations were run with the original, reordered projection data (size 128^2) as well as reduced projection data (size 80^2) (described in §*6.2.2*). The partial reconstruction size was 128^3 and 80^3 for these data respectively.

Thresholding

We investigated the effect of thresholding the acquired projections. All pixels less than the threshold value were zeroed to exclude low-count background from contributing to, and potentially biasing, the cost function. The thresholds applied were 10% and 20% of the maximum count in a projection. Note that thresholding had the additional effect of sharpening projection edges. Thresholding was not applied to the reprojections since the smoothing effect of reprojection already averaged away any low-count background to zero.

Smoothing

The influence of smoothing the acquired projections was investigated. Projections were convolved with a 2D, 2-pixel (7mm) FWHM, normalised Gaussian kernel prior to being compared with the corresponding reprojections. Since reprojection itself is a smoothing operation, reprojections were not further smoothed.

Scatter Correction

We investigated the effect of scatter correction on the estimation of motion. Scatter correction of projection data was described in §*6.2.2*. Once scatter-corrected, the original or reduced-size projection data were reconstructed to give a scatter-corrected partial reconstruction, the location of which was then optimised.

Note that the thresholding, smoothing and scatter correction modifications were applied for the optimisation stage only. Once the single motions were estimated, motion correction was performed using the original projection data, ie. non-thresholded, unsmoothed, and non scatter-corrected data.

Truncation

In our motion correction approach, motion-corrupted projection data are grouped into projection sets according to identified movements. Reconstructing any such set produces a ‘partial’ estimate. Although the cost function calculation is performed on projections, motion correction essentially represents optimal 3D alignment of one partial reconstruction to another. Transforming a partial reconstruction as part of the directed search for motion parameters can truncate part of the activity distribution. This in turn can lead to a large, but artificial, MSD between projections. The result is a biasing of the search toward a solution that optimally compensates for non-overlapping activity between the images caused by truncation. How influential this truncation effect is depends upon the actual head movement and the proximity of the brain to the FOV edge. The latter is exacerbated by the commonplace practice of positioning the camera above

the patient shoulders in order to minimise the radius of rotation. Thus the brain often sits close to the bottom of the FOV in the projections, with facial activity extending to the bottom edge. On account of this, any translation up the patient bed, or significant rotation about the x-axis, generates a truncated image. An example of the truncation effect for one of the volunteer studies is illustrated in Figure 6.4a. It can be seen that optimisation of the reconstruction has proceeded to a point where registration of the projection images is being strongly biased by non-overlapping facial activity.

Methods to Avoid Truncation Effects: Fixed Mask and Roving Mask

To minimise this effect it is necessary to select a reduced area on the projections over which the cost function is calculated. This area should identify a volume of activity common to the projections being compared. A simple implementation excludes everything below a fixed horizontal cut-off line. We implemented this approach, approximating the level of the line from the cine. The cost function was modified to compute the MSD only over pixels above the line. Figures 6.5 illustrates the principle and Figure 6.4b shows how this was applied in a volunteer study to exclude non-overlapping facial counts from the MSD calculation.

Clearly this simple method is limited by actual brain movement: severe rotations preclude a fixed horizontal line being used to exclude biasing regions of the activity distribution whilst at the same time maximising inclusion of brain tissue.

We developed a more general approach using a roving mask. This involved tracking the 3D orientation of the plane of truncation and restricting the cost function calculation to boundaries determined by the projection of this plane. The main steps of this algorithm were:

(i) Define the centre of transformation
(ii) Define a point on the plane
(iii) Define the direction of the normal to the plane
(iv) Define the level (vertical pixel index), in the projection space, of the projection of this plane
(v) Combine the transformation matrices for the 3D reorientation/translation (Simplex) and the 2D (axial) rotation (to align the desired projection direction with the positive vertical axis). Note that application of this matrix enables the desired projection to be obtained by summing up columns
(vi) Compute the location of the transformed initial point
(vii) Compute the (normalised) normal to the transformed plane
(viii) Compute the intersection of the transformed plane with the image volume faces
(ix) Compute the projection of the upper-most boundary of this plane
(x) Form the 2D mask containing all points above this projected boundary. If the computed boundary line is below the initial level, use the initial level as the boundary

Operations (v)-(x) are executed for each reprojection generated. In contrast to the fixed mask, the roving mask is defined by a variable cut-off line that is dependent on both the actual brain movement and the angular position of the projection required. Figure 6.6 illustrates the principle and Figure 6.4c shows an example of its application in a volunteer study.

Various other authors in the literature have described what is essentially a tomographic data truncation problem. Klein et al (2001) used the COM of list mode data to measure the respiratory cycle. They describe two masking techniques to selectively include cardiac events and thereby avoid biasing due to liver activity moving in and out of the FOV. The first excludes a fixed segment of data that does not contain the object of interest (heart). The second involves drawing a volume of interest on the reconstructed data and projecting this volume into sinogram space. These two methods are analogous to the two masking methods used in our work. Thielemans and Mustafovic (2003) describe normalising for sensitivity changes in list mode PET due to activity contributing to the FOV for only part of the time.

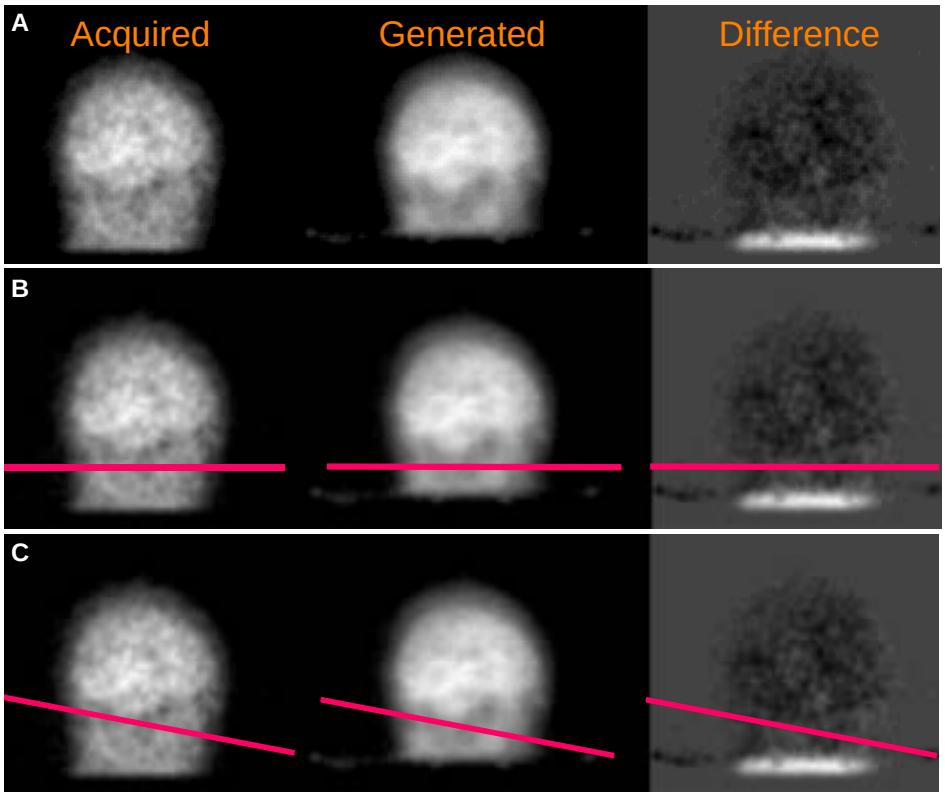

Figure 6.4. (a) Truncation occurring during the optimisation of the second volunteer study. Translation of the head upward (axially) has resulted in truncation of facial activity in the reprojection and caused an artificial mismatch between the acquired and reprojected images. The mismatch is indicated by the difference image (right column); (b) the fixed mask cut out everything below a horizontal line in the projection space; (c) the roving mask cut out everything below a (varying) sloped line in the projection space.

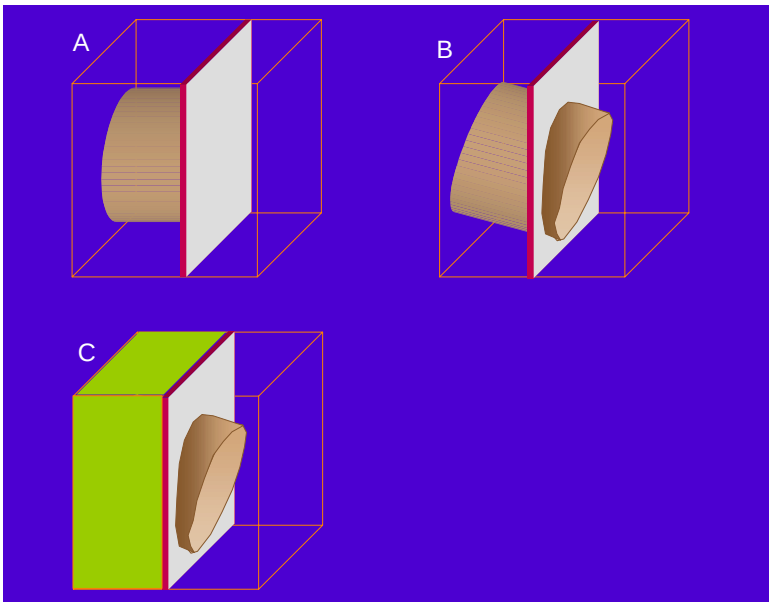

Figure 6.5. Principle of the fixed mask. (a) The plane (axial slice) corresponding to the base of the useful image; (b) a plane is estimated from the cine such that the object is never truncated superior to this plane; (c) the region superior to this plane is projected and the MSD calculated over the projected area.

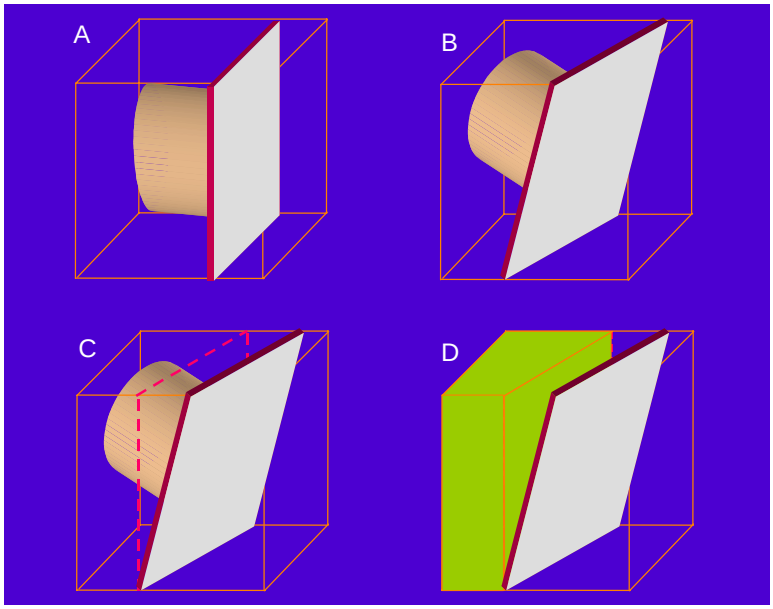

Figure 6.6. Principle of the roving mask. (a) The plane (axial slice) corresponding to the base of the useful image is defined; (b) this plane is reoriented according to the transformation applied to the object; (c) the most superior points of intersection are projected as the boundary of the required region; (d) the region superior to this boundary is projected and the MSD calculated over the projected area.

6.2.6 Motion Correction

Changing Coordinate Systems

If we denote the initial and final location of any particular brain movement by $\mathbf{T}_1{}^c$ and $\mathbf{T}_2{}^c$ respectively, where the superscript indicates measurement with respect to the original camera coordinate system, *C*, the change in location is given by

$$\mathbf{T}_{12}^{c} = \mathbf{T}_2^{c}\left[\mathbf{T}_1^{c}\right]^{-1} . \tag{28}$$

For reconstructions generated from non-reduced projection data, $\mathbf{T}_{12}{}^c$ is obtained directly from the downhill-Simplex optimisation procedure. However, if reduced square projections are extracted from the original projections as per §*6.2*, and if this extraction is asymmetric in the axial direction, the reconstructed matrix centre (camera frame origin) shifts axially. In this case the change we measure is given by

$$\mathbf{T}_{12}^{c'} = \mathbf{T}_2^{c'}\left[\mathbf{T}_1^{c'}\right]^{-1} , \tag{29}$$

where $\mathbf{T}_1{}^{c'}$ and $\mathbf{T}_2{}^{c'}$ denote the initial and final brain positions as measured with respect to the new shifted system, *C′*. Letting $\mathbf{T}^{c'\text{-}c}$ denote the matrix operator to convert from the new camera system to the original system, $\mathbf{T}_1{}^c$ and $\mathbf{T}_2{}^c$ may be expressed with respect to the new system as

$$\begin{aligned} \mathbf{T}_1^{c} &= \mathbf{T}^{c'-c}\mathbf{T}_1^{c'} \\ \mathbf{T}_2^{c} &= \mathbf{T}^{c'-c}\mathbf{T}_2^{c'} , \end{aligned} \tag{30}$$

and substituting into (29) we obtain:

$$\mathbf{T}_{12}^{c} = \mathbf{T}^{c'-c}\mathbf{T}_{12}^{c'}\left[\mathbf{T}^{c'-c}\right]^{-1} . \tag{31}$$

Thus motion estimates in the new frame can be converted to estimates in the original frame provided $\mathbf{T}^{c'\text{-}c}$ is known. The shift magnitude is simply the difference between the y-index of the extracted matrix centre and the y-index of the non-reduced matrix centre (63.5). Directionally the shift is positive (+z) since it corresponds to translation down the patient bed (Figure 6.7). Therefore $\mathbf{T}^{c'\text{-}c}$ is given by

$$\mathbf{T}^{c'\text{-}c} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & \mathit{shift} \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{32}$$

Motion estimates obtained from reduced projections were converted using (31). The new estimates were then directly comparable with the processed Polaris output (also a measurement of $\mathbf{T}_{12}^{C}$).

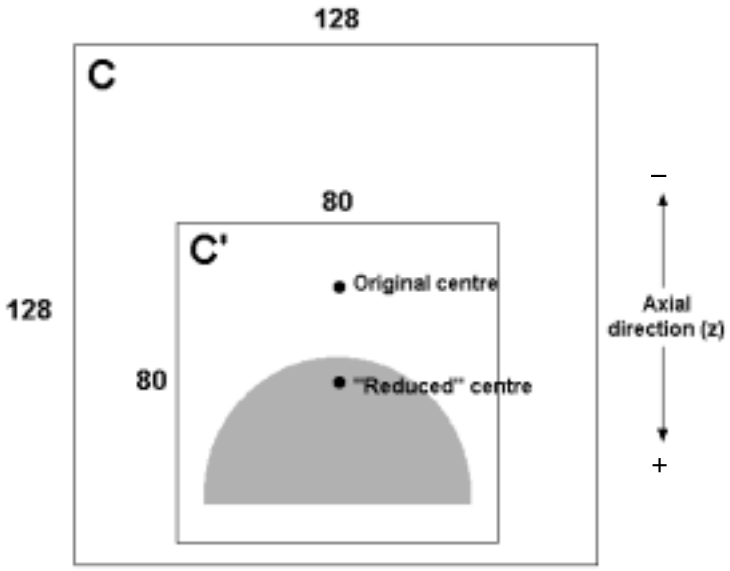

Figure 6.7. Shift in coordinate system due to the reduction of the projection matrix size. The *C'*-*C* transformation is a positive shift since the relative brain location in the extracted projections is more centralised compared to the original images where the brain sits near the bottom of the FOV.

Reconstruction

Another advantage of the conversion just described is that motion-corrected reconstructions could always be generated from the original (non-reduced) projection data. These acquired projections contained excess space around the brain and thereby minimised edge effects during motion correction.

During OSEM reconstruction, any voxels that get zeroed remain fixed at this value even if non-zero data are subsequently back-projected through these voxels. This is a consequence of the multiplicative iterative updating step. To compensate for the effect and preserve total counts, counts that would have filled a zeroed voxel are shunted to the edges of the voxel (or voxel neighbourhood), thus creating artificially high-count regions. Typically these high-count regions are produced outside the tissue of interest (brain and skull) and are therefore inconsequential to diagnosis and easily avoided in any quantitative measurement. Our motion correction approach involves transformation of the partial reconstruction prior to reconstruction updating. If these transformations are based on nearly exact estimates of the motion, the transformed reconstruction will in theory always serves as a good starting object. However, if there is a moderate error in the motion estimates, the transformed partial reconstruction will be slightly misaligned with respect to the projection data to be added. The updating

data can then be back-projected into zeroed regions and induce truncation or high-count artifacts. In our previous experience we have observed truncation of brain and skull as well as high-count edge effects close to the tissue of interest. To avoid this, once the current estimate was transformed, circular masks were created to surround the brain in each slice, and a pixel-wise logical operation was executed to impose a value of unity in the brain wherever a zero existed within the region defined by the mask. This ensured that the projection data used in an update would contribute all relevant counts to the new estimate. Although the procedure added counts to the starting object, these counts were subsequently removed from the brain during the update if they were inconsistent with the measured data.

Motion Correction

Once motion parameters had been estimated (and converted to the original coordinate system if necessary), they were applied in motion correction using the masking technique above. No attenuation correction was included in the estimation of motion or in the formation of the final motion-corrected result. All volunteer and semi-simulated studies had single, discrete corrupting movements and therefore required a single OSEM update for correction. Once corrected, the inverse movement was applied to return the corrected reconstructions back to their starting locations, i.e. the location of the initial partial reconstruction.

6.2.7 Pre-Analysis

This section details important results and pre-processing steps necessary before a thorough quantification of motion correction could be made. A description of the actual analyses performed is provided in the next section.

Equivalent Smoothing

It was necessary to determine the 3D Gaussian kernel (FWHM), which, when convolved with a noisy image, would reduce noise by the same extent as the interpolative smooth incurred in one application of a typical 3D transformation using our transformation routine. We generated a uniform, digital 'block' phantom in a 64^3 matrix and formed the set of projections of this phantom using an analytical projector. High-resolution parallel-hole collimation and depth-dependent resolution were modelled in the projection process. Attenuation was modelled by incorporating a uniform attenuation map (μ = 0.12). Noise was added to the attenuated projections at a level similar to the volunteer studies. The projections were then reconstructed with attenuation correction, and various arbitrary 3D transformations were applied to the reconstructed block using our transformation routine. For each transformation the noise was calculated in a uniform 20^3 region (8000 voxels) by dividing the standard deviation of counts in the region by the mean count in the region. The average of these measurements was 0.09 and was taken to be the noise level after a single interpolative smooth induced by 3D transformation. For comparison, we performed 3D convolution of the reconstructed block estimate using different FWHM of the Gaussian kernel. In each case the noise level was computed post-smooth. Based on these results (Figure 6.8) the interpolative

smooth from 3D transformation was shown to be equivalent to 3D[1] convolution with a 1.1-pixel FWHM Gaussian kernel.

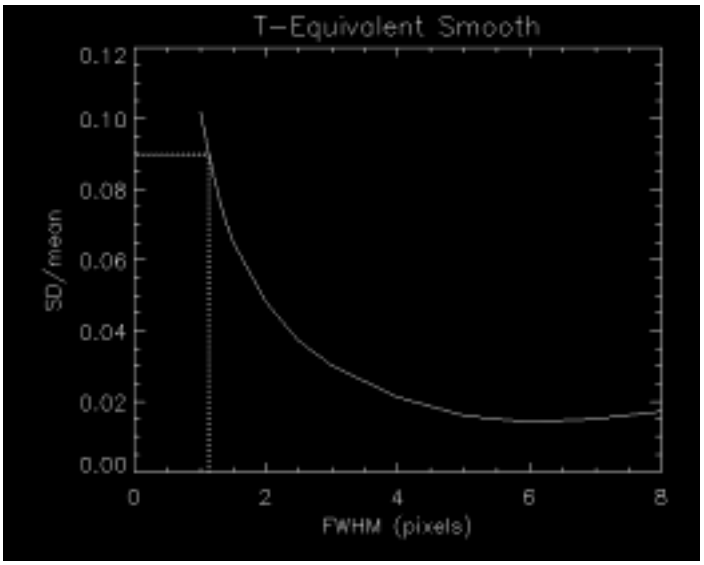

Figure 6.8. Noise level after applying 2D convolution with a Gaussian kernel of different FWHM. The kernel FWHM that resulted in a degree of smoothing equivalent to the interpolative smoothing from one 3D transformation was 1.1 pixels.

Motion-Corrected Reconstructions

The motion-corrected reconstructions generated as per §*6.2.6* were 128^3 in size and positioned in the location of their respective (initial) partial reconstructions. In all cases this was the location consistent with the majority of projections. Each 128^3 reconstruction was prepared for visual comparison with other motion-corrected studies by 3D Gaussian filtering using a 2-pixel FWHM kernel. An 80^3 block of voxels containing the whole brain was extracted from the unsmoothed reconstruction for use in an independent registration to determine the alignment between the reference and corrected reconstructions (described below). In the case of the semi-simulated studies, the same 80^3 block was extracted from the smoothed reconstruction for use in the computation of a mean square difference improvement ratio (MSDR). Note that the smoothing kernel was chosen to significantly reduce noise in the reconstructed image and thereby prevent biasing of the MSDR values.

[1] A 2D symmetric Gaussian kernel was applied in-plane followed by a 1D kernel across-plane.

Uncorrected Reconstructions

The reordered movement study for each volunteer was reconstructed without attenuation correction to produce a 128^3 voxel estimate. This was 3D Gaussian filtered using a 2.28-pixel[1] FWHM kernel in order to achieve a smoothness equivalent to the combination of one interpolative smooth (FWHM = 1.1-pixels, incurred in returning the motion-corrected reconstruction to its starting location), plus one post-correction smooth (FWHM = 2-pixels, see above). The filtered reconstruction was used for visual assessment. In the case of the semi-simulated studies, an 80^3 block of voxels (the same block as extracted from the corrected reconstructions) was extracted from the smoothed uncorrected reconstructions for use in the computation of MSDR values.

Reference Reconstructions

The reordered reference study for each volunteer was reconstructed without attenuation correction to produce a 128^3 voxel estimate. This was 3D Gaussian filtered using a 2-pixel FWHM kernel. An 80^3 block of voxels containing the whole brain was extracted from the unsmoothed reference reconstruction and input into an independent registration to determine the alignment between the reference and corrected reconstructions (described below). The output parameters of this registration were adjusted for the change in coordinates caused by the reduction (§*6.2.6*) and applied to the 128^3 smoothed reference reconstruction to align it with the smoothed corrected reconstruction. Note that after this alignment the net smoothness of the corrected, uncorrected and reference (128^3) reconstructions was equivalent. Finally, in the case of the semi-simulated studies, the same 80^3 block was extracted from the smooth, aligned reference reconstruction and normalised (decay adjusted) to the same total counts as the corrected study. It was then ready for the MSDR computation.

Independent Registration

For the first three volunteer studies, an independent registration of the motion-corrected and reference reconstructions was performed to align the two studies and hence enable a visual comparison of transverse slices. The registration was done using an in-house 3D medical image registration package based on the maximisation of normalised mutual information. We modified the package to use our own transformation routine and hence output the registration parameters according to the same order of operation as the downhill-Simplex routine.

As described above, an 80^3 block of voxels containing the whole brain was extracted from both the unsmoothed motion-corrected reconstruction and the unsmoothed reference reconstruction. We loaded the reduced motion-corrected volume as the target image and the reduced reference volume as the floating image.

[1] The net FWHM was calculated using the formula: $FWHM_{net} = \sqrt{(FWHM_1^2 + FWHM_2^2)}$.

No independent registrations were required for the semi-simulated clinical studies since in this case the reference and movement projection data shared a common location. The motion-corrected reconstruction could be transformed back to exactly the same location as the reference reconstruction for the visual comparison.

6.2.8 Analysis

Qualitative Analyses

We assessed the graphical Polaris motion traces visually for each volunteer. For the reference studies the assessment criterion was high stability over the entire acquisition – that is, minimal drift in all DOF. For the movement studies the criterion was presence of a single, sharp, discrete shift involving multiple DOF, with high stability in all DOF on either side of this shift.

Visual assessment was also made of the motion-corrected, motion-corrupted, reference, and Polaris-corrected reconstructions to compare the extent of motion artifacts before and after correction.

Comparison of Parameters

All motion estimates derived from our downhill-Simplex optimisations were compared with the corresponding motion measurement recorded by the Polaris. Note that estimates derived from optimisations involving reduced projection data were converted to the original camera coordinate system before this comparison was made.

MSDR Computation

The semi-simulated studies were used to generate a quantitative figure of improvement from motion correction. As described in preceding paragraphs, an 80^3 block containing the whole brain was extracted from the motion-corrected, uncorrected, and reference reconstructions. These volumes had approximately the same smoothness, equivalent to one interpolative smooth from 3D transformation and one direct smooth. Moreover, since the alignment parameters were known exactly, the reference and corrected reconstructions were perfectly registered. The overall improvement in motion artifacts post-correction was quantified by measuring the MSD of both the uncorrected and corrected reconstructions with respect to the reference reconstruction. Formally, this mean square difference improvement ratio, MSDR, was calculated according to (12) (§*4.5*):

$$MSDR = \frac{C\left(\hat{g}_r\left(\mathbf{s}_M\right), \hat{g}_u\left(\mathbf{s}_M\right)\right)}{C\left(\hat{g}_r\left(\mathbf{s}_M\right), \hat{g}_c^{(M)}\left(\mathbf{s}_M\right)\right)}, \qquad (12)$$

where $\hat{g}_u$, $\hat{g}_r$, and $\hat{g}_c$ denote the uncorrected, reference and corrected reconstructions respectively, and other terms are defined in §*3.1*. The MSDR was computed over central brain slices.

A flow chart of the methods in overview (§*6.2.1*-§*6.2.8*) is shown in Figure 6.9.

6.3 Results and Discussion

Recorded Polaris motion traces for the initial three volunteers are shown in Figure 6.10. The reference studies (left column) reflect a high degree of stability in all but one DOF, the y-translation (red). This means that for volunteers 2 and 3 in particular, there was a slow, vertical head drift of a few pixels over the course of their scans. Despite the fluctuation of this parameter, it did not result in obvious artifacts in the reference reconstructions (Figure 6.14). Traces from the movement scans (right-hand column of Figure 6.10) reflect remarkably stable head positions either side of sharp, discrete movements. These plots indicate that our assumption of single movements for these studies was valid. Fluctuations in the reference scan must be considered with respect to the movement scan, and a straightforward comparison of magnitudes in these plots indicates that the reference scan movements were negligible compared to the requested movement. Traces measured for volunteer 4 and the two semi-simulated traces derived from these data are shown in Figure 6.11. Again, the overall stability of the parameters was excellent. Comparing the reference fluctuation and actual (simulated) movement is simpler and more direct in the semi-simulated traces since they contain part of the reference data and part of the motion data. The glitch in x-translation (white) near sample 60 is clearly seen to be small in comparison to the movement.

In Figure 6.12 we present the Polaris data of two minimally constrained (elastic strap across forehead) patients over the course of a 30-minute SPECT brain scan (Pandos et al 2002). Figure 6.12a serves as a counter to the objection that movements in our clinical data represent an idealistic scenario. By the same token however, Figure 6.12b indicates that drifting motion also occurs during real studies, a pattern of motion we have not sought to correct in this validation. As described in §*2B.1.1*, patient motion is often characterised by both of these patterns. We have focused on sharp, discrete movements because they are better suited to our data-driven approach. However, the approach does have the potential to correct drifting motion provided such motion is approximated as a series of step-wise discrete movements (eg. dataset 7 in §*4.2*).

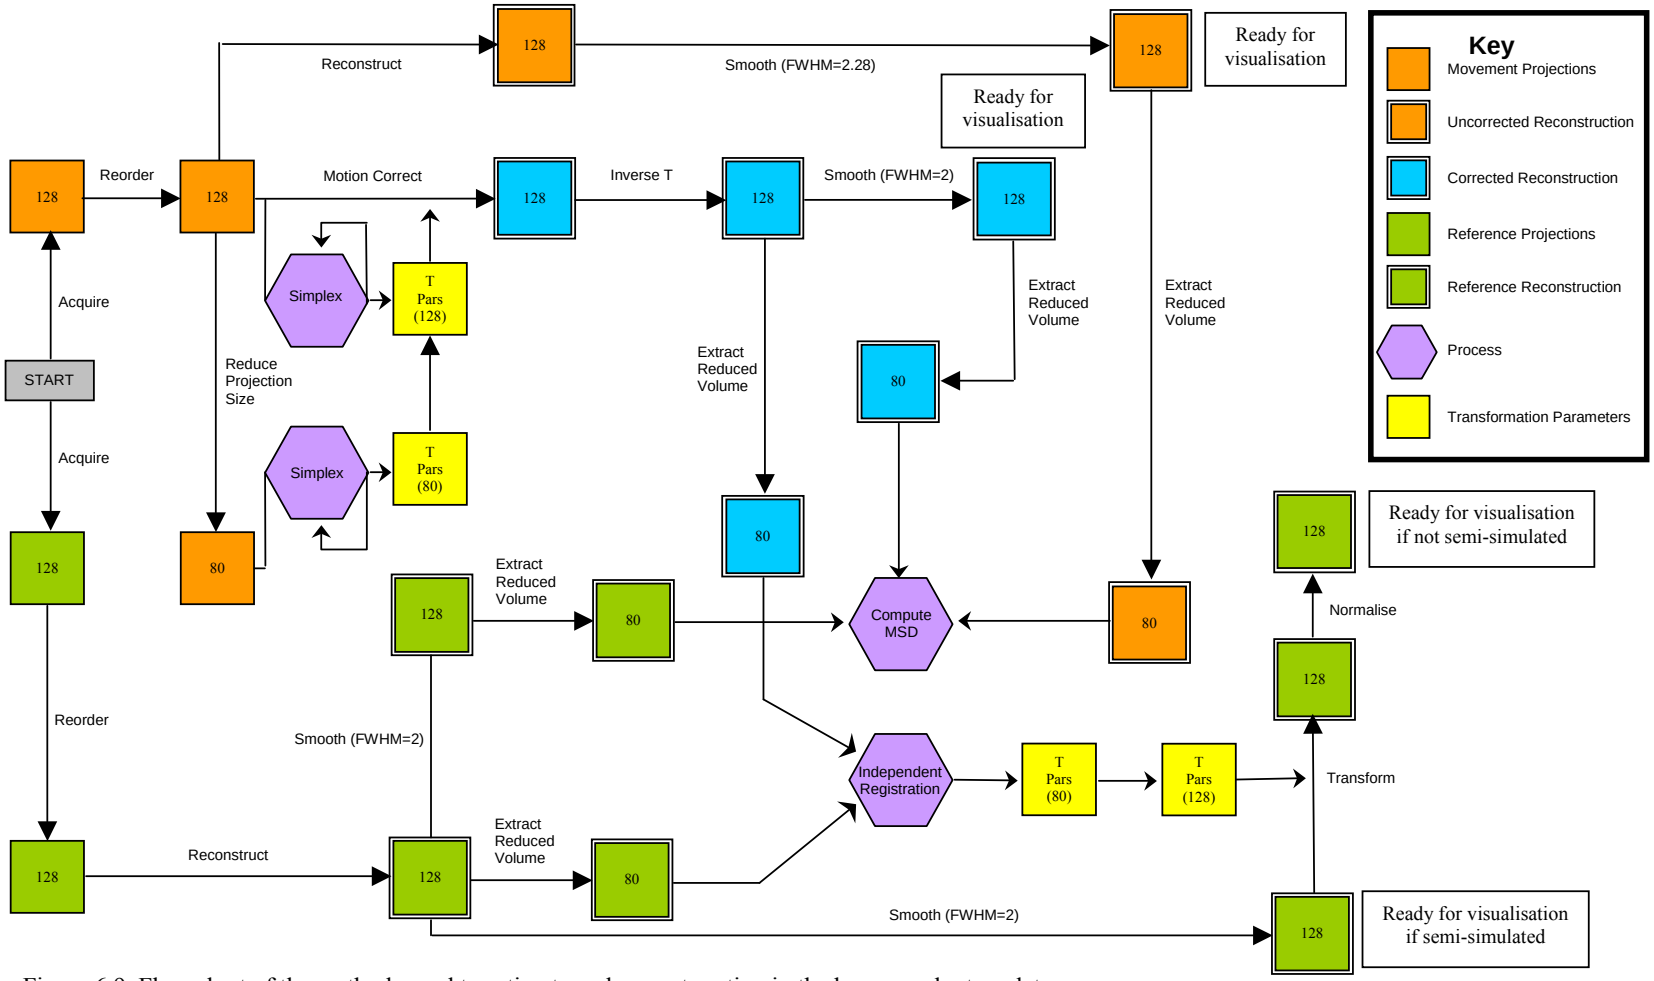

Figure 6.9. Flow chart of the methods used to estimate and correct motion in the human volunteer data.

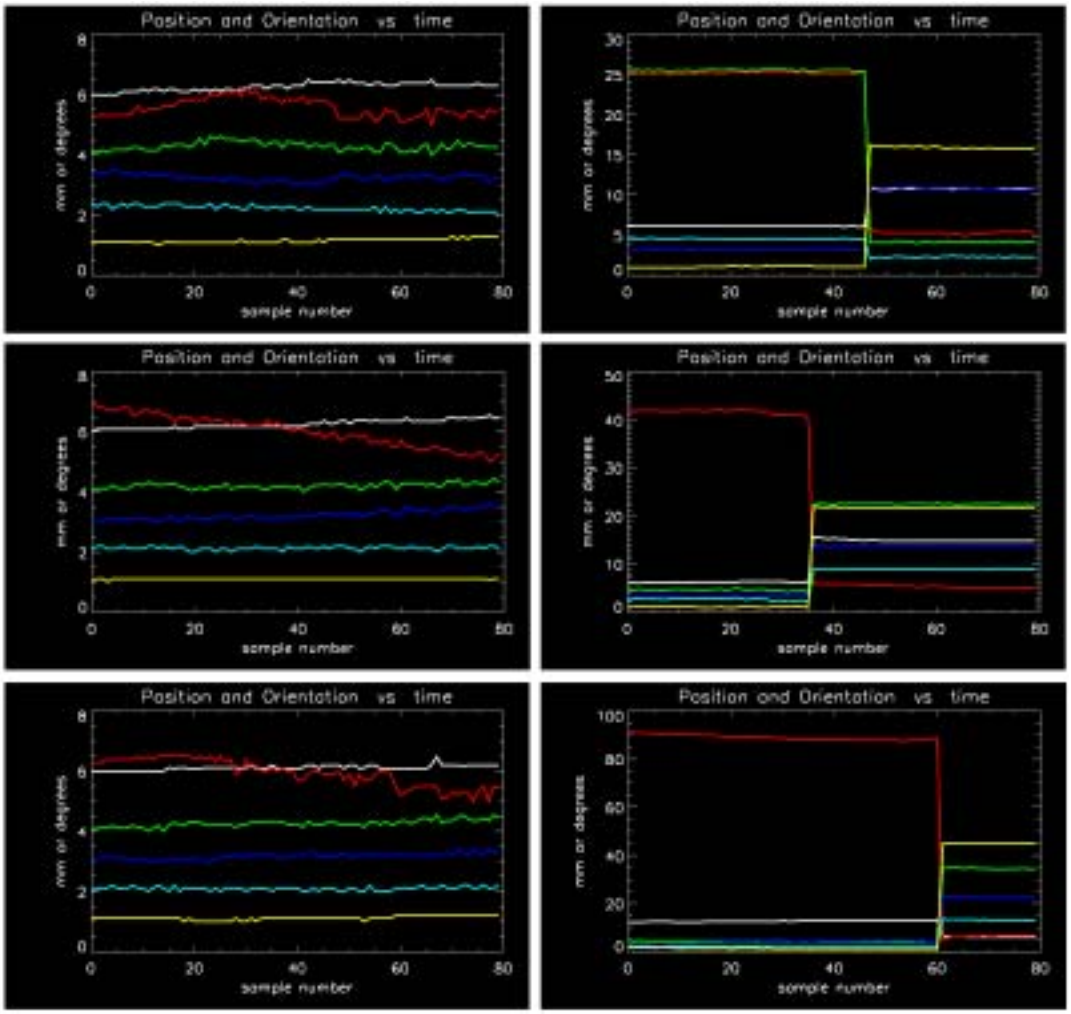

Figure 6.10. Polaris motion traces for the reference (left) and motion (right) scans of volunteers 1 to 3 (top to bottom respectively). 80 samples were recorded for each scan. The colour coding is: cyan, x-rotation; blue, y-rotation; orange, z-rotation; white, x-translation; red, y-translation; green, z-translation. It can be seen from the traces in the left column that there was minimal fluctuation of the six DOF during the reference scans. Even the drifting of the y-translation is seen to be negligible in comparison to the requested motions represented in the right column. The movement scans (right column) reflect excellent stability either side of a sharp, discrete change in location.

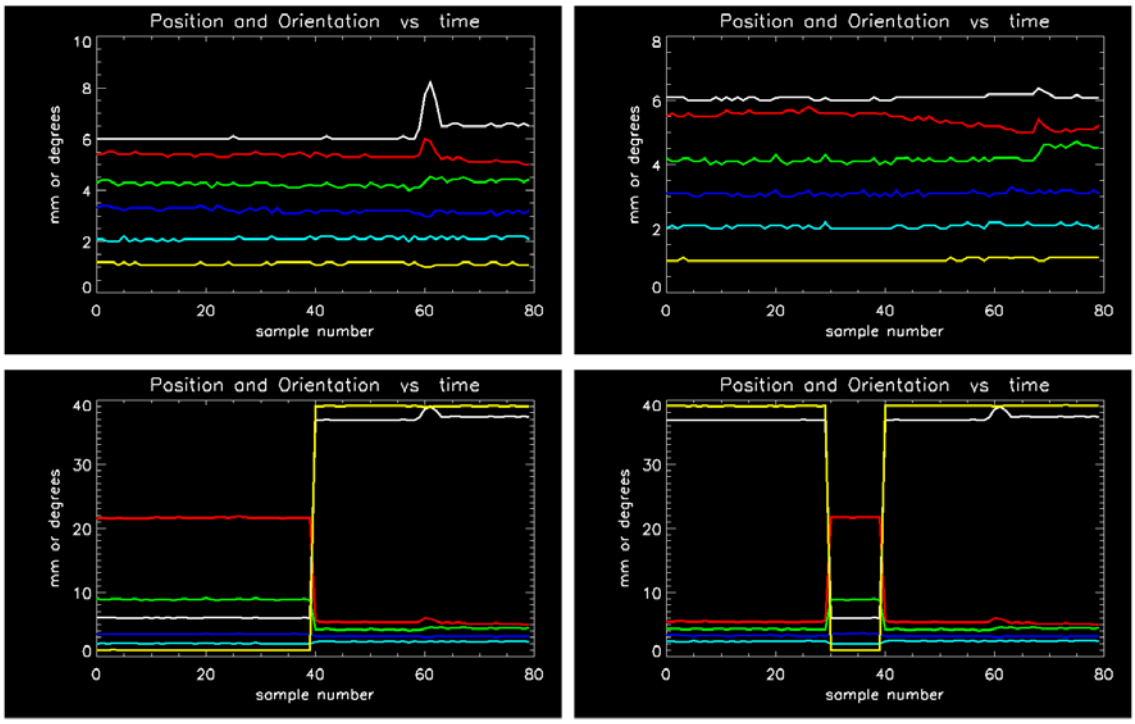

Figure 6.11. Polaris motion traces for the reference (top left) and motion (top right) scans of volunteer 4, plus the two semi-simulated traces – one with $^1/_2$ the projections affected (bottom left) and the other with $^1/_8$ of the projections affected (bottom right). 80 samples were recorded for each scan. The colour coding is the same as in Figure 6.10. The top traces show that volunteer 4 held a single position stably for each scan, with minimal fluctuation. The semi-simulated scans (bottom row) reflect a clean, sharp movement and indicate that this was much more significant than the involuntary fluctuations in the source data (eg. the x-translation (white) glitch near frame 60).

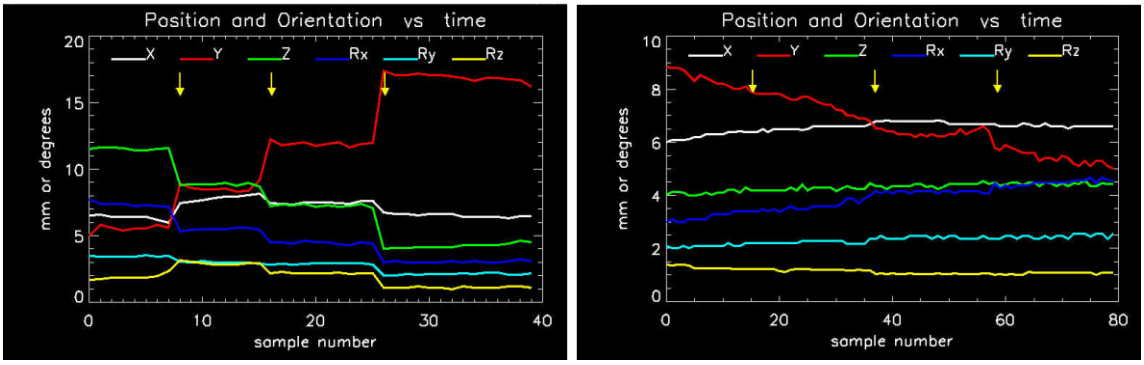

Figure 6.12. Polaris motion traces measured for two minimally constrained patients during a 30-minute brain SPECT study. The colour coding is the same as that in Figures 6.10 and 6.11. Left: The first patient demonstrated distinct changes in mean head location at multiple times throughout the acquisition. It is this type of motion we have sought to reproduce and correct for in our volunteer studies. The yellow arrows indicate user-selected samples between which the change in head location will be computed. Right: the second patient demonstrated gradual drifting motion (particularly y-translational) similar to that observed in our volunteers. Here, yellow arrows indicate samples at which the study may be segmented in order to approximate motion as a series of discrete steps. Both of these types of motion present themselves in the clinical setting. Data is courtesy of Leighton Barnden, Queen Elizabeth Hospital, Adelaide.

Motion estimates extracted from our optimisations for each volunteer study are given in Table 6.1. The estimates recorded by the Polaris are shown in Table 6.2. Measurements in Table 6.1 for a particular volunteer are directly comparable with the corresponding measurement in Table 6.2 since all extracted estimates were converted to the original camera coordinate system. These data are presented graphically in Figure 6.13, blue lines representing sets of estimates obtained using our approach, the white line representing the Polaris measurement. In the case of volunteer 3 (Figure 6.13 (c)), the four extracted 6-parameter estimates with least deviation from the Polaris measurement have been coloured green for clarity. For each set of data it is clear that we were able to generate multiple sets of motion estimates that followed the trend of the Polaris closely. The likelihood of outliers was greatest for volunteer 3. Importantly, the volunteer 3 study contained the largest angular and translational movements as recorded by the Polaris (and verified visually using the cine). The smallest spread in extracted results was obtained for the two semi-simulated data sets (Figure6.13 (d) and (e)). This was expected because of the additional control enforced when the source data for these sets were acquired (§*6.2.3*).

Quantitative figures of improvement (MSDR) for the two semi-simulated data sets are shown in the right-hand column of Table 6.1. Motion correction resulted in an improvement of the uncorrected study (closer to the reference) by a factor of >4 for the first set ($^1/_2$ of projections corrupted), and a factor of approximately 2 for the second set ($^1/_8$ of projections corrupted). This difference is not surprising since a greater degree of improvement is expected for greater corruptions. Note this result is in accord with the digital phantom experiments (§*4.6*).

An important aim of the clinical experiments was to determine the conditions and constraints that would enable reliable motion estimates to be generated from data-driven optimisation. We assumed the Polaris data trend to be a reasonable reflection of patient movement for each volunteer. Data in Tables 6.1 and 6.2 were subjected to three basic analyses as a means of drawing general conclusions regarding optimisations on clinical scans:

(i) Observation of extracted estimates consistent with the Polaris trend

(ii) Determination of the apparent cause of outlier data points

(iii) Comparison of MSDR values obtained from optimisations for the two semi-simulated sets

The conclusions were:

- Reducing matrix size so that a greater proportion of the projection image was filled by brain resulted in estimates as good as or better than when a larger matrix with a lot of empty space around the brain was used. In particular, maximising the amount of brain in the projection appears necessary to obtain sensible estimates for large motions. For example, no reasonable estimates could be obtained for volunteer 3 using the full projection matrix size. Moreover, in general, when a smaller matrix was used, motion estimates appeared less sensitive to other parameters being varied.

- Failed (outlier) motion estimates generally exhibited significant deviation from the Polaris estimate in most DOF. In many cases these failed estimates were characterised by severely wrong registrations involving a large overlap of low-count background area (contributing minimally to the MSD), and a much smaller, very poorly registered high-count area of brain. This indicates that the typical clinical cost function is susceptible to local minima, though we suggest such solutions are far from the true solution and readily excludable by visual assessment. Thresholding was found to be effective in confining the solution to the region in which the sensible solution exists by removing background counts. This was of particular importance for volunteer 3 where the z-rotation was >40°.
- For most volunteers the fixed-mask performed just as well as a roving mask. However, it is preferable to use a roving mask where large movements are suspected or detected from the cine/comparison of projections. For volunteer 3, use of a fixed-mask excluded too much brain and therefore prevented robust optimisation. The roving mask permits the maximum amount of brain to be included in the projection comparison.
- Scatter correction had little effect on the estimates (and MSDR values). One exception was for volunteer 1 where it appears scatter correction (with reduced matrix) led to the outlier. Scatter subtraction does increase noise in the projections, though this did not appear to have an adverse effect on motion estimation. This is in accord with previous results suggesting the data-driven approach is very robust with respect to noise.
- Although smoothing did not generally result in outliers, it tended to reduce the accuracy of estimates by a small degree. This should be further tested on studies with better count statistics eg. patient scans where the administered dose is higher.

It should also be noted that optimisation can fail if the fixed mask is not chosen properly. We estimated the appropriate boundary of this mask from the cine: for all likely transformations tried by the downhill-Simplex routine, there should be no truncated activity distribution above this boundary. From the projection images it is not possible to guarantee this; however, if motion is not too severe, and if chosen conservatively, the boundary will likely be valid. Choosing the boundary too low in an attempt (apparently) to include more brain can result in biasing and be counter-productive. For example, consider the two fixed-mask optimisations for volunteer 2 in Table 6.1. Here, projections of size 128^2 were used with no thresholding, smoothing or scatter correction. A mask excluding the bottom 16 rows resulted in an outlier set of motion parameters, whereas the mask excluding the bottom 30 rows resulted in a set of motion parameters consistent with the majority of other optimisations.

Table 6.1 (continued over page). State of optimisation variables and extracted parameters for all optimisations performed on the volunteer and semi-simulated data. For the semi-simulated sets, an MSDR describing the improvement due to correction is given in the far right column.

Study	Projection Size	Scatter Correction	Threshold (% of max. count)	Smoothing FWHM (pixels)	Truncation Solution	Motion Parameters					
						x-rot. (°)	y-rot. (°)	z-rot. (°)	x-trans. (pixels)	y-trans. (pixels)	z-trans. (pixels)
Volunteer 1	128	X	10	2	Roving	4.03	-2.88	-24.88	0.23	0.49	-5.11
	128	Yes	10	2	Roving	3.70	-2.93	-24.59	0.16	0.30	-5.21
	128	X	X	X	Fixed	5.53	-3.34	-24.93	0.88	1.17	-5.17
	128	X	10	X	Fixed	3.55	-2.50	-25.43	-0.06	0.34	-4.80
	128	X	10	2	Fixed	3.83	-0.82	-24.94	-0.93	0.88	-5.01
	128	X	10	X	Roving	0.96	-2.23	-22.77	-0.90	-1.06	-5.09
	128	X	X	X	Roving	26.44	-40.51	-37.16	20.77	0.33	4.10
	80	X	10	X	Roving	3.75	-2.57	-24.25	-0.37	0.44	-5.06
	80	X	10	2	Roving	2.83	-3.30	-23.59	0.14	-0.24	-5.08
	80	Yes	10	2	Roving	9.57	-9.28	-22.41	4.93	2.34	-4.40
	80	X	X	X	Fixed	2.99	-2.88	-25.08	-0.07	-0.07	-5.20
	80	X	10	X	Fixed	2.43	-2.28	-24.32	-0.48	-0.24	-4.80
	80	X	10	2	Fixed	2.83	-2.49	-24.79	-0.31	-0.03	-5.05
Volunteer 2	128	X	10	X	Roving	-10.90	2.75	22.30	2.84	-3.99	-0.82
	128	X	10	2	Roving	-11.02	1.50	22.62	3.63	-3.83	-0.71
	128	X	X	X	Fixed	-11.62	0.89	22.73	4.10	-3.94	-0.83
	128	X	10	2	Fixed	-6.80	7.49	21.83	-0.51	-2.80	-1.39
	128	Yes	10	X	Roving	-11.33	2.25	21.86	3.17	-4.23	-0.86
	128	Yes	10	2	Roving	-12.02	2.27	22.83	3.48	-4.47	-0.74
	128	X	X	X	Fixed	-7.43	6.33	21.25	0.01	-2.97	-1.41
	80	X	X	X	Roving	-10.30	0.61	21.35	3.89	-3.45	-0.75
	80	X	10	X	Roving	-11.35	3.94	21.08	2.07	-4.69	-0.83
	80	X	10	2	Roving	-11.98	0.89	22.35	4.11	-4.26	-0.48
	80	Yes	10	X	Roving	-12.79	0.86	21.47	4.13	-4.74	-0.52
	80	Yes	10	2	Roving	-12.15	0.98	21.04	3.90	-4.54	-0.57
	80	X	X	X	Fixed	-10.92	2.21	21.79	3.10	-3.97	-0.84
	80	X	10	2	Fixed	-11.66	2.93	21.04	2.76	-4.63	-0.61

Study	Projection Size	Scatter Correction	Threshold (% of max. count)	Smoothing FWHM (pixels)	Truncation Solution	Motion Parameters						MSDR
						x-rot. (°)	y-rot. (°)	z-rot. (°)	x-trans. (pixels)	y-trans. (pixels)	z-trans. (pixels)	
Volunteer 3	128	X	X	X	Fixed	20.17	-11.58	-46.30	3.89	6.82	0.27	-
	128	X	10	X	Fixed	14.27	-6.45	-50.59	-1.09	6.41	-0.03	-
	128	X	10	X	Roving	23.46	4.93	-44.96	-1.33	16.52	0.71	-
	128	X	10	X	Roving	18.50	-11.53	-53.39	2.48	5.46	0.58	-
	128	X	20	X	Roving	23.97	-14.52	-49.40	6.36	6.50	2.45	-
	128	X	X	X	Roving	57.77	-38.22	-83.22	23.30	-3.77	15.31	-
	128	Yes	X	X	Roving	49.65	-10.53	-75.36	16.69	12.81	6.47	-
	80	X	10	X	Roving	0.05	29.74	-34.49	-22.27	11.81	4.88	-
	80	X	X	X	Roving	7.87	43.33	-58.40	-19.06	29.39	8.78	-
	80	Yes	X	X	Roving	-3.62	59.26	-46.46	-30.77	24.02	16.47	-
	80	X	10	2	Roving	13.87	22.19	-44.77	-12.81	19.37	2.06	-
	80	X	20	X	Roving	14.97	23.02	-41.18	-13.09	19.37	2.92	-
	80	Yes	10	2	Roving	11.60	21.92	-46.13	-13.52	18.30	2.14	-
	80	Yes	20	X	Roving	8.66	24.40	-50.05	-15.48	19.03	2.67	-
Semi-sim. 1	128	X	10	X	Roving	-2.31	-1.58	36.53	5.59	6.04	1.56	4.3
	128	X	X	X	Fixed	-2.12	-2.04	36.62	5.76	6.31	1.43	4.3
	128	X	10	2	Roving	-6.93	15.32	28.57	-1.87	-1.35	3.09	1.4
	80	X	10	X	Roving	-3.06	-1.80	35.59	5.79	5.58	1.61	4.3
	80	X	10	2	Roving	-2.93	0.34	34.66	4.80	4.52	1.72	3.6
	80	Yes	10	X	Roving	-4.44	-3.97	35.61	7.33	5.79	1.90	4.4
	80	Yes	10	2	Roving	-1.80	-4.45	36.55	6.72	7.17	1.37	4.0
	80	X	X	X	Fixed	-2.31	-2.03	35.94	5.74	6.07	1.50	4.3
	80	X	10	X	Fixed	-3.88	-2.41	36.02	6.43	5.37	1.96	4.3
	80	X	10	2	Fixed	-0.97	-3.21	36.22	5.84	7.07	1.46	4.1
Semi-sim. 2	128	X	10	X	Roving	-5.86	-5.67	34.25	8.19	5.55	2.43	1.9
	128	X	10	2	Roving	-8.21	-5.99	32.27	8.93	4.24	2.84	1.7
	128	Yes	10	X	Roving	-1.60	-2.36	32.79	5.45	6.38	1.43	2.0
	128	Yes	10	2	Roving	-0.72	-0.76	33.86	4.48	6.51	1.50	1.8
	128	X	X	X	Fixed	-3.32	-3.80	33.72	6.73	6.14	1.82	2.1
	80	X	10	X	Roving	-3.17	-2.89	33.67	6.30	5.87	1.97	2.0
	80	X	10	2	Roving	-3.12	-2.19	32.65	5.78	5.58	1.80	2.0
	80	Yes	10	2	Roving	-2.14	-4.37	33.29	6.58	6.80	1.49	2.0
	80	X	X	X	Fixed	-2.16	-3.64	33.64	6.30	6.64	1.62	2.0
	80	X	10	X	Fixed	-3.56	-2.85	34.21	6.38	5.69	2.02	2.0
	80	X	10	2	Fixed	-3.70	-3.47	33.62	6.66	5.78	1.96	2.1

Table 6.2. Projections corresponding to each major brain location in the volunteer studies and the Polaris measurement of the major motions. An MSDR value is provided in the right column for the semi-simulated datasets.

Study	Location	Acquired Projections			Change in Position (Polaris)						MSDR
		Det 1	Det 2	Det 3	x-rot. (°)	y-rot. (°)	z-rot. (°)	x-trans. (pixels)	y-trans. (pixels)	z-trans. (pixels)	
Volunt. 1	1	[1-4], [21-40]	[41-44], [61-80]	[81-84], [101-120]	7.2	-4.0	-14.8	-0.29	1.46	-5.42	
	2	[5-20]	[45-60]	[85-100]							
Volunt. 2	1	[1-20], [38-40]	[41-60], [78-80]	[81-100], [118-120]	-12.1	1.6	19.6	2.92	-5.10	-1.58	
	2	[21-37]	[61-77]	[101-117]							
Volunt. 3	1	[1-11], [21-40]	[41-51], [61-80]	[81-91], [101-120]	12.4	17.5	-40.9	-13.7	17.17	1.61	
	2	[12-20]	[52-60]	[92-100]							
Semi-sim. 1	1	[1-20]	[41-60]	[81-100]	0.3	-0.3	38.0	7.45	7.68	1.26	0.6
	2	[21-40]	[61-80]	[101-120]							
Semi-sim. 2	1	[1-35]	[41-75]	[81-115]	0.3	-0.3	38.0	7.45	7.68	1.26	0.3
	2	[36-40]	[76-80]	[116-120]							

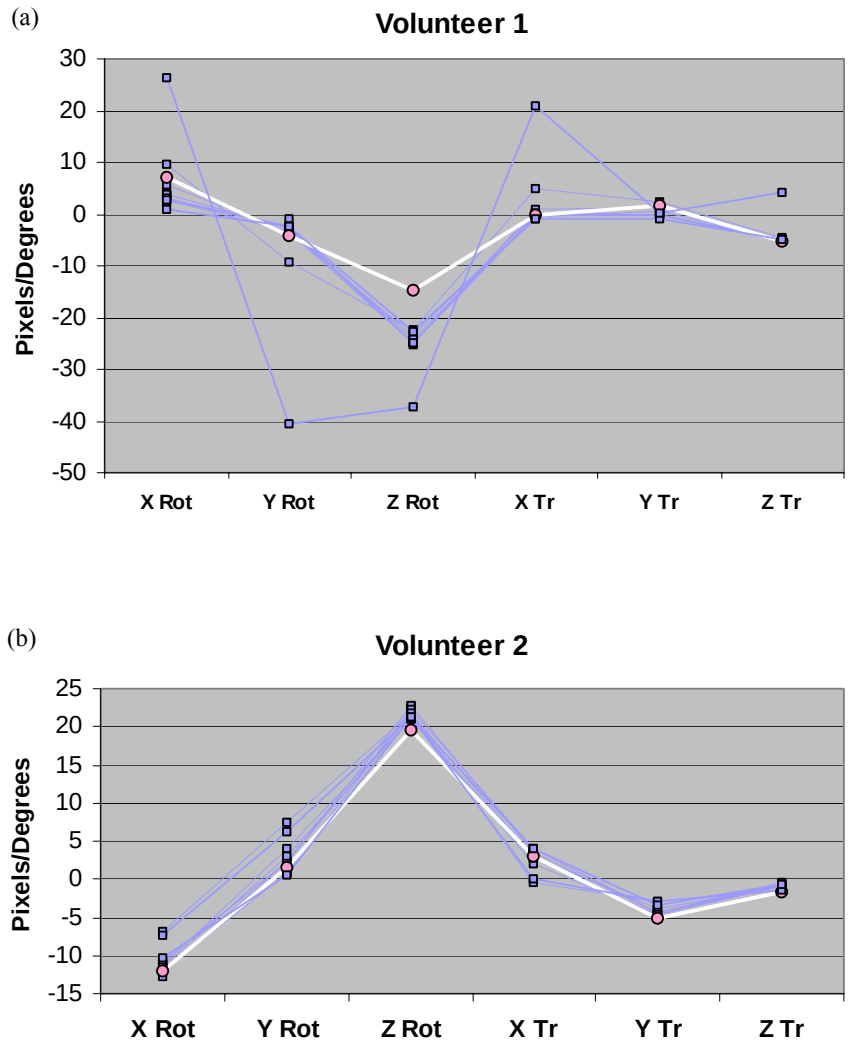

Figure 6.13 (a)-(b). Motion parameter sets (blue curves, blue points) extracted for (a) volunteer 1 and (b) volunteer 2. For both volunteers the majority of extracted parameter sets closely followed the trend of the Polaris (white curve, pink points) even though some parameter sets differed noticeably from the actual Polaris values in terms of magnitude.

(c)

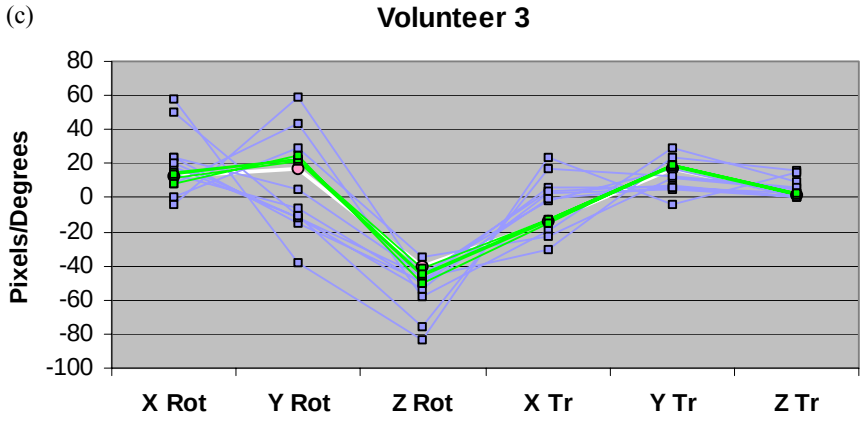

(d)

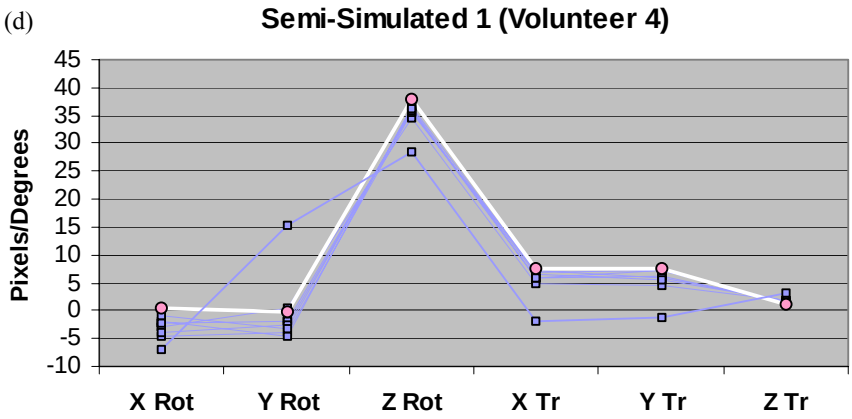

(e)

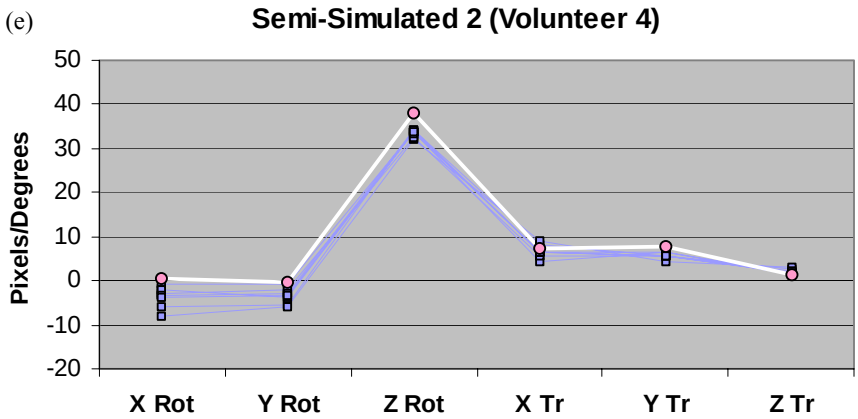

Figure 6.13 (cont.) (c)-(e). Motion parameter sets (blue curves, blue points) extracted for (c) volunteer 3, (d) semi-simulation 1, and (e) semi-simulation 2. The Polaris data is superimposed in each case (white curve, pink points). The optimisations for volunteer 3 were the most prone to failure due to a large rotation and translation. In this case the four parameter sets closest to the Polaris trend have been coloured green for clarity. For the semi-simulated sets there was only one failure out of all optimisations performed. These optimisations exhibited the least dispersion, and, although closely following the Polaris trend, were offset slightly from the Polaris values in most DOF.

A series of smoothed central brain slices from the uncorrected, reference, and corrected reconstructions are shown in Figure 6.14. Slices corrected using our method and the Polaris method are shown. Without motion correction there were clear motion artifacts manifested as distortion and smearing, asymmetry, and areas of low perfusion. This is in agreement with what has been described in the literature. After correction using both methods there was a clear reduction of these artifacts and an improvement in image contrast. Some corruption was still evident after applying motion correction. For volunteer 3 and both semi-simulated studies this residual corruption appeared to be worse for the Polaris correction.

For all volunteer studies we extracted motion estimates that closely followed the trend of the Polaris. However, in each case there was systematic variation from the Polaris measurements. This variation was further motivation for utilising semi-simulated data enabling rigorous quantification of motion correction. Variability between parameters was smallest for the semi-simulated sets, and extracted parameters obtained for semi-simulated data followed the Polaris parameters more closely than the initial 3 volunteer studies. However, the same systematic variation from the Polaris values was still present. For example, in the first semi-simulated set, the x and y-rotations were consistently larger and the x and y-translations consistently lower than those measured by the Polaris.

These results are evidence both that a semi-simulation approach is an effective and rigorous means of testing and validating the approach on clinical data, but also that the current implementation of the Polaris for measuring motion requires some refinement. We believe that the primary consideration should be the head target attachment. Operators at the institution from which we received the data were aware of the potential for the forehead strap to lift from the head and for the posterior strap to ‘catch’ on the patient bed, both scenarios causing head and target movement to become momentarily independent. Even though precautions were taken to minimise this for the semi-simulated studies, it is possible that such slippage still occurred and went unnoticed. Although the Polaris provided improvement, with the current implementation it cannot be assumed an accurate gold standard. Recently, staff have developed a new design of head target attachment aimed at minimising contact between the attachment sections and the bed and maintaining good, close contact with the head at all times. This design (Figure 6.15) has yet to be tested.

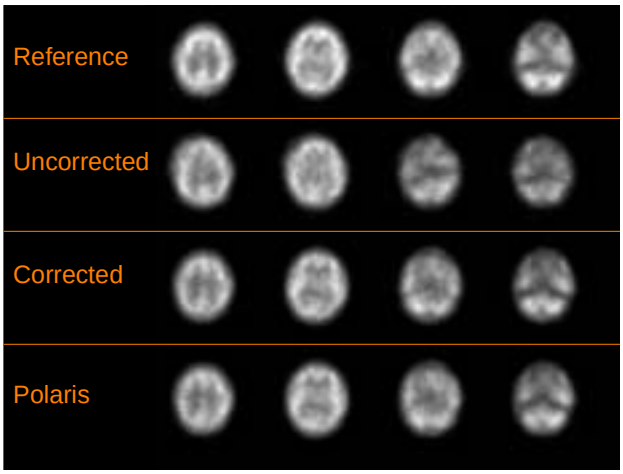

Figure 6.14a. Upper, mid, and lower brain slices for the reference, uncorrected, and corrected reconstructions of volunteer 1. Images were filtered to the same degree of smoothness (see main text). In the uncorrected slices there is a general loss of contrast and a severe smearing artifact in the antero-lateral cortex of the upper brain (left column). The cortical artifact was removed and contrast and symmetry much improved after correction using both methods.

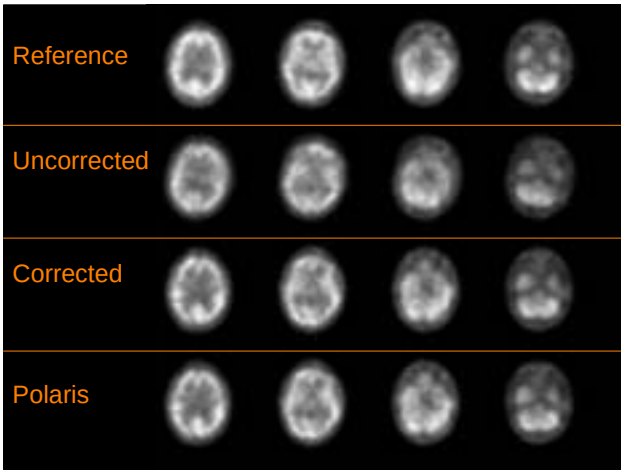

Figure 6.14b. Images for the second volunteer study. In the uncorrected slices there is a general blurring and loss of fine detail due to motion. Contrast is clearly improved after motion correction using both methods, however there is still distortion and loss of detail anteriorly and centrally.

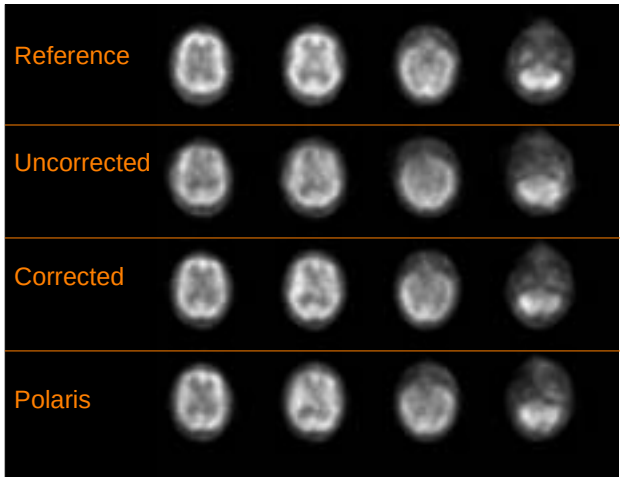

Figure 6.14c. Images for the third volunteer study. This study had the most severe motion out of the initial three volunteers. Severe 'ghosting' and distortion is apparent in the uncorrected slices. This was significantly reduced after correction. The symmetry of the cerebellum was better after data-driven correction than Polaris correction.

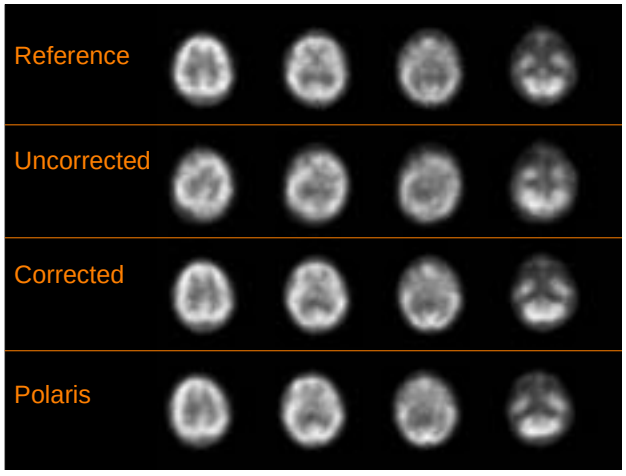

Figure 6.14d. Images for the first semi-simulated study. Severe distortion is apparent in the uncorrected slices since the study was evenly divided between two quite distinct head locations. The data-driven motion-corrected slices show better contrast and alignment than the Polaris-corrected slices.

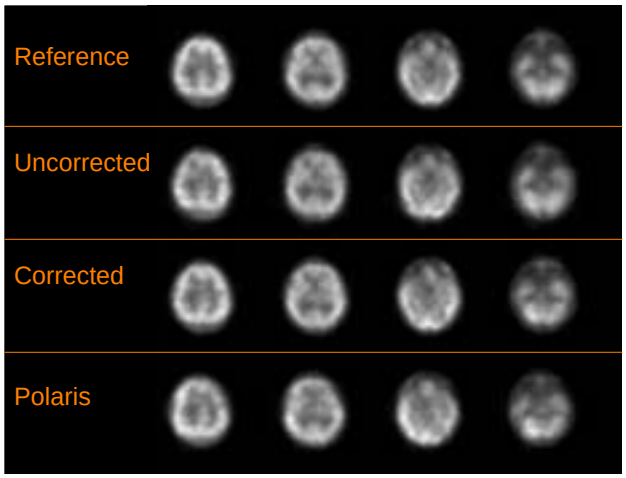

Figure 6.14e. Images for the second semi-simulated study. Compared to the first semi-simulated dataset, motion corruption in the uncorrected slices was far less apparent since only $^{1}/_{8}$ of the projections were affected by movement. The data-driven motion-corrected slices are clearly very similar to the reference slices and better than the Polaris-corrected slices.

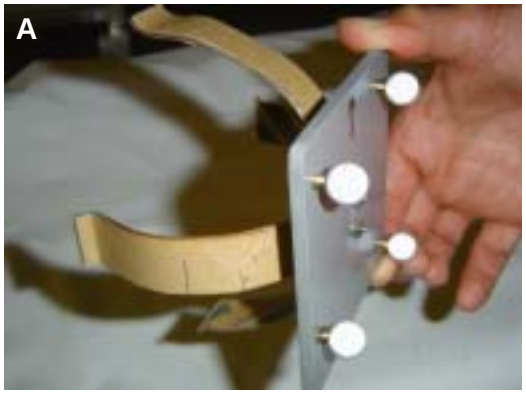

Figure 6.15. (a) New design for securing the head target (grey with white spherical reflectors) to the head. The design consists of four perspex 'clasps' that ensure constant contact with the head whilst minimising contact with the bed; (b) the design in action with an additional strap to provide extra stability of the target.

CHAPTER 7

SUMMARY, CONCLUSIONS, AND FURTHER WORK

7.1 Summary

This thesis describes a method for patient head motion correction that is aimed at improving the diagnostic quality and accuracy of brain SPECT images. We sought to demonstrate the feasibility of six DOF rigid-body motion correction in clinical brain SPECT studies using the novel data-driven methodology and to provide a thorough treatment of practical considerations when implementing this approach in the clinical setting.

Despite the use of restraining measures, patient head motion occurs frequently in brain SPECT studies and is a well-recognised source of artifacts in reconstructed data (Silver et al 1994, Cooper and McCandless 1995, Botvinick et al 1993, Green et al 1994). Apart from performing repeat studies, the alternative is to apply motion correction. Numerous strategies have been proposed in the literature to achieve this. Some methods lack data efficiency – for example those that involve multiple acquisitions (eg. Passalaqua and Narayanaswamy 1995, Britten et al 1998) or which fail to use all the acquired data (Germano et al 1994, Ivanovic et al 1998). Other methods involve marker sources being placed on the patient (Groch et al 1985, Ivanovic et al 2001). Methods requiring additional hardware for independent detection or measurement of motion add calibration and maintenance complications and often add extra time to already lengthy procedures (eg. Goldstein et al 1997, Green et al 1994, Lopresti et al 1999, Fulton et al 1999, Bloomfield et al 2003). The biggest shortcoming of most published methods is that they do not correct for motion in all six rigid-body DOF (eg. Arata et al 1995, Lee and Barber 1998, Li et al 1995).

Among the published data-driven methods, none provides six DOF motion correction using a single scan, and in a potentially fully automated fashion. A novel data-driven approach is described in this work which satisfies these criteria. Motion correction is based on measured projection data being incorporated into a reconstruction consistent with the 3D orientation at which they were acquired. The method brings together various elements present in published techniques: a comparison of measured projections with projections generated from the reconstruction (Arata et al 1995, Lee and Barber 1998); an optimisation of the reconstruction (Surova-Trojanova et al 2000); and a 3D implementation of OSEM reconstruction (Fulton 2000). It has long been recognised that the radon (projection) transform relates projection movement to object movement (eg. Milanfar 1999, Robinson and Milanfar 2001). To our knowledge, this thesis is the first time a data-driven projection registration has been applied in brain tomography for the estimation of six DOF motion.

Initially the data-driven approach was validated in computer simulations. Datasets containing multiple, six DOF motions were simulated. Motions varied in magnitude, angular location, and angular extent. Accurate motion parameters were extracted after a single iteration of the approach. Corrected slices displayed a clear reduction in symmetry and perfusion defects. This was supported quantitatively. There was no requirement that attenuation be incorporated in the estimation process, enabling considerable benefits in performance.

The second stage of validation was testing in real studies of the Hoffman brain phantom. Two single-movement studies and one double-movement study were acquired by arbitrarily reorienting (six DOF) the phantom during a scan. The Polaris optical tracking device was used to provide an independent measurement of the applied movements. This device uses infrared reflections to determine the position and orientation of markers attached to the object of interest. Since data were received from a separate institution, modifications were made to obtain consistency between the camera coordinate system relevant to the acquisition and the system defined for motion correction. The x and y-rotation and z-translation parameters extracted using the data-driven approach were in good agreement with the Polaris measurements, however a systematic discrepancy was observed for the z-rotation and x and y-translation parameters. In the case of the double-movement study, all parameters were in good agreement, and the motion-corrected slices closely resembled those obtained using the Polaris. Symmetry in the activity distribution was shown to cause extended, shallow troughs in the cost function; this appeared to be the cause of erroneous motion estimations obtained in the high-quality physical Hoffman phantom data.

The final stage of testing involved volunteer human subject data and constituted a preliminary clinical validation. Three volunteers underwent two standard triple-detector SPECT brain scans. For the first scan volunteers attempted to maintain a single head location; for the second scan they performed a single (six DOF) change in head location when prompted. A fourth volunteer attempted to maintain a single (but different) head location in each of two scans. Data from the fourth volunteer were used to simulate two additional single-movement studies. The Polaris provided an independent measure of subject motion. For all volunteers, sets of motion parameters that followed the trend of the Polaris closely could be generated. However, there was a systematic discrepancy between data-driven values and the Polaris values. Use of scatter correction had little influence on motion estimation. Use of smaller projections tended to provide estimates at least as good as larger projections. Pre-smoothing of acquired data generally led to less accurate motion estimates. Thresholding removed background counts and helped prevent wild estimates. Also, the need to account for data truncation during the optimisation was demonstrated, particularly in the case of large movements. For all volunteers, image symmetry and contrast improved after applying data-driven correction. In certain cases, correction was better than that obtained from the Polaris. The method used for attaching the target to the patient head is an important consideration. Poor attachment can lead to decoupling of target and head movement as is likely the cause of poor Polaris results.

In summary, the contributions of this thesis are:

- Development of a novel method capable of correcting for six DOF head motion in SPECT using the measured data alone
- Development of methodology for its implementation in SPECT
- Development and validation of methodology for six DOF data-driven motion correction in SPECT phantom studies
- Demonstration of the feasibility of six DOF data-driven motion correction in SPECT scans of human subjects
- Development of methodology for a clinical implementation of six DOF data-driven SPECT motion correction

7.2 Conclusions

We have described a novel data-driven approach capable of correcting for general rigid-body motion occurring during brain SPECT acquisitions. Practical aspects of the approach have been treated in detail: detection of motion, estimation of motion, reconstruction, attenuation correction, and validation. Successful correction of multiple complex movements in simulated data and real phantom data has been demonstrated. The feasibility of correcting for fully 3D head movements in SPECT scans of human subjects has also been demonstrated. We have considered discrete motions occurring between projections. The method is best suited to, but not limited to, this type of motion. Our technique is free of external gadgetry and is potentially fully automated. No prior knowledge of motion is required. It is easily adapted to different types of collimation and multi-detector geometries.

7.3 Further Work

(1) Clinical validation

A thorough clinical validation of data-driven motion correction for human subjects should be carried out. In particular, multiple controlled motions of volunteer subjects, as well as motion in minimally constrained patients, should be considered. A robust independent measurement of the motion using the Polaris should be available. This would require a modification to the head target attachment used in this work (eg. that of Figure 6.15). A quality reference study (motion-free) would be required for each motion study, acquired back-to-back with the motion study in order that the Polaris measurements from the two scans be compatible. Qualitative assessment of motion correction should take the form of a receiver operator characteristic (ROC) study.

(2) Improvement in Efficiency and Automation

There is a lot of scope to improve the efficiency of the technique. Currently the "bottle-neck" of the procedure in terms of time is the 3D transformation of the reconstruction during optimisation. An implementation of OSEM allowing projections to be incorporated at an arbitrary 3D orientation would significantly reduce the number of processor operations. Alternative optimisation algorithms may also provide an increase in speed.

The motion estimation stage was fully automated in this work. Automating the updating stage is a simple extension. Automating the motion detection phase poses more of a challenge and would require well-defined criteria for isolating and categorising the motion groups. The latter would enable the entire approach to be user-independent.

(3) Alternative reconstruction algorithms

It is recognised that identification of large angle groups could be less likely if motion is slow and progressive. In terms of motion correction this means more optimisations, the extreme case being a separate optimisation for each simultaneously-acquired angle group. For dual-90° detector geometry, this in turn would require reconstruction updates using subsets of two (pairs). Though OSEM may be limited here by the subset balance condition, the more general form of this iterative reconstruction algorithm, rescaled block iterative (RBI) algorithm, should have no such limitation. Further investigation of this is required. In theory then, the identification of angle groups is not a limitation, though at the current speed, failure to identify reasonably sized groups would make the algorithm prohibitively slow for practical use.

An additional complicating factor when finding motion relationships between angles is the choice of reconstruction subsets. Our observation is that this choice influences the quality of the partial reconstruction and the successive reconstruction updates. Further investigation into how influential OSEM subset ordering is in distinguishing angle groups and estimating motion is necessary. Again, subset-based reconstruction algorithms not reliant on subset size or ordering may provide improvements.

Although we have described methodology for multi-detector (non-opposite) configurations, we postulate that the approach could be implemented using a single detector gamma camera. Motion components orthogonal to the detector, although 'missed' in terms of motion detection, would not contribute corrupt information at that projection angle and therefore should not be of concern. Therefore, it is feasible that data-driven motion correction could be applied one projection at a time. This, however, would require the use of a reconstruction algorithm permitting singly populated subsets. The RAMLA algorithm is an option for investigation here.

(4) Understanding the effect of motion

Clearly the number of possible movements and the pattern of patient movement is prohibitively large to test exhaustively. Nevertheless, in this work a range of motions were tested to enable some assessment of how the type, amplitude, angular location, and angular duration of motion influenced performance of the algorithm. Both a broader range of movements in real subjects, and an investigation into what qualifies as 'significant' movement should be tested. To our knowledge, no comprehensive studies on the latter have been performed for brain SPECT. It is particularly important to assess the significance of motion as a source of artifacts visible over and above the noise likely to be encountered in the clinical setting.

(5) Application to other organs

A natural extension would be to apply the method in cardiac SPECT studies. The presence of significant activity in surrounding organs such as the liver will influence correction, but the nature of this problem is similar to the truncation problem already encountered in this work; a solution is therefore feasible.

BIBLIOGRAPHY

Anderson AW and Gore JC. Analysis and correction of motion artifacts in diffusion weighted imaging. *Magn. Reson. Med*., 32:379-387, 1994.

Arata LK, Pretorius PH, and King MA. Correction of organ motion in SPECT using reprojection data. In *Proceedings of the 1995 IEEE Nuclear Science Symposium and Medical Imaging Conference*, volume 3, pages 1456-1460, San Francisco, CA, USA, 1995.

Ardekani BA, Braun M, Hutton BF, Kanno I, and Iida H. A fully automatic multimodality image registration algorithm. *J. Comput. Assist. Tomog.*, 19:615-623, 1995.

Bailey DL. Transmission scanning in emission tomography. *Eur. J. Nucl. Med*., 25:774-785, 1998.

Barber DC. Registration of low resolution medical images. *Phys. Med. Biol.*, 37:1485-1498, 1992.

Bergstrom M, Boethius J, Eriksson L, Greitz T, Ribbe T, and Widen L. Head fixation device for reproducible position alignment in transmission CT and positron emission tomography. *J. Comput. Assist. Tomog.*, 5:136-141, 1981.

Bidaut LM and Vallee JP. Automated registration of dynamic MR images for the quantification of myocardial perfusion. *J. Magn. Reson. Imag*., 13:648-655, 2001.

Blodgett T, et al. The effect of respiratory motion on PET/CT image quality. *J. Nucl. Med.*, 43 (5 suppl.), 58P, 2002. (Abstract.)

Bloomfield PM, Spinks TJ, Reed J, Schnorr L, Westrip AM, Livieratos L, Fulton R, and Jones T. The design and implementation of a motion correction scheme for neurological PET. *Phys. Med. Biol.*, 48:959-978, 2003.

Boes JL and Meyer CR. Multi-variate mutual information for registration. In Taylor C and Colchester A, editors, *Medical Image Computing and Computer-Assisted Intervention, Second International Conference*, pages 606-612. Springer-Verlag, Berlin, 1999.

Botvinick EH, Zhu YY, O'Connell WJ, and Dae MW. A quantitative assessment of patient motion and its effect on myocardial perfusion SPECT images. *J. Nucl. Med.*, 34:303-310, 1993.

Boxt LM. Intravenous digital subtraction angiography of the thoracic and abdominal aorta. *Cardiovascular Interventional Radiology*, 6:205-213, 1983.

Britten AJ, Jamali F, Gane JN, and Joseph AEA. Motion detection and correction using multi-rotation 180° single-photon emission tomography for thallium myocardial imaging. *Eur. J. Nucl. Med.*, 25:1524-1530, 1998.

Browne J and Di Pierro AR. A row-action alternative to the EM algorithm for maximising likelihood in emission tomography. *IEEE Trans. Med. Imag.*, 15:687-699, 1996.

Bruyant PP, King MA, and Pretorius HP. Correction of the respiratory motion of the heart by tracking of the centre of mass of thresholded projections: a simulation study using the dynamic MCAT phantom. *IEEE Trans. Nucl. Sci.*, 49:2159-2166, 2002b.

Bruyant PP. Analytic and iterative reconstruction algorithms in SPECT. *J. Nucl. Med.*, 43:1343-1358, 2002a.

Buzug TM, Weese J, and Strasters KC. Motion detection and motion compensation for digital subtraction angiography image enhancement. *Phillips J. Research*, 51:203-227, 1998.

Byrne CL. Accelerating the EMML algorithm and related algorithms by rescaled block-iterative methods. *IEEE Trans. Imag. Processing*, 7:100-109, 1998.

Byrne CL. Block-iterative methods for image reconstruction from projections. *IEEE Trans. Imag. Processing*, 5:792-794, 1996.

Byrne CL. Convergent block-iterative algorithms for image reconstruction from inconsistent data. *IEEE Trans. Imag Processing*, 6:1296-1304, 1997.

Capek M, Mroz L, and Wegenkittl R. Robust and fast medical registration of 3D-multi-modality data sets. In *Proceedings of Medicon 2001, IX Mediterranean Conference on Medical and Biological Engineering and Computing*, pages 515-518, Pula, Croatia, 2001.

Cho K, Kumiata S, Okada S, and Kumazaki T. Development of respiratory gated myocardial SPECT system. *J. Nucl. Cardiol*., 6:20-28, 1999.

Chornoboy ES, Chen CJ, Miller MI, Miller TR, and Snyder DL. An evaluation of maximum likelihood reconstruction for SPECT. *IEEE Trans. Med. Imag.*, 9:99-110, 1990.

Collins DL, Holmes CJ, Peters TM, and Evans AC. Automatic 3-D model-based neuroanatomical segmentation. *Human Brain Mapping*, 3:190-208, 1995.

Cooper JA and McCandless BK. Preventing patient motion during tomographic myocardial perfusion imaging. *J. Nucl. Med.*, 36:2001-2005, 1995.

Cooper JA, Neumann PH, and McCandless BK. Detection of patient movement during myocardial perfusion imaging. *J. Nucl. Med.*, 34:1341-1348, 1993.

Cunneen ML, Schleyer PJ, and Rowe CC. Motion correction induced artifacts in myocardial imaging: A phantom study. ANZ Nucl. Med., (March suppl.), page 11, 2002. (Abstract.)

D. J. de Vries and S. C. Moore, Registration of CT and SPECT volumes using CT reprojections and planar transmission images. *Conference Record of the Nuclear Science Symposium 2001*, volume 3, pages 1441-1443. IEEE 2001.

Dann R, Hoford J, Kovacic S, Reivich M, and Bajcsy R. Evaluation of elastic matching system for anatomic (CT, MR) and functional (PET) cerebral images. *J. Comput. Assist. Tomog.*, 13:603-611, 1989.

Daube-Witherspoon ME, Yan YC, Green MV, and Kempner KM. Correction for motion distortion in PET by dynamic monitoring of patient position. *J. Nucl. Med.*, 31:186 (abstract), 1990.

Defrise M, Kinahan PE, and Michel C. Image reconstruction algorithms in PET. In Valk PE, Bailey DL, Townsend DW, and Maisey MN, editors, *Positron emission tomography: basic science and clinical practice*, pages 91-114. Springer, 2002.

Depuey GE and Garcia EV. Optimal specificity of thallium-201 SPECT through recognition of imaging artifacts. *J. Nucl. Med.*, 30:441-449, 1989.

Dhanantwari AC, Stergiopoulos S, Zamboglou N, Baltas D, and Vogt H. Correcting organ motion artifacts in x-ray CT systems based on tracking of motion phase by the spatial overlap correlator. II. Experimental study. *Med. Phys*., 28:1577-1596, 2001.

DiBella E, Barclay A, Eisner R, and Schaefer R. A comparison of rotation-based methods for iterative reconstruction algorithms. *IEEE Trans. Nucl. Sci.*, 43:3370-3376, 1996.

Eberl S, Kanno I, Fulton RR, Hutton BF, and Fulham MJ. Automated interstudy image registration technique for SPECT and PET. *J. Nucl. Med.*, 37:137-145, 1996.

Ebert M. Optimisation in radiotherapy II: programmed and inversion optimisation algorithms. *Aust. Phys. Eng. Sci. Med*., 20:220-230, 1997a.

Ebert M. Optimisation in radiotherapy III: stochastic optimisation algorithms and conclusions. *Aust. Phys. Eng. Sci. Med*., 20:231-241, 1997b.

Eisner RL, Churchwell A, Noever T, et al., Quantitative analysis of the tomographic thallium-201 myocardial bullseye display: critical role of correcting patient motion. *J. Nucl. Med.*, 29:91-97, 1988.

Eisner RL, Noever T, Nowak D, et al. Use of cross-correlation function to detect patient motion during SPECT imaging. *J. Nucl. Med.*, 28:97-101, 1987.

Enzmann DR and Freimarck R. Head immobilisation for digital subtraction angiography. *Radiology*, 151:801, 1984.

Faber TL, McColl RW, Opperman RM, Corbett JR, and Peshock RM. Spatial and temporal registration of cardiac SPECT and MR images: methods and evaluation. *Radiology*, 179:857-861, 1991.

Fahey FH and Harkness BA. SPECT imaging with rotating gamma camera systems. In Henkin RE, editor, *Nuclear Medicine, volume 1*. St Louis, Mosby-year book, 1996.

Freire L, Roche A, and Mangin J-F. What is the best similarity measure for motion correction in fMRI time series? *IEEE Trans. Med. Imag.*, 21:470-484, 2002.

Friedman J, Van Train K, Maddahi J, et al. Upward creep of the heart: a frequent source of reversible defects during thallium-201 stress-redistribution SPECT. *J. Nucl. Med*., 30:1718-1722, 1989.

Fulton RR, Eberl S, Meikle SR, Hutton BF, and Braun M. A practical 3D tomographic method for correcting patient motion in clinical SPECT. *IEEE Trans. Nucl. Sci.*, 46:667-672, 1999.

Fulton RR, Hutton BF, Braun M, Ardekani B, and Larkin R. Use of 3D reconstruction to correct for patient motion in SPECT. *Phys. Med. Biol.*, 39:563-574, 1994.

Fulton RR, Hutton BF, M. Braun, and B. Ardekani, Correction of gross patient motion artifacts in positron emission tomography. In Wilson JN and Wilson DC, editors, *Mathematical Methods in Medical Imaging II*, volume 2035, pages 166-170. Proc. SPIE, 1993.

Fulton RR, Meikle SR, Eberl S, Pfeiffer J, Constable CJ, and Fulham MJ. Correction for head movements in positron emission tomography using an optical motion-tracking system. *IEEE Trans. Nucl. Sci.*, 49:116-123, 2002.

Fulton RR, Tellmann L, Pietrzyk U, and Herzog H. Compensation for head movement in 3D PET. *Conference Record of the Nuclear Science Symposium 2001*, volume 3, pages 1580-1584. IEEE 2001.

Fulton RR. *Correction for patient movement in emission tomography*. PhD thesis, University of Technology, Sydney, 2000.

Geckle WJ, Frank TL, Links JM, and Becker LC. Correction for patient motion in SPECT: application to exercise thallium-201 cardiac imaging. *J. Nucl. Med.*, 29:441-450, 1988.

Gennert MA, Bruyant PP, Narayanan MV, and King MA. Calibrating optical images and gamma camera images for motion detection. *J. Nucl. Med.*, 43 (5 suppl.), page 223P, 2002a. (Abstract.)

Gennert MA, Bruyant PP, Narayanan MV, and King MA. Detecting patient motion in SPECT imaging using stereo optical cameras. *J. Nucl. Med.*, 43 (5 suppl.), page 223P, 2002b. (Abstract.)

Germano G, Chua T, Kavanagh P, Kiat H, and Berman DS. Detection and correction of patient motion in dynamic and static myocardial SPECT using a multi-detector camera. *J. Nucl. Med.*, 34:1349-1355, 1993.

Germano G, Kavanagh PB, Kiat H, Van Train K, and Berman DS. Temporal image fractionation: rejection of motion artifacts in myocardial SPECT. *J. Nucl. Med.*, 35:1193-1197, 1994.
Germano G, Nichols KJ, Cullom SJ, Faber TL, and Cooke CD. Gated perfusion SPECT: technical considerations. In Depuey EG, Garcia EV, and Berman DS, editors, *Cardiac SPECT Imaging, second edition*. Lippincott Williams and Wilkins, 2001.

Goldstein SR, Daube-Witherspoon ME, Green MV, and Eidsath A. A head motion measurement system suitable for emission computed tomography. *IEEE Trans. Med. Imag*., 16:17-27, 1997.

Green MV, Seidel J, Stein SD, Tedder TE, Kempner KM, Kertzman C, and Zeffiron TA. Head movement in normal subjects during stimulated PET brain imaging with and without head restraint. *J. Nucl. Med.*, 35:1538-1546, 1994.

Groch MW, Erwin WD, Turner DA, and Domnanovich JR. Dual-isotope motion correction technique for gated exercise scintigraphy. *J. Nucl. Med.*, 26:1478-1484, 1985.

Hamilton WR. *Elements of quaternions*. Chelsea Publishing Company, New York, 1969.

Harrington DP. Renal digital subtraction angiography. *Cardiovascular Interventional Radiology*, 6:214-223, 1983.

Hill DL, Batchelor PG, Holden M, and Hawkes DJ. Medical image registration. *Phys. Med. Biol.*, 46:R1-R45, 2001.

Hillier D and Wallis J. Myocardial SPECT motion artifact as a function of motion type and number of camera detectors. *J. Nucl. Med.*, 40 (suppl):126P (abstract), 1999.

Hoffman EJ, Cutler PD, Digby WM, and Mazziotta JC. Phantom to simulate cerebral blood flow and metabolic images for PET. *IEEE Trans. Nucl. Sci.*, 37:616-620, 1990.

Hoh CK, Dahlbom M, Harris G, Choi Y, Hawkins RA, Philps ME, and Maddahi J. Automated iterative three-dimensional registration of positron emission tomography images. *J. Nucl. Med.*, 34:2009-2018, 1993.

Huang SC and Yu DC. Capability evaluation of a sinogram error detection and correction method in computed tomography. *IEEE Trans. Nucl. Sci.*, 39:1106-1110, 1992.

Hudson HM and Larkin RS. Accelerated image reconstruction using ordered subsets of projection data. *IEEE Trans. Med. Imag.*, 13:601-609, 1994.

Hutton BF, Hudson MH, and Beekman FJ. A clinical perspective of accelerated statistical reconstruction. *Eur. J. Nucl. Med.*, 24:797-808, 1997.

Hutton BF, Kyme AZ, Lau YH, Skerrett DW, and Fulton RR. A hybrid 3D reconstruction/registration algorithm for correction of head motion in emission tomography. *Trans. Nucl. Sci.*, 49:188-194, 2002.

Ivanovic M, Pellot-Barakat C, Weber DA, Loncaric S, and Shelton DK. Effects of patient motion in coincidence studies on hybrid PET/SPECT system. *Conference Record of the Nuclear Science Symposium IEEE 2000*, volume 3, pages 16/49-53. IEEE 2000.

Ivanovic M, Weber DA, Loncaric S, Pellot-Barakat C, and Shelton DK. Patient motion correction for multicamera SPECT using 360° acquisition/detector. *IEEE Trans. Nucl. Sci.*: 989-993, 1998.

Jones WF. Real-time event stream correction for patient motion in clinical 3-D PET. Conference Record of the *Nuclear Science Symposium 2001*, volume 4, pages 2062-2064. IEEE 2001.

K. Thielemans and S. Mustafovic. Image reconstruction of motion corrected sinograms. Private communication, 2003.

Kiat H, Van Train KF, Friedman JD, et al. Quantitative stress-redistribution thallium-201 SPECT using prone imaging: methodologic development and validation. *J. Nucl. Med.*, 33:1509-1515, 1992.
Klein GJ, Reutter BW, and Huesman RH. Continuously rebinned gating for respiratory motion compensation in cardiac PET. *J. Nucl. Med.*, 43 (5 suppl.):144P (abstract) 2002.

Klein GJ, Reutter BW, Botvinick EH, Budinger TF, and Huesman RH. Fine-scale motion detection using intrinsic list mode PET information. In Staib L, editor, *Proceedings of the 2001 IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 71-78, 2001.

Klein GJ, Reutter BW, Ho MH, Reed JH, and Huesman RH. Real-time system for respiratory-cardiac gating in positron tomography. *IEEE Trans. Nucl. Sci.*, 45:2139-2143, 1998.

Kyme AZ, Hutton BF, Hatton RL, and Skerrett DW. Evaluation of factors that influence the performance of an automatic motion correction algorithm for brain SPECT. *Nucl. Med. Commun*., 22: 923 (abstract), 2001.

Kyme AZ, Hutton BF, Hatton RL, Skerrett DW, and Barnden LR. Practical aspects of a data-driven motion correction approach for brain SPECT. *IEEE Trans. Med. Imag*., 22:722-729, 2003.

Lalush DS. Iterative image reconstruction. In Wernick and Aarsvold, editors, *Emission Tomography: The fundamentals of PET and SPECT*. Elsevier (to appear).

Lee CC, Jack CR, Grimm RC, Rossman PJ, Felmlee JP, Ehman RL, and Riederer SJ. Real-time adaptive motion correction in functional MRI. *Magn. Reson. Med*., 36:436-444, 1996.

Lee KJ and Barber DC. Use of forward projection to correct for patient motion during SPECT imaging. *Phys. Med. Biol.*, 43:171-187, 1998.

Leslie WD, Dupont JO, McDonald D, and Peterdy AE. Comparison of motion correction algorithms for cardiac SPECT. *J. Nucl. Med.*, 38:785-790, 1997.

Levin DC, Shapiro RM, Boxt LM, Dunham L, Harrington DP, and Ergun DL. Digital subtraction angiography: Principles and pitfalls of image enhancement techniques. *Amer. J. Roentgenol.*, 143:447-454, 1984.

Li J, Jaszczak RJ, and Coleman RE. A filtered backprojection algorithm for axial head motion correction in fan-beam SPECT. *Phys. Med. Biol.*, 40:2053-2063, 1995b.

Li J, Jaszczak RJ, Wang H, and Coleman RE. A filtered backprojection algorithm for fan-beam SPECT which corrects for patient motion. *Phys. Med. Biol.*, 40:283-294, 1995a.

Linney NC and Gregson PH. Organ motion detection in CT images using opposite rays in fan-beam projection systems. *IEEE Trans. Med. Imag.*, 20:1109-1122, 2001.

Lopresti BJ, Russo A, Jones WF, et al. Implementation and performance of an optical motion tracking system for high resolution brain PET imaging. *IEEE Trans. Nucl. Sci.*, 46:2059-2067, 1999.

Lu W and Mackie TR. Tomographic motion detection and correction directly in sinogram space. *Phys. Med. Biol.*, 47:1267-1284, 2002.

Maas LC, Frederick B, and Renshaw P. Recoupled automated rotational and translational registration for functional MRI time series data: the DART registration algorithm. *Magn. Reson. Med*., 37:131-139, 1997.

Maes F, Vandermeulen D, and Suetens P. Comparative evaluation of multiresolution optimisation strategies for multimodality image registration by maximisation of mutual information. *Med. Imag. Anal*., 3:373-386, 1999.

Maintz JBA and Viergever MA. A survey of medical image registration, *Med. Imag. Anal*., 2:1-36, 1998.

Maintz JBA, Meijering EHW, and Viergever MA. General multimodal elastic registration based on mutual information. In Hanson KM, editor, *Proceedings SPIE, Medical Image Processing*, volume 3338, pages 144-154. SPIE Press, Bellingham, WA, 1998.

Malison RT, Miller EG, Greene R, McCarthy G, Charney DS, and Innis RB. Computer-assisted coregistration of multislice SPECT and MR brain images by fixed external fiducials. *J. Comput. Assist. Tomog.*, 17:952-960, 1993.

Matsumoto N, Berman DS, Kavanagh PB, et al. Quantitative assessment of motion artifacts and validation of a new motion-correction program for myocardial perfusion SPECT. *J. Nucl. Med.*, 42:687-694, 2001.

Mawlawi O, Weiss R, Shinn AK, Pidcock J, Slifstein M, and Laruelle M. Performance characteristics of a head immobilisation device for PET imaging. *J. Nucl. Med.*, 40:281P, 1999.

Meijering EHW, Niessen WJ, Pluim JPW, and Viergever MA. Quantitative comparison of sinc-approximating kernels for medical image interpolation. In Taylor C and Colchester A, editors, *Medical Image Computing and Computer-Assisted Intervention*, pages 210-217. Springer-Verlag, Berlin, 1999.

Meltzer CC, Bryan RN, Holcomb HH, et al. Anatomical localisation for PET using MR imaging. *J. Comput. Assist. Tomog.*, 14:418-426, 1990.

Menke M, Atkins SM, and Buckley KR. Compensation methods for head motion detected during PET imaging. *IEEE Trans. Nucl. Sci.*, 43:310-316, 1996.

Milanfar P. A model of the effect of image motion in the Radon transform domain. *IEEE Trans. Imag. Processing*., 8:1276-1281, 1999.

Nehmeh SA, Rosenzweig KE, Squire OD, et al. Effect of respiratory gating on reducing lung motion artifacts in PET imaging of lung cancer. *Med. Phys*., 29:366-371, 2002.
Nelder JA and Mead R. A simplex method for function minimization. *Computer Journal*, 7:308-313, 1965.

O'Connor MK, Kanal KM, Gebhard MW, and Rossman PJ. Comparison of four motion correction techniques in SPECT imaging of the heart: a cardiac phantom study. *J. Nucl. Med.*, 39:2027-2034, 1998.

Pandos G, Barnden L, Lineage H, Smith T, and Unger S. A survey of head movement during clinical brain SPECT using an optical tracking system. ANZ Nucl. Med., (March suppl.), page 30, 2002. (Abstract.)

Pappata S, Dehaene S, Poline JB, et al. In vivo detection of striatal dopamine release during reward: a PET study with [^{11}C]raclopride and a single dynamic scan approach. *Neuroimage*, 16:1015-1027, 2002.

Passalaqua AM and Narayanaswamy R. Patient motion correction of SPECT images: dual scan approach. *IEEE Trans. Nucl. Sci.*: 1270-1274, 1995.

Pelizzari CA, Chen GTY, Spelbring DR, Weichselbaum RR, and Chen C-T. Accurate three-dimensional registration of CT, PET, and/or MR images of the brain. *J. Comput. Assist. Tomog.*, 13:20-26, 1989.

Pellot-Barakat C, Ivanovic M, Herment A, Erlandsson K, and Shelton DK. Detection of motion in hybrid PET/SPECT imaging based on the correlation of partial sinograms. *IEEE Trans. Med. Imag.*, 20:1072-1083, 2001.

Pellot-Barakat C, MIvanovic, Weber DA, Herment A, and Shelton DK. Motion detection in triple scan SPECT imaging. *IEEE Trans. Nucl. Sci.*, 45:2238-2244, 1998.

Penczek PA, Grasucci RA, and Frank J. The ribosome at improved resolution: new techniques for merging and orientation refinement in 3D cryo-electron microscopy of biological particles. *Ultramicroscopy*, 53:251-270, 1994.

Penney GP, Weese J, Little JA, Desmedt P, Hill DLG, and Hawkes DJ. A comparison of similarity measures for use in 2D-3D medical image registration. *IEEE Trans. Med. Imag.*, 17:586-595, 1998.

Picard Y and Thompson CJ. Motion correction of PET images using multiple acquisition frames. *IEEE Trans. Med. Imag.*, 16:137-144, 1997.

Pickens DR and Price RR. Digital image motion correction by spatial warp methods. *Med. Phys*., 14:56-61, 1987.

Press WH, Flannery BP, Teukolsky SA, and Vetterling WT. Numerical Recipes in C: the art of scientific computing. Second edition. Cambridge University Press, 1992.

Prigent FM, Hyun M, Berman DS, and Rozanski A. Effect of motion on thallium-201 SPECT studies: a simulation and clinical study. J. Nucl. Med., 34:1845-1850, 1993.

Qi J and Huesman RH. Correction of motion in PET using event-based rebinning method: pitfall and solution. *J. Nucl. Med.*, 43 (5 suppl.), page 146P, 2002. (Abstract.)

Rabe FE, Yune HY, Klatte EC, and Miller RE. Efficacy of glucagon for abdominal digital angiography. *Amer. J. Roentgenol.*, 139:618-619, 1982.

Radau PE, Linke R, Slomka PJ, and Tatsche K. Optimization of automated quantification of ^{123}I-IBZM uptake in the striatum applied to parkinsonism. *J. Nucl. Med.*, 41:220-227, 2000.

Rahmim A, Bloomfield P, Houle S, et al. Practically feasible histogram-mode and list mode EM reconstruction with full motion compensation. (In press).

Reichmann K, Roede R, Boschen F, Kuehn KU, Oertzen JV, and Biersack HJ. Detection of head movement in CBF-SPECT with an annular gamma camera. *J. Nucl. Med.*, 43 (5 suppl.), page 223P, 2002. (Abstract.)

Richie CJ, Godwin JD, Crawford CR, Stanford W, Anno H, and Kim Y. Minimum scan speeds for suppression of motion artifacts in computer tomography. *Radiology*, 185:37-42, 1992.

Robinson D and Milanfar P. Efficiency and accuracy tradeoffs in using projections for motion estimation. In *Conference Record of the Thirty-Fifth Asilomar Conference on Signals, Systems and Computers*, volume 1, pages 545-550, 2001.

Roche A, Malandain G, Ayache N, and Prima S. Towards a better comprehension of similarity measures used in medical image registration. In *Second International Conference on Medical Robotics, Imaging and Computer-Assisted Surgery*, pages 555-566, Cambridge, UK. Springer-Verlag, 1999.

Rosenthal MS, Cullom J, Hawkins W, Moore SC, Tsui BMW, and Yester M. Quantiative SPECT imaging: a review and recommendations by the focus committee of the society of nuclear medicine computer and instrumentation council. *J. Nucl. Med*., 36:1489-1513, 1995.

Ruttiman UR, Andreason PJ, and Rio D. Head motion during positron emission tomography: is it significant? *Psychiatry Res*., 61:43-51, 1995.

Shepp LA and Vardi Y. Maximum likelihood reconstruction for emission tomography. *IEEE Trans. Med. Imag.*, 1:113-122, 1982.

Shepp LA, Vardi Y, Ra JB, Hilal SK, and Cho ZH. Maximum likelihood PET with real data. *IEEE Trans. Nucl. Sci.*, 31:910-913, 1984.

Silver KM, Currie GM, and McLaughlin AF. Patient motion characterisation in cerebral SPECT acquisition. ANZ Nucl. Med., June:22-25, 1994.

Slomka PJ, Gilbert AH, Stephenson J, and Cradduc T. Automated alignment and sizing of myocardial stress and rest scans to three-dimensional normal templates using an image registration algorithm. *J. Nucl. Med.*, 36:1115-1122, 1995.

Soares EJ, Byrne CL, and Glick SJ. Noise characterisation of block-iterative reconstruction algorithms: I. Theory. *IEEE Trans. Med. Imag*., 19:261-270, 2000.

Surova-Trojanova H, Barker CW, Carrasquillo JA, and Bacharach SL. Registration of planar emission images with reprojected CT data. *J. Nucl. Med.*, 41:700-705, 2000.

Tsui BMW, Segars WP, and Lalush DS. Effects of upward creep and respiratory motion in myocardial SPECT. *IEEE Trans. Nucl. Sci.*, 47:1192-1195, 2000.

Tsui BMW. Quantitative SPECT. In Henkin RE, editor, *Nuclear Medicine, volume 1*. St Louis, Mosby-year book, 1996.

Turkington TG, Hoffman JM, Jaszczak RJ, MacFall JR, Harris CC, Kilts CD, Pelizzari CA, and Coleman RE. Accuracy of surface fit registration for PET and MR brain images using full and incomplete brain surfaces. *J. Computer Assisted. Tomography*, 19:117-124, 1995.

Ulug AM, Barker PB, and Zijl PCM. Correction of motion artifacts in diffusion-weighted images using a reference phase map. *Magn. Reson. Med*., 34:476-480, 1995.

V. Calhoun, T. Adali, and G. Pearlson, A frequency-space approach for motion correction in functional MRI. In Proceedings of 10th Image and Multidimensional Digital Signal Processing Workshop IMDSP '98, 12-16 July, Alpbach, Austria. (Abstract.)

Van den Elsen PA, Maintz JBA, and Viergever MA. Geometry driven multimodality image matching. *Brain Topogr*., 5:153-158, 1992.

Van den Elsen PA, Pol ED, and Viergever MA. Medical image matching – a review with classification. *IEEE Eng. Med. and Biol*., March:26-39, 1993.

Van den Elsen PA, Pol ED, Sumanaweera TS, Hemler PF, Napel S, and Adler JR. Grey value correlation techniques used for automatic matching of CT and MR brain and spine images. In *Visualisation in Biomedical Computing*, Proc. SPIE, volume 2359, pages 227-237. SPIE Press, Bellingham, WA, 1994.

Wallis JW and Miller TR. An optimal rotator for iterative reconstruction. *IEEE Trans. Med. Imag.*, 16:118-123, 1997.

Watabe H, Choi E, Sato N, and Iida H. Evaluation of accuracy of motion correction with optical motion tracking system and PET. *J. Nucl. Med.*, 43 (5 suppl.), page 208P, 2002a. (Abstract.)

Watabe H, Sato N, Choi E, Bloomfield P, and Iida H. Development of motion correction technique for brain PET with optical motion tracking system and list mode acquisition. *J. Nucl. Med.*, 43 (5 suppl.), page 204P, 2002b. (Abstract.)

Weinzapfel BT and Hutchins GD. Automated PET attenuation correction model for functional brain imaging. *J. Nucl. Med.*, 42:483-491, 2001.

Wells WM, Viola P, Atsumi H, Nakajima S, and Kikinis R. Multi-modal volume registration by maximization of mutual information. *Med. Imag. Anal*., 1:35-51, 1996.

Woo SK, Watabe H, et al. Sinogram-based motion correction of PET images using optical motion tracking and list-mode data. *Conference Record of the Nuclear Science Symposium 2002*, volume 2, pages 830-834. IEEE 2002.

Woods RP, Cherry SR, and Mazziotta JC. MRI-PET registration with an automated algorithm. *J. Comput. Assist. Tomog.*, 17:536-546, 1993.

Woods RP, Cherry SR, and Mazziotta JC. Rapid automated algorithm for aligning and reslicing PET images. *J. Comput. Assist. Tomog.*, 16:620-633, 1992.

Yang C, Orphanoudakis SC, and Strohnbehn JW. A simulation study of motion artifacts in computed tomography. *Phys. Med. Biol.*, 27:51-61, 1982.

Zeffiro T. Clinical functional image analysis: artifact detection and reduction. *Neuroimage*, 4:S95-100, 1996.