

Alignment of Contrast Enhanced Medical Images

Andrew Melbourne

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
of
University College London.

Centre for Medical Image Computing
University College London

2009

I, Andrew Melbourne, confirm that the work presented in this thesis is my own.

Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

To mum

Abstract

The re-alignment of series of medical images in which there are multiple contrast variations is difficult. The reason for this is that the popular measures of image similarity used to drive the alignment procedure do not separate the influence of intensity variation due to image feature motion and intensity variation due to feature enhancement. In particular, the appearance of new structure poses problems when it has no representation in the original image. The acquisition of many images over time, such as in dynamic contrast enhanced MRI, requires that many images with different contrast be registered to the same coordinate system, compounding the problem. This thesis addresses these issues, beginning by presenting conditions under which conventional registration fails and proposing a solution in the form of a 'progressive principal component registration'. The algorithm uses a statistical analysis of a series of contrast varying images in order to reduce the influence of contrast-enhancement that would otherwise distort the calculation of the image similarity measures used in image registration. The algorithm is shown to be versatile in that it may be applied to series of images in which contrast variation is due to either temporal contrast enhancement changes, as in dynamic contrast-enhanced MRI or intrinsically in the image selection procedure as in diffusion weighted MRI.

Acknowledgements

I would like to thank my principal supervisor David Atkinson for help and advice over the last three years. In addition I would like to thank him for being remarkably patient with the many fairly abstract and unlikely ideas that we have discussed over that period. I also thank my second supervisor, David Hawkes, whose enthusiasm and overview of the direction of my thesis has been both insightful and motivating.

I also offer my thanks to those at the Institute of Cancer Research, with whom I worked in collaboration on the grant that funded my PhD (EPSRC grant no GR/T 20434/01). Thank you to David Collins, Martin Leach, Dow-Mu Koh and Keiko Miyazaki for their help in the many project meetings over the course of the grant. Thanks especially to Matt Orton for his help during our many discussions not only on Dynamic Contrast Enhanced MRI, but on various registration algorithms and many other tangential subjects. Thanks to Mark White, his in depth knowledge of MRI image formation and obscure programming languages has proved invaluable over the last few years and also to Bill Crum for his advice on fluid registration.

Thanks should also go to those who shared many drinks and listened to my problems during my *dottorato commedia*; particularly Payal, Darren, Tim, Ged, Irina and many others both within the Centre for Medical Image Computing and beyond. I would also like to thank my parents and my sister Clare for their support through school and university, who have always helped and encouraged me, in particular my mum who would have loved to see me complete my studies.

Finally I would like to thank my wife Jenny for her patience and support over the last few years, without her I am not sure I would have started a PhD, let alone finished one.

Contents

1	Introduction	21
1.1	Introduction	21
1.2	Chapter Summary	22
1.3	Magnetic Resonance Data used in this Thesis	25
2	Literature Review	28
2.1	Images	28
2.2	Image Registration	28
2.2.1	Cost Functions	29
2.2.2	Transformation Models	32
2.2.3	Recent Developments in Image Registration	36
2.3	Magnetic Resonance Imaging	36
2.3.1	Dynamic Contrast Enhanced MRI	38
2.3.2	Registration of DCE-MRI	43
2.3.3	Diffusion Weighted MRI	47
2.3.4	Registration of Diffusion Weighted MRI	49
3	Creation of Simulated Dynamic Contrast Enhanced MRI Data	50
3.1	Introduction	50
3.1.1	Finding an intrinsic T1 map	51
3.2	Developing A Liver Model	52
3.2.1	Liver Deformation	52
3.2.2	A Contrast Enhancement Model	56
3.2.3	A Tumour Model	59
3.2.4	Proposed Volume Preservation Modification	62
3.3	Motion Model Examples	65
3.4	Enhancement Model Examples	67
3.5	Conclusion	69
4	Cost Functions and Contrast Enhancement	70
4.1	Introduction	70

4.2	Conventional Cost-Functions	70
4.2.1	Method	70
4.2.2	Results	71
4.2.3	Conclusion	73
4.3	Cost Function Minimisation	79
4.3.1	Method	79
4.3.2	Results	80
4.3.3	Conclusion	80
4.4	The Cost Function Matrix Mean (CFMM)	83
4.4.1	Using Simulated DCE-MRI to investigate the CFMM	83
4.5	Choice of Anchor Image in Conventional Registration	85
4.6	Conclusion	87
5	Progressive Principal Component Registration (PPCR)	88
5.1	Principal Components Analysis	89
5.1.1	PCA for Functional Analysis	89
5.1.2	PCA Used for Data Compression	90
5.1.3	Formation of the Covariance Matrix	91
5.2	The PPCR Algorithm	93
5.3	Registration of Simulated Data using PPCR	96
5.3.1	Method	96
5.3.2	Results	97
5.3.3	Conclusion	101
5.4	The Effect of PPCR on Cost-Function Minimisation Space	102
5.4.1	Results	102
5.4.2	Conclusion	103
5.5	Choice of Anchor Image in Conventional Registration - Revisited	106
5.6	Analysis of the PPCR Algorithm Progress	108
5.6.1	Changing the number of Principal Components	108
5.6.2	Changing the number of Registration Iterations	110
5.7	Future Adjustment for large variations in Contrast Enhancement	112
5.8	Conclusion	113
6	Registration of Breath-hold Dynamic Contrast Enhanced MRI	114
6.1	Registration of 2D DCE-MRI Using PPCR	115
6.1.1	Introduction	115
6.1.2	Method	117
6.1.3	Results	118
6.1.4	Conclusion	123

6.2	Registration of 3D DCE-MRI Using PPCR	125
6.2.1	Introduction	125
6.2.2	Method	125
6.2.3	Analysis by MRIW	126
6.2.4	Analysis by Cost-Function Matrix Mean	132
6.2.5	Conclusion	132
7	Registration of Diffusion Weighted MRI	133
7.1	Introduction	133
7.2	Method	135
7.3	Results	137
7.3.1	Visual Inspection of Fractional Anisotropy Maps	137
7.3.2	Analysis of Fractional Anisotropy Variation	143
7.3.3	Analysis of Tensor Fitting Residuals	144
7.4	Conclusion	145
7.5	Diffusion MRI of the Liver	146
7.5.1	Introduction	146
7.5.2	Method	146
7.5.3	Results	149
7.5.4	Conclusion	149
8	Kullback Leibler Assisted Image Matching and Patching (KLAMP)	152
8.1	Introduction	152
8.2	Method	152
8.3	Results	156
8.3.1	Inspection of Driving Force Gradients using KLAMP	156
8.3.2	Setting the KLD Threshold in KLAMP	159
8.3.3	Float Image Pre-Processing using KLAMP	160
8.3.4	Joint Entropy Recalculation Using KLAMP	165
8.4	Conclusion	169
9	Conclusion	170
9.1	Summary	170
9.2	Future Work	173
	Appendices	174
A	Formulation of Fluid Equation	175

B	Formulation of Cost-Function Gradients	177
B.1	Change in Sum of Squared Difference with Pixel Displacement	177
B.2	Change in Cross Correlation with Pixel Displacement	177
B.3	Change in Joint Entropy with Pixel Displacement	178
B.4	Change in Mutual Information with Pixel Displacement	179
B.5	Change in Normalised Mutual Information with Pixel Displacement	180
B.6	Alternative Information Based Cost-Functions	180
C	Formation of a Fluid-based Image Registration Algorithm	182
D	Formation of a KLAMP capable Fluid-based Image Registration Algorithm	183
E	List of Movies included on supplementary CD	184
	Bibliography	184

List of Figures

2.1	Illustrative example timing diagram for a Gradient Echo sequence. Top , Standard gradient echo protocol. Bottom Spoiled Gradient Echo including a spoiler gradient for more rapid acquisition, as used in Gd-DTPA contrast-enhancement imaging.	39
2.2	Illustrative example timing diagram for a Spin Echo acquisition sequence.	40
2.3	Example schematic timing diagram for a Diffusion Weighted Spin Echo acquisition sequence. Diffusion gradients (along any desired direction, but shown here along G_{ss}) are added either side of the 180° RF pulse, causing spin phase shifts that are refocussed dependent on the position and motion of the spins. For acquisition time reasons, the read-out block is likely to consist of an EPI sequence.	48
3.1	Example force seeds added to an underlying base image (left) and resulting deformation field (right). The green points represent centres of superior-inferior driving force from which the image displacement is calculated using an elastic equation to give a corresponding deformation field (right hand image). Breathing is modelled by varying the magnitude of the force centres by a raised sinusoid.	54
3.2	Using a spline model to describe breathing variation for six consecutive breath-holds. An initial spline (red) (shown slightly saw-tooth so that breathing in takes longer) is defined. The initial spline nodes (green/yellow) are allowed to vary from their defined locations in subsequent breaths with a Gaussian distribution in both time and magnitude.	55
3.3	Example of 3D Gross Abdominal Segmentation overlaid on anatomical reference images. Segmented by hand into liver (white), kidney (green), aorta & vascular features (yellow), heart (left & right side (red & blue respectively)). Pathology may also be marked (magenta).	58
3.4	2D pathology construction. a Circular boundary. b distorted boundary, smoothed with splines. c boundary thickness definition. d labelling pharmacokinetic features (colour coding interior and exterior).	59
3.5	Example tumours added to underlying (pre-enhancement) image. Segmentation colours correspond to different pharmacokinetic parameters for tumour rim and tumour core. For tumour generation process see text.	60

3.6	2D pathology examples over an approximately 3 minute time period for top) peripheral tumour enhancement; middle) filling tumours where the rim enhances quickly and later the tumour core enhances; bottom) necrotic core tumour where the rim enhances quickly and the core does not enhance.	61
3.7	Example intensity-time curves used to simulate the uptake profile of major organs (left) and curves used in Section 3.2.3 to model pathology. Signal is generated for $\alpha = 24^\circ$, $TR = 4ms$ and for intrinsic tissue $T1$ found from Table 3.1. Note that these curves have an unrealistically rapid wash-out due to an implementation error that should be resolved prior to additional work regarding biological pharmacokinetic parameter extraction. See text for parameter choice and further clarification.	61
3.8	Maintaining rigid object shape in an elastic deformation: a Original Square. b Square moved downwards using elastic deformation. c Unmodified elastic deformation. d affine deformation patched into elastic deformation, note the unacceptable deformation discontinuity at the boundary.	63
3.9	Maintaining rigid object shape in an elastic deformation: a Original Square. b Square moved downwards using elastic deformation. c Unmodified elastic deformation. d affine deformation patched into elastic deformation, now interpolated between affine and elastic deformations for comparison with Figure 3.8 using the method described in Table 3.5	64
3.10	Example deformation for superior-inferior <i>breathing motion</i> with 10 image cycle length with additional random medial-lateral deformations.	65
3.11	Example difference images (with original un-deformed image) for superior-inferior <i>breathing motion</i> with 10 image cycle length for comparison with 3.10.	66
3.12	Example contrast enhancement time course for coronal images.	67
3.13	Example contrast enhancement time course for sagittal images.	68
4.1	Example real contrast-enhanced images b and c and pre-enhancement image a demonstrating passage of bolus through right and left sides of the heart and aorta. Difference images for b-a and c-a are also shown to illustrate differences in position of abdominal wall and superior liver.	71
4.2	Derived image-similarity local gradient images for Least Squares Cost Function. Top row pre-enhancement image and two post-enhancement images. Second row corresponding (normalised) force gradients in <i>x-direction</i> . Third row corresponding (normalised) force gradients in <i>y-direction</i> . Bottom row corresponding displacement vector fields.	74
4.3	As Figure 4.2 but for Cross-Correlation Cost Function.	75
4.4	As Figure 4.2 but for Joint Entropy Cost Function.	76
4.5	As Figure 4.2 but for Mutual Information Cost Function.	77
4.6	As Figure 4.2 but for Normalised Mutual Information Cost Function.	78

- 4.7 Two identical images **a** and **b**, float image is deformed by a known force and then the cost-function value is found between deformed float and anchor. **c**: difference image between **a** and **b**. **d**: difference image between **a** and **b** at maximum deformation. 81
- 4.8 Cost-function minimisation space for labelled similarity measure for images corresponding to Figure 4.7, identical images where the float is deformed by a known force of varying magnitude and direction. Distance from centre on X and Y axes represents force strength in that direction (zero force at centre). 81
- 4.9 Pre and Post Contrast Enhancement images **a** and **b**, float image is deformed by a known force and then the cost-function value is found between deformed float and anchor. **c**: difference image between **a** and **b** with no deformation (there is small existing misalignment). **d**: difference image between **a** and **b** at maximum deformation. 82
- 4.10 Cost-function minimisation space for labeled similarity measure for images corresponding to Figure 4.9, pre and post enhancement images where the float is deformed by a known force of varying magnitude and direction. Distance from centre on X and Y axes represents force strength in that direction (zero force at centre). 82
- 4.11 Plot of NMI-CFMM values for varying sinusoidal-motion magnitude in linear steps between minimum and maximum force strength and contrast agent 'dose' varied in linear steps before conversion to signal by the spoiled gradient echo equation (Equation 2.26). 84
- 4.12 Assessment of registration result of variations in target image selection using Cost Function Matrix Mean (see Section 4.4). For four separate datasets from Table 1.1, fluid registration proceeds using the *n*th image as the anchor image. The NMI Matrix Mean is shown for the result of registration using Cross-Correlation (blue) and Normalised Mutual Information (red). Also shown is the original NMI Matrix Mean before registration (magenta). 86
- 5.1 First four normalised principal components for DCE-MRI dataset. PC1 is a general enhancement profile incorporating the mean intensity change over a time-scale of a few minutes. PC2 appears to act to correct those pixels that are not enhancing. PC3 and PC4 appear to enhance PC1 in areas of rapid initial enhancement, further describing differences between pixels in the wash out phase. Extracting pharmacokinetic parameters from these components is likely to be difficult despite the fact that these four components contain 97% of the dataset variance. 90
- 5.2 **Top row**: Images from a DCE-MRI dataset for pre-enhancement, bolus arrival in the left heart, bolus arrival in the liver and late post-enhancement. **Subsequent Rows**: Images rebuilt using, on each row, 1,2,3,4 components respectively. 91

- 5.3 Illustration of the PPCR Algorithm for a dataset with five images. Principal Components Analysis of 5 images produces 5 Principal Components and 5 Weighting Maps. The algorithm increments the number of Principal Components used to generate target images at each iteration, recalculating the PCA after each iteration. The last (fifth) Principal Component is not used as this will result in five anchor images that are identical to their float images. 95
- 5.4 Graphs of absolute image residual displacement with time for varying motion (levels 1-4 corresponding to 9mm, 18mm, 26mm, 35mm maximum displacements. **a-d** Cyclic breathing motion for contrast-enhancement level 3 (step-like curve for PPCR) shows separation of final registration position between two locations *see text*. **e-h** Breath-hold depth inconsistency for contrast-enhancement level 3 showing artefacts for direct fluid registration under increasing enhancement. 97
- 5.5 **a)** Source image from dataset with motion level 4 (up to 35mm displacement) and enhancement level 3. **b)** PPCR generated target image from first iteration for model with cyclic breathing motion for comparison with **c)** PPCR generated target image from first iteration for model with breath-hold depth inconsistency. The ill-defined nature of **b** results in separation of final registration position *see text*. 98
- 5.6 Graphs plotting statistics for variations in **a)** tumour volume change for breathing-motion extent and enhancement level, note the trend for decreasing tumour size in the gold standard. **b)** average joint entropy of target and source images per simulation for direct registration (to pre-contrast image) and to first set of target images generated by PPCR for breathing motion. **c)** for breath-hold depth inconsistency. Note increasing joint entropy with motion *and* enhancement. 99
- 5.7 Effect of PCA on joint image histogram formation from a real dataset of 20 2D images. **Column 1** original (real) dataset (images 1-10), **Column 2** first image in dataset, **Columns 3-5** images 1-10 rebuilt using 1-3 principal components. **Column 6** joint image histograms (x-axis float, y-axis anchor image intensities) of images in column 1 with those in column 2, **Columns 7-9** joint image histograms of Column 1 with Columns 3-5. Dispersion in Column 6 is the result of both motion and contrast enhancement, using PCA allows some removal of enhancement (lobe-like) dispersion, although the effect is reversed for the pre-enhancement images in Column 7. Inclusion of further principal components in Columns 8-9 removes. 100

- 5.8 Direct registration cost-function minimisation spaces for the first ten images from a DCE-MRI dataset. For ease of presentation, the negative-log values for MI, NMI and CC are shown; here light corresponds to good image alignment and dark to poor. The x-axis corresponds to a large medial-lateral force applied negative through positive from left to right, the y-axis is the equivalent for the superior-inferior force. The cost function comparison is taken between the corresponding image on the far left, and the float in the neighbouring column subject to the given deformation (*see text for clarification*). 104
- 5.9 PPCR cost-function minimisation spaces for the first ten images from a DCE-MRI dataset. For ease of presentation, the negative-log values for MI, NMI and CC are shown; here light corresponds to good image alignment and dark to poor. The x-axis corresponds to a large medial-lateral force applied negative through positive from left to right, the y-axis is the equivalent for the superior-inferior force. The cost function comparison is taken between the corresponding image on the far left and its neighbour, subject to the given deformation (*see text for clarification*). 105
- 5.10 Assessment of registration result of variations in target image selection using Cost Function Matrix Mean (see Section 4.4). For four separate datasets from Table 1.1, fluid registration proceeds using the *n*th image as the anchor image. The NMI Matrix Mean is shown for the result of registration using Cross-Correlation (blue) and Normalised Mutual Information (red). Also shown are the original NMI Matrix Mean before registration (magenta) and after registration using PPCR (green). 107
- 5.11 Example curves showing the approach of the PPCR deformation field towards the final deformation. Graphs show the per-pixel residual and the approach of the residual towards the final deformation position for four separate 2D DCE-MRI datasets from Table 1.1. 109
- 5.12 Example curves showing the approach of the PPCR deformation field towards the maximum deformation found when varying the number of internal registration iterations up to a maximum of 400 iterations in the component fluid registration algorithm. Graphs show the absolute residual (divided by total number of pixels) and the approach of the residual towards the maximum deformation position for two DCE-MRI datasets from Table 1.1. 111
- 5.13 Illustration of the PPCR adjustment process. After the formulation of principal components and weighting maps in the first iteration, we form target images from the first principal component. **Step 1** we find the mean (and standard deviation) sum-of-squared differences (SSD) between all float-anchor pairs. **Step 2** if any of the float-anchor pairs have an SSD value outside of one standard deviation from the mean SSD, we add further principal components until the SSD of that float-anchor pair is below one standard deviation from the mean float-anchor SSD. 112

- 6.1 Absolute-difference images demonstrating registration failure for fluid registration to first image in dataset and correct PPCR, of two post-contrast images. Images **a** and **e**, anatomical images for reference. The drawn region is the same across a row and provides a visual guide. Images **b** and **f**, no registration for comparison. Images **c** and **g**, registration to first image in dataset, **c**, correct gross-registration of the liver and **g** artefactual tumour motion. Images **d** and **h**, registration by PPCR with improved motion correction and reduced artefacts. 118
- 6.2 Plots of pixel intensity with time for pixels selected from anatomical images (left-hand column). The unregistered cases (NoReg) are shown for comparison. Artefactual oscillations in fluid timecourse for a) tumour artefacts due to nearby boundary motion & b) tumour artefacts due to tumour mis-registration (see Figure 6.1g). c) corrected motion of bright region within liver & d) corrected diaphragm position with large inferior displacement at timepoint 26. In this case both registration methods (Fluid & PPCR) identified the large displacement. Again for different patients, e) a correctly registered bright region and f) correct registration by both methods of inferior liver motion. Sagittal images, g) correct registration by both methods and h) correct registration by PPCR and mis-registration close to a tumour boundary by the simple fluid registration scheme. . . . 120
- 6.3 Results of fluid registration to the first image in the data set or PPCR. The comparison uses the sum of squared differences (SSD) between the image pixel data and the decaying exponential model of the post-enhancement phase. Results are shown as percentages of the respective SSD in the case of no registration. 122
- 6.4 Corresponding percentage change to the Cost-Function Matrix Mean (CFMM) values for Unregistered Data, Registration by direct-fluid registration and by PPCR. Also shown for PPCR is the highest CFMM found during the iterations. 123
- 6.5 Example Contrast Enhanced Images of the Abdomen. Pre-enhancement image **a** for comparison with image showing passage of bolus through heart **b** and subsequent portal enhancement of the liver **c**. 125
- 6.6 Patient 2: Model fitting results using dual-cosine arterial input function showing parameter maps for (**rows**): $K^{trans}(min^{-1})$ (range 0-3), v_e (range 0-1), HPI (range 0-1) and Pixel Residual (range 0-8). Each column represents registration by **Column 1**: Unregistered Data, **2**: Direct Fluid Registration and **3**: PPCR. 128
- 6.7 Patient 5: Model fitting results using dual-cosine arterial input function showing parameter maps for (**rows**): $K^{trans}(min^{-1})$ (range 0-1), v_e (range 0-1), HPI (range 0-1) and Pixel Residual (range 0-0.5). Each column represents registration by **Column 1**: Unregistered Data, **2**: Direct Fluid Registration, **3**: PPCR and **4**: Automatic registration on scanner using unknown algorithm. 129

6.8	Patient 6: Model fitting results using dual-cosine arterial input function showing parameter maps for (rows): $K^{trans}(min^{-1})$ (range 0-1), v_e (range 0-1), HPI (range 0-1) and Pixel Residual (range 0-0.8). Each column represents registration by Column 1 : Unregistered Data, 2 : Direct Fluid Registration, 3 : PPCR and 4 : Automatic registration on scanner using unknown algorithm.	130
6.9	Selection of local regions of interest for analysis of changes to pharmacokinetic parameters before and after registration. See Figure 6.8 for statistics over the whole liver. . . .	130
6.10	$K^{trans}(min^{-1})$ parameter distribution histograms for each of the three regions (rows) in Figure 6.9. Columns: results after each registration method as labelled.	131
7.1	Example DWI Slice, comprised of a B0 Image and fifteen diffusion direction images (with gradient direction labelled)	134
7.2	Top row : Images from a DW-MRI dataset for gradient directions 1 to 3. Subsequent Rows : Images rebuilt using, on each row, 1,2,3,4 components respectively.	139
7.3	(Subject 175) Demonstration of apparent improved pixel resolution and feature demarcation with registration method in Fractional Anisotropy (FA) and Principal Diffusion Direction Images. Note increased demarcation of features in the anterior brain, with increasing improvement using affine, fluid and PPCR registration protocols respectively. Also note a decrease in spurious features in the ventricles with registration method. . . .	140
7.4	(Subject 175) Demonstration of apparent improved pixel resolution and feature demarcation with registration method in Fractional Anisotropy (FA) and Principal Diffusion Direction Images. Note improved resolution of features throughout this brain slice with increasing registration algorithm complexity. Serious motion artefacts associated with this particular patient have been removed most successfully using PPCR, making the dataset suitable for further analysis.	141
7.5	(Subject 52) Demonstration of apparent improved pixel resolution and feature demarcation with registration method in Fractional Anisotropy (FA) and Principal Diffusion Direction Images. Fractional Anisotropy Images a to d : note disruption of corpus colosum using simple fluid registration, but preservation under PPCR. This disruption is likely to be due to erroneous through plane registration of the corpus colosum edges in the diffusion direction images, e to h . Note also reduction in noise in ventricles under PPCR in comparison to other methods. Definition of lateral brain features is enhanced in both non-rigid registration algorithms.	142
7.6	Analysis of variation in Fractional Anisotropy value for a leave-one-out analysis of the fifteen diffusion directions. Values are expressed as a percentage of the original FA standard deviation for each registration method <i>see text</i>	143
7.7	Analysis of tensor fitting residual after fitting to 15 diffusion directions. Values are expressed as a percentage of the original residual of the unregistered images for each registration method.	144

7.8	Example Abdominal DWI Slice, comprised of a B0 Image and fifteen diffusion images corresponding to different b values and orthogonal directions prs (labelled). The imaging signal decreases with increasing b -value (Equation 7.3), liver tumours appear as focal increases in the signal.	147
7.9	ADC maps for slices from Patient 1 from Table 1.4: Column 1 before registration, Column 2 after fluid registration using NMI cost function, Column 3 after registration by PPCR.	150
7.10	ADC residual maps for slices from Patient 1 from Table 1.4: Column 1 before registration, Column 2 after fluid registration using NMI cost function, Column 3 after registration by PPCR. Showing some evidence of reduction to model-fit residuals after registration.	151
8.1	Formation of joint image histograms between a pre enhancement images and either a pre or post enhancement image. Force-gradients that reduce the dispersion in the histogram formed using the post enhancement image may result in shrinkage of enhancing features.	153
8.2	Selection of most-likely pixel intensity values for the masked image patch. Masked pixels have a new value calculated by considering their most likely value from the new joint image histogram between pre and masked post-enhancement images (as demonstrated by the green line markings, arrows show the original intensity values, the green dot the replaced intensity value). a Original joint image histogram and b Joint image histogram after masking.	155
8.3	Formation of joint image histograms between two pre enhancement images and a masked post enhancement image. Formation of the image mask as discussed in the text reduces force-gradients that are likely to contribute to shrinkage of enhancing features.	156
8.4	Formation of Normalised Mutual Information force gradients between unaltered post-enhancement and pre-enhancement images. Note the pinching gradients around the heart and aorta in addition to abdominal wall and diaphragm position corrections.	157
8.5	Formation of Normalised Mutual Information force gradients between masked post enhancement image and pre-enhancement image. Note the absence of pinching force gradients in the heart and aorta but appearance of flat regions with visible gradients at the boundaries. The mask used is identical to that shown in Figure 8.3.	158
8.6	Formation of Normalised Mutual Information force gradients between masked and patched post-enhancement image and pre-enhancement image. Note the absence of pinching force gradients and its replacement with local noise in masked regions. Abdominal wall and liver position corrections are maintained. The mask used is identical to that shown in Figure 8.3.	158
8.7	KLAMP mask formation for threshold values between 0 and -2 using step size 0.25 (Pixels are masked if $\log \frac{m_i}{n_i} < threshold$).	159

- 8.8 Segmentation of enhancing features and the liver for analysis of registration performance. Image 1 represents the anchor image to which we register enhancing float images 5 to 7. Image 2 is the additional training image used in the KLAMP method. Motion artefacts between images 1 and 2 are used to generate the training joint image histogram. 160
- 8.9 Example images from Table 8.2 Row 8. **a)** anchor image. **b)** original un-deformed image to compare with **c)** registration without KLAMP and **d)** registration with KLAMP. **Row 2:** Corresponding segmentation images after registration. **Row 3:** Segmentation difference images with gold-standard segmentation for each registration method. 163
- 8.10 Cost-function minimisation spaces of a pre and post enhancement image. Top row: Unmodified images deformed by a force varying in x and y . Bottom row: Corresponding minimisation spaces for pre-enhancement image and KLAMPed post-enhancement image demonstrating a flattening of the cost-function space. 164
- 8.11 Example images from Table 8.2 Row 8. **a)** anchor image. **b)** original un-deformed image to compare with **c)** registration without KLAMP and **d)** registration with KLAMP. **Row 2:** Corresponding segmentation images after registration. **Row 3:** Segmentation difference images with gold-standard segmentation for each registration method. 168

List of Tables

1.1	DCE-MRI Patient 2D Scan Data	25
1.2	DCE-MRI Patient 3D Scan Data	26
1.3	Body Diffusion Patient Data (3D Axial)	26
1.4	Body Diffusion Patient Data (3D Axial)	27
3.1	Typical T1 values for organs at 1.5T and 3T from (de Bazelaire et al., 2004). All values in <i>ms</i>	52
3.2	Modelled Arterial Input Function Parameters for Dual Cosine Input Model (see (Orton et al., 2008) and (Parker et al., 2006) for source.)	57
3.3	Modelled Pharmacokinetic Parameters for given organs from consideration of vascular properties (see text for description) (Parker et al., 2006)	57
3.4	Modelled Pharmacokinetic Parameters for Simulated Pathology (see text for description and compare with Table 3.3). We increase the K^{trans} of the tumour boundary to mimic the expected rate-constant increase due to angiogenesis and add some delay in the onset of enhancement in a filling region.	60
3.5	Algorithm for patching affine transformation into a global elastic deformation (see text for discussion)	63
5.1	PPCR Registration for varying maximum number of internal registration iterations	110
6.1	Blind-evaluation of different registration methods. Each row contains scores comparing two registration types, representing the number of cases each registration was preferred over the other, and the number of cases in which there was no preference. Permutations are between either the unregistered case (NoReg), registration to first image in the dataset (Fluid) and to PPCR (PPCR).	121
6.2	Change in total model fitting residual error for each dataset in Table 1.2 after each registration method for the segmented regions in Figures 6.6 to 6.8.	127
6.3	DCE-MRI Patient 3D Scan Data. NMI Cost function matrix mean values after registration by each method (the percentage improvement relative to the unregistered case is shown in brackets).	132
7.1	Total ADC Residual Before and After Image Registration	149

8.1	Segmentation statistics for the change in area of heart and aorta and the intersection of the liver after image registration with and without KLAMP. Columns 7 and 8: Absolute Residual Deformation (as percentage of original deformation) after image registration with and without KLAMP. *see Figure 8.9.	162
8.2	Segmentation statistics for the change in area of heart and aorta and the intersection of the liver after image registration with and without internal-KLAMP. Columns 7 and 8: Absolute Residual Deformation (as percentage of original deformation) after image registration with and without internal-KLAMP. *see Figure 8.11	167
C.1	Algorithm for fluid registration (see text for discussion)	182
D.1	Algorithm for KLAMPed fluid registration (see text for discussion)	183

Chapter 1

Introduction

1.1 Introduction

This thesis addresses the problem of the registration of images containing local contrast changes. In particular it outlines the problems of using conventional registration techniques in the presence of contrast enhancement. The thesis develops a solution to the problem of the correction of motion artefacts when local contrast-change is occurring.

The desire for accurate quantification of diagnosis and therapy related to widespread diseases such as cancer has led to the formation of many new imaging techniques allowing opportunities to explore advances in treatment monitoring. The increased sophistication of these new imaging modalities has brought additional challenges to the area of image registration. The use of contrast agents to alter image contrast in areas of interest and the use of diffusion weighted MRI to assess dominant diffusion directions are two examples of imaging techniques that have seen increased use in recent years. The length of the imaging procedure in both these cases leads to the formation of images that may contain both intra and inter image artefacts due to motion of the subject. This thesis will focus on the second kind of artefact, those due to motion between images.

The use of image registration is important when seeking to extract information from multiple images. The images must be in good feature alignment so that equivalent imaging pixels represent the same structures and may be compared or combined. Non-rigid image registration procedures seek to maximise the image similarity as defined by a particular measure. The deformation of one image so that it more closely matches another (as defined by the image similarity measure) becomes an optimisation procedure; the deformation field of the image is iteratively refined toward a maximum of image similarity. The re-alignment of a series of medical images that encode multiple contrast variations as a result of either exogenous contrast agents or intrinsic temporal or directional contrast change is difficult. Popular measures of image similarity used to drive the alignment minimisation procedure do not separate the influence of intensity variation due to image feature motion and intensity variation due to feature enhancement. This is true of many cost-functions including information based image similarity measures such as mutual information. The appearance of new structure may pose problems for image registration when it has no representation in the original image. Changes to the intensity of some parts of an image relative to others are also problematic since this violates the one-to-one intensity matching assumed by

many cost-functions.

The acquisition of many images over time, such as in dynamic contrast enhanced MRI, requires that many images with different contrast be registered to the same coordinate system. Since these images are acquired over a time scale of many minutes, patient motion is likely to be a problem, requiring image registration before further analysis can be carried out. In the case of dynamic contrast enhanced MRI, regions of interest are likely to be areas that are enhancing. The correct registration of these regions is crucial when extracting pharmacokinetic information.

This thesis addresses these issues, beginning by presenting some conditions under which conventional registration fails and proposing a solution in the form of a progressive principal component registration (Melbourne et al., 2007b). The thesis proposes the modification of conventional registration algorithms when considering the registration of large groups of contrast enhanced images. This is distinctly different from some of the previous methods of registration in this area. The algorithm uses a statistical analysis of a series of contrast varying images in order to reduce the influence of contrast-enhancement that would otherwise distort the calculation of the image similarity measures used in image registration. The algorithm is shown to be versatile in that it may be applied to a series of images in which contrast variation is due to either temporal contrast enhancement changes, as in dynamic contrast-enhanced MRI (Melbourne et al., 2007a) or intrinsically in the image selection procedure as in diffusion weighted MRI (Melbourne et al., 2008b).

1.2 Chapter Summary

Chapter Two: Literature Review

This chapter introduces the development of image registration algorithms. The focus is particularly on non-rigid, intensity-based methods such as b-spline and fluid registration algorithms incorporating information theoretic image similarity measures. A history of their application to medical images and their eventual application to dynamic contrast enhanced images is discussed. The failure of current registration methods to properly accommodate contrast-variations is outlined, alongside some recent efforts to address this problem. The chapter also discusses the development of MRI and the importance of dynamic contrast enhanced MRI in oncology, presenting some of the challenges associated with the technique. The development of diffusion weighted MRI is also discussed alongside the analogous problems impacting image registration due to local gradient-influenced contrast variation.

Chapter Three: Creation of Simulated Dynamic Contrast Enhanced MRI Data

This chapter discusses the development of fully simulated DCE-MRI. The chapter includes three main areas: the development of a global elastic force model of breathing deformation, incorporating a modification to allow certain regions to deform rigidly; the development of a model of contrast-enhancement processes of both major organs and of pathology using recent work on hepatic contrast enhancement; the importance of the influence of contrast-agent dose and MR signal generation parameters. Part of this work was presented at MICCAI 2008 in (Melbourne et al., 2008a).

Chapter Four: Cost-Functions and Contrast Enhancement

This chapter presents a discussion of current cost-functions and their limitations when contrast-enhancement is present; motivating the need for either contrast-enhancement invariant cost-functions or a method to allow conventional cost-functions to be used. The novel Cost-Function Matrix Mean (CCFM) method for analysing registration performance for groups of images is introduced alongside a method of visualising the potential minimisation space of a cost-function under particular conditions of contrast-enhancement.

Chapter Five: Progressive Principal Component Registration (PPCR)

The use of principal components analysis with medical images is discussed, alongside the difficulty of extracting physiological information from principal components. The use of principal components analysis during an iterative registration procedure is developed, resulting in the Progressive Principal Component Registration algorithm published in (Melbourne et al., 2007b) (patented: see Image Registration Method PCT/GB2008/001520, Filed on 2 May 2008). The conditions under which PPCR will provide an advantage are also discussed in this chapter. The simulated abdominal dynamic contrast enhanced MRI developed in Chapter 3 and the PPCR algorithm discussed are used to investigate the performance of image registration under varying motion and enhancement characteristics. Inspection of the changes to joint-entropy as a function of motion artefact and contrast enhancement are used to infer registration performance. It is shown that for contrast enhanced data, PPCR provides an advantage by allowing conventional cost-functions to be minimised (or maximised) in cases where minimisation is not necessarily possible using conventional post-enhancement to pre-enhancement image registration. Part of this work is the basis of a presentation at MICCAI 2008, (Melbourne et al., 2008a).

Chapter Six: Registration of Breath-hold Dynamic Contrast Enhanced MRI

The Progressive Principal Component Registration algorithm developed in Chapter Four is now applied to real data. The algorithm is applied to both 2D and 3D dynamic contrast enhanced MRI datasets acquired under repeated end-exhale breath-hold. The performance of the 2D registration is analysed by expert visual assessment, by intensity-time curve fitting (published as part of (Melbourne et al., 2007b) and at ISMRM 2007 (Melbourne et al., 2007a)) and by the Cost-Function Matrix Mean. The 3D data are analysed using software developed by the Institute of Cancer Research (MRI Workbench (d'Arcy et al., 2006)) for the extraction of pharmacokinetic parameters. An assessment of pharmacokinetic model-fit residuals both before and after registration reveals an improvement using PPCR compared to conventional image registration (submitted to ISMRM 2009).

Chapter Seven: Registration of Diffusion Weighted MRI

The Progressive Principal Component Registration algorithm developed in Chapter 4 is now applied to a different application. Analogous to dynamic contrast enhanced MRI, diffusion weighted MRI acquires many images analysing diffusion along different directions. The registration of contrast variations between diffusion directions presents the same problems to conventional registration as found in DCE-MRI. Local contrast changes due to diffusion gradient direction invalidate the assumptions of registration

cost-functions. The PPCR algorithm provides a way of incorporating the direction dependent contrast-variations and allowing improved registration performance. The method is applied to 3D datasets and the improvement in registration is analysed using visual image inspection, inspection of fractional anisotropy variability under a leave-one-out analysis and inspection of tensor fit residuals. Part of this work was presented at ISMRM 2008 (Melbourne et al., 2008b). Preliminary work on the registration of diffusion-weighted MRI under varying b -value of the liver is also presented (submitted to ISMRM 2009).

Chapter Eight: Kullback Leibler Assisted Image Matching and Patching (KLAMP)

This chapter discusses the development of a novel method of directly influencing the formation of cost-function gradients during image registration in order to reduce artefacts due to contrast-enhancement. Analysis of the Kullback-Leibler divergence between joint image histograms in which both contain motion, but only one contains contrast-enhancement, allows the removal of contrast-enhancement by image matching and patching. The method is embedded into a fluid registration algorithm. The resulting deformations can be analysed using simulated data, analysis of pre and post registration segmentation and cost-function gradient analysis (Part of this work has been submitted to ISMRM 2009).

1.3 Magnetic Resonance Data used in this Thesis

This section briefly describes the data used in this thesis, consisting of both 2D and 3D liver DCE-MRI studies and diffusion weighted studies of both the brain and liver.

Institute of Cancer Research *Livdt* Study

This study consists of multiple abdominal (liver) dynamic contrast enhanced MRI scans for use in a clinical study. Seven patients are considered and repeat scans are made after a given period of time as shown in Table 1.1. The majority of datasets are a coronal orientation. Each dataset consists of three spatially separated slices anterior-posterior and hence we only consider these data to be 2D in further analysis. Data are acquired with a TR of $11ms$, TE of $4.7ms$ with a flip angle of $\alpha = 3^\circ$. Images are acquired in $3 \times 2s$ with a $7s$ breathing interval. Approximately 40 frames were acquired for each dataset. The Gadolinium based contrast agent, Magnevist, is injected after the 5th image acquisition at $3.5mls^{-1}$. Scans take approximately 9 minutes to perform; misalignment between scans represents a measure of the consistency of the depth of the breath-holds with any additional motion due to the abdominal walls and nearby organs. The length of time between image acquisitions means that it is unlikely that there will be any periodic motion in the sequence.

Table 1.1: DCE-MRI Patient 2D Scan Data

Patient	Number of Scans	Follow Up (Days)	Follow Up (Days)	Follow Up (Days)	Follow Up (Days)	View
1	4	+2	+7	+36		coronal
2	4	+2	+9	+44		sagittal-oblique
3	4	+54	+89	+112		coronal
4	3	+7	+33			coronal
5	2	+2				coronal
6	4	+2	+16	+44	+72	coronal
7	4	+2	+9	+37		coronal

Institute of Cancer Research *Neuro-endocrine* Study

This study consists of six patients with full 3D abdominal (liver) datasets with either 20 or 40 timepoints. This data was acquired by the Institute of Cancer Research on a Siemen's Avanto 1.5T MRI scanner. These datasets are at a temporal resolution of 12s consisting of a 6s held-breath volume acquisition and a further 6s breathing interval. The acquisition TR is $4ms$ with a flip angle of $\alpha = 24^\circ$, a further low flip angle image ($\alpha = 2^\circ$) is also acquired for use in $T1$ estimation. Again these data are taken under repeated breath-hold and particularly evident in these datasets is timing of the acquisition to show the bolus arrival in the heart (Table 1.2). For Patients 5 and 6, also included are the results of the MRI scanner manufacturer's in-built proprietary registration algorithm.

Table 1.2: DCE-MRI Patient 3D Scan Data

Patient	Number of Scans	Volume	Timepoints	View
1	1	256x256x20	20	coronal
2	1	256x256x20	20	coronal
3	1	256x256x20	20	coronal
4	1	256x256x12	40	coronal
5	1	256x256x12	40	coronal
6	1	256x256x12	40	coronal

IXI Brain Data

This series of data consists of 12 volunteer studies over a range of ages. Each volunteer dataset contains 15 diffusion directions acquired with diffusion b -value of $1000s.mm^{-2}$ and a b_0 volume, with volumes of $128 \times 128 \times 64$ pixels from an axial perspective. The datasets are part of the larger IXI dataset which as of writing, is still available here: fantal.doc.ic.ac.uk. The fifteen normalised gradient directions are shown in Table 1.3.

Table 1.3: Body Diffusion Patient Data (3D Axial)

x	y	z
1	0	0
0	1	0
0	0	-1
-0.18	-0.11	0.98
-0.06	0.38	0.92
0.71	0.05	0.70
0.62	-0.44	0.65
0.24	0.78	0.57
-0.26	-0.62	0.74
-0.82	0.17	0.55
-0.84	0.53	0.11
-0.26	0.96	0.14
0	0.97	-0.25
0.75	0.67	-0.02
0.97	0.23	-0.02

Institute of Cancer Research Body DWI Data

These data consist of two datasets of three orthogonal diffusion directions for use in abdominal (liver) oncology taken at multiple b -values including $[0, 50, 100, 150, 250, 500, 750]s.mm^{-2}$ (three each for

each orthogonal direction), see Table 1.4.

Table 1.4: Body Diffusion Patient Data (3D Axial)

Patient	Size	Number of Images	b-Values					
				(x3)	(x3)	(x3)	(x3)	(x3)
1	128x128x18	16	0	50	100	250	500	750
2	256x256x12	7	0	150	500			

Chapter 2

Literature Review

2.1 Images

Within this thesis we consider purely medical images, and in particular those produced using Magnetic Resonance Imaging. Particular features in an MRI image can be enhanced by intrinsically altering scanner parameters or by adding exogenous contrast enhancement or highlighting the freedom of the water to diffuse along a particular direction (as used in Diffusion Weighted MRI).

If we take two images, we often expect there to be changes between them. For example, if two images are taken of a subject breathing in and then breathing out, organs such as the heart and the liver are in different places. If images are taken at different times, we might also expect things to change: tumours growing or shrinking; the heart in systole or diastole; the brain changing in Alzheimer's. The process of finding the spatial alignment between two images is known as *image registration* and describes how well we can write one image in the co-ordinate system of a second (in this case the images are said to be *registered*). More simply, it describes the changes you have to make to one image to produce a second. Image registration can be defined between two images of the same type (mono-modal) or between images of different types (multi-modal) such as the registration of an MR image to an X-ray CT image. Image registration is, in general, a mathematically ill-defined problem. The algorithms we shall see in the subsequent sections are, despite their mathematical complexity, quite simple in the behaviour they can describe: they are remarkably good when two images contain the same features, but if objects between images appear, disappear or change intensity, they often struggle to find an alignment; if registration is defined as describing the features in one image in terms of another, if features have moved out of the image, the registration will not be well-defined. Throughout this thesis we consider the registration of multiple images (inter-image registration), assuming there has been little or no motion during the acquisition of the individual image. Intra-image motion can also be corrected using image registration methods but these processes are beyond the direct scope of this thesis.

2.2 Image Registration

This section will introduce a development of image registration, particularly the cost-functions and the deformation techniques used to maximise those cost functions, presented in a pedagogical fashion. The use of the word *registration* in this thesis should briefly be discussed as it will be interchangeably pre-

sented as a verb (*to register*) to describe the process of aligning two images and as an attributive adjective describing the state of two images irrespective of any image processing (*to be in registration*). For instance, we improve the registration of two images by using a registration tool and the word for this process is also (image) registration. This phrasing is not ideal since it is purely context dependent. Throughout, we will consider the registration of two images (which may be considered volumes if required). We seek to deform one image, which we will call the **float** image (sometimes written as a source image in the literature) so that it resembles as closely as possible (depending on our requirements of similarity) an **anchor** image (sometimes written as a target image in the literature). Conceptually, the anchor image is *fixed* and the float image deforms until it matches the anchor. For shorthand and in equations we represent the **float** image by \mathbf{F} and its individual pixel values by \mathbf{F}_{ij} (for two dimensions) and the **anchor** image by \mathbf{A} and its individual pixel values by \mathbf{A}_{ij} . The fundamental registration equation is provided in Equation 2.1 where we maximise (or more generally extremise) an image similarity measure (*cost*) between *anchor* and *float* image subject to deforming the float image in space \mathbf{r} by a deformation $T(\mathbf{r})$.

$$\max[\text{cost}(\mathbf{A}(\mathbf{r}), \mathbf{F}(T(\mathbf{r})))] \quad (2.1)$$

2.2.1 Cost Functions

The choice of cost function is of importance to the final registration result, selection of a suitable cost-function is crucial to the success of the registration algorithm. It is important to choose a similarity measure that is best-suited to the images that are being registered. For this reason a pedagogical development of cost functions is presented here. An example movie of different cost-function values with (translational) displacement is included on the supplementary CD (see Appendix E).

Mono-Modal Images

Often the most basic cost function presented is the Sum of Squared Differences between anchor and float image intensities. This is appropriate when considering images whose intensity profiles differ by Gaussian noise only; the intensities in the anchor image are expected to be identical in the float image, with the exception of Gaussian noise. Its use in MRI is often limited by fluctuating contrast variations between different MRI images as a result of the large number of tunable parameters. The cost function, C , can clearly be seen to arise from the standard Gaussian distribution for mean μ and standard deviation σ (Equation 2.2 for each of N datapoints x_i) where the product of the different intensity distributions can be written as the minimisation of a sum of the exponent (Equation 2.3 and 2.4). The method is often used in the testing of new algorithms ((Christensen et al., 1996), (Cahill et al., 2007b)).

$$C = \prod_{i=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(x_i - \mu)^2}{2\sigma^2}\right] \quad (2.2)$$

$$C = \frac{1}{(2\pi\sigma^2)^{N/2}} \exp\left[-\frac{1}{2\sigma^2} \sum_{i=1}^N (x_i - \mu)^2\right] \quad (2.3)$$

$$SSD = \sum_{i=1}^N (x_i - \mu)^2 \quad (2.4)$$

Multi-Modal Images

The assumption of images differing only by Gaussian noise often cannot be made. It may then be appropriate to consider that the intensities differ in a linear way; that the intensities in the float image can be described by a global scalar multiplication of intensities in the anchor image. If we consider our float \mathbf{F} and anchor \mathbf{A} images as vectors of length $n \times m$, our image similarity measure is simply (with the proviso that our images have the mean intensity subtracted) the angle between these two vectors (Equation 2.5).

$$CrossCorrelation = \frac{\mathbf{A} \cdot \mathbf{F}}{\|\mathbf{A}\| \|\mathbf{F}\|} \quad (2.5)$$

For perfectly aligned images, all pixels will have the same values in both images and therefore the value of the Cross Correlation (CC) will be 1. If the float image was the photo-negative of the anchor image, the value of the Cross Correlation would be -1 . Work by Hermosillo provides details on the implementation of cross-correlation and information theoretic cost functions (Hermosillo, 2002).

Information Theory Based Cost Functions

More generally we might suppose that there is a relationship between pixel intensities that is not reliant on any presumed intensity function. Since it is particularly illustrative, first we will derive the Joint Entropy cost function (JE), as first proposed by Hill (Hill et al., 1994). If we take a given intensity in the float image, we then look at how many times it corresponds to all other available intensities in the anchor image (for this reason we consider intensity discretised images, for example with 256 possible intensities). It is possible to imagine that some of the results will represent a true alignment of pixels and some the result of unregistered, misaligned pixels. By doing this for every intensity value in the float image, we build up a joint image histogram. We now propose that the histogram counts represent the probability of a particular pixel combination occurring. We can now see that for a good alignment there will be very high numbers of pixels corresponding to a one-to-one relationship in the joint histogram, meaning that a given pixel intensity in the float image is *highly likely* to correspond to one pixel intensity in the anchor image. Poor alignment would see a one-to-many spread of a particular intensity value in the float image to values in the anchor image (an important aside, and limitation of this method, is that local intensity changes other than those due to motion are likely to lead to a valid spread in the joint histogram counts). The extent of this dispersion can be summarised by the *entropy* of the joint image histogram (Equation 2.6) for the distributions of probabilities p_{ij} in the histogram. The entropy of the joint image histogram for anchor and float images will be denoted H_{AF} . As discussed, a smaller value of joint entropy should correspond to better image alignment.

$$JointEntropy = - \sum_i \sum_j p_{ij} \log p_{ij} \quad (2.6)$$

The probability distribution used above for the joint histogram may also be applied to a single image histogram, where the histogram counts represent the number of occurrences of a particular intensity within a given image (either the anchor or the float). As above, we can convert these histogram counts to probabilities and find the entropy of a particular image (Equation 2.7). These entropies are often referred to as *marginal entropies* and will be denoted as H_A , for the marginal anchor entropy and H_F , for the

marginal float entropy. Images having equal numbers of pixels at each intensity will contain the most information.

$$\text{MarginalEntropy} = - \sum_i p_i \log p_i \quad (2.7)$$

A further potential cost function is the Kullbach Leibler Distance (KLD). This is a measure of the 'distance' between one probability distribution, which can be written as p_i , and another, written as q_i (the KLD value will be zero for $p_i = q_i \forall i$). For images these distributions are calculated as discussed previously. The distance between the probability distributions of an anchor and a float can be written as Equation 2.8.

$$\text{KLDivergence} = \sum_i p_i \log \frac{p_i}{q_i} \quad (2.8)$$

$$\text{JSDivergence} = \frac{1}{2} \text{KLD}(P \| M) + \frac{1}{2} \text{KLD}(Q \| M) \quad (2.9)$$

$$M = \frac{1}{2}(P + Q) \quad (2.10)$$

What is important with regards to image registration is, if this was used as an image similarity measure, that it is not symmetric. The Kullbach-Leibner Distance between the anchor and the float is *not the same* as the distance between the float and the anchor. For this reason it is not a distance and should be referred to as a *divergence*. Also we have no reason to suppose that the forward distance is more appropriate for aligning images than the backwards distance. Some authors seek to symmetrise the measure by considering both the forward and backward distances as in the case of the Jensen-Shannon Divergence (Equations 2.9 and 2.10). Chiang (Chiang et al., 2008) uses the symmetrised KL-Divergence in the registration of Diffusion Tensor MRI to re-orientate tensors according to the Gaussian Probability Density Functions (PDF) of the diffusion tensors. In this case the measure is applied between two images with the same expected PDFs, hence the measure may be considered appropriate.

The Mutual Information Cost Function (MI) combines information from both the entropy of the joint image histogram, H_{AF} , and the individual entropies (marginal entropies H_A and H_F) of the separate images (Equation 2.11) ((Viola & Wells, 1997), (Pluim et al., 2000), (Pluim et al., 2003)). The advantage over joint entropy is the inclusion of the marginal entropies. Not only do we seek to minimise the joint entropy, we seek to maintain the amount of entropy (information) in the individual anchor and float images. This has the effect of counteracting a situation where the joint entropy falls by tending toward a situation which reduces the spread in the individual intensity distributions. This would reduce the joint entropy and could be caused if the image overlap begins to decrease, causing a large number of pixels to align to a background intensity value.

$$MI = H_A + H_F - H_{AF} \quad (2.11)$$

A related formulation of mutual information is to divide the sum of marginal entropies by the joint entropy (Equation 2.11) to give the Symmetric Uncertainty Coefficient, otherwise known as Normalised Mutual Information (NMI) (Studholme et al., 1999). Mutual Information does not completely solve the overlap problem. This modification further removes the problem of images being driven away from one

another by normalising to the amount of information in the joint image histogram and the amount of overlap (Hill et al., 2001). The value of NMI is normalised in the sense that the maximum and minimum MI values are dependent on the number of intensity bins and the distribution of the image intensities. Although NMI was not developed in this way by Studholme *et al*, the expression for NMI is equivalent to the Symmetric Uncertainty Coefficient (the average of how well A can be used to predict F and vice versa). The expression must be made symmetric because the uncertainty coefficients themselves are not: $U(A|F) \neq U(F|A)$, so one possible symmetry is to weight by the marginal entropies (Equation 2.13) giving Equation 2.14.

$$U(A|F) = \frac{H_A + H_F - H_{AF}}{H_A} \quad (2.12)$$

$$SUC = \frac{H_A}{H_A + H_F} U(A|F) + \frac{H_F}{H_A + H_F} U(F|A) \quad (2.13)$$

$$SUC = 2\left(1 - \frac{H_{AF}}{H_A + H_F}\right) \quad (2.14)$$

$$NMI = \frac{H_A + H_F}{H_{AF}} \quad (2.15)$$

2.2.2 Transformation Models

In addition to an appropriate measure of image similarity selected from the previous section, in order to maximise the cost-function we need to deform our image in a well-defined way so that the information in the image is not degraded. Methods for maintaining constancy of the deformed image are now outlined in this section.

Before an in-depth discussion of non-rigid deformation models, it is appropriate to mention transformation models with much lower degrees of freedom. A rigid transformation is described by only three translations and three rotations, one along each axis, and the entire coordinate system of the image is transformed accordingly. Extending the rigid transformation to include scaling and shearing we include 6 further degrees of freedom and the transformation is now affine. Image transformations using an affine model keep parallel lines parallel. Further degrees of freedom may be added by including projective transformations or by allowing the transformation to be described as a polynomial function.

Optical Flow Registration

An early image matching algorithm was proposed by Horn and Schunck (Horn & Schunck, 1981). Due to the ease of coding and the simple conceptual nature of the result, the method is still to be found in many publications ((Hayton et al., 1997), (Alvarez, 2000), (Martel et al., 2007)) and is the basis for a further registration method known as the Daemons Algorithm ((Thirion, 1998)) and subsequent work. The algorithm is based on the assumption of moving points in the image having constant image intensity, hence the cost function in this case is implicit to the transformation model, and therefore, for points in an image, \mathbf{F} , we have Equation 2.16 and Equation 2.17. Equation 2.17 comes about because we are following the trajectory of a *particular piece* of intensity, not considering the intensity change at a fixed point, as per the assumption.

$$\frac{d\mathbf{F}}{dt} = 0 \quad (2.16)$$

$$\frac{\delta \mathbf{F}}{\delta x} \frac{\delta x}{\delta t} + \frac{\delta \mathbf{F}}{\delta y} \frac{\delta y}{\delta t} + \frac{\delta \mathbf{F}}{\delta z} \frac{\delta z}{\delta t} + \frac{\delta \mathbf{F}}{\delta t} = 0 \quad (2.17)$$

It can be seen from Equation 2.17 that we have a linear system in the component velocities ($\frac{\delta y}{\delta t}$ etc.). It is not possible to determine a velocity along a brightness contour so we must include a smoothness constraint of the form, $\nabla^2 \mathbf{v}$, where \mathbf{v} is representative of the velocities present in Equation 2.17. Including this constraint allows us to produce a smooth deformation that will restore intensity discrepancies due to motion. The complexity of the required algorithm is relatively low, particularly if the smoothness constraint is approximated using the difference of a point in the velocity field from its adjacent neighbours and included directly in finding the solution to Equation 2.17, also making the algorithm fast. For this reason, optical flow algorithms are often used when testing modifications to the registration paradigm, such as in Hayton (Hayton et al., 1997) who applied an optical flow algorithm to registration of DCE-MRI (using a model-fitting cost function), and Martel (Martel et al., 2007) who applied the optical flow algorithm to Dynamic Contrast Enhanced MRI Data, including a piecewise linear intensity change constraint due to Gennert (Gennert & Negahdaripour, 1987). Vercauteren (Vercauteren et al., 2007) used Efficient Second-Order Minimisation to analyse an optical flow algorithm in the forms of a daemons algorithm as preliminary work towards including modifications to the algorithm to ensure diffeomorphic transformations. A diffeomorphic transformation is one that is invertible - or more strictly that the inverted deformation is also differentiable. This may be preferable since the registration solution is ill-posed, we at least have a solution that has workable mathematical properties.

An extension of the optical flow method was proposed by Thirion in 1998 (Thirion, 1998) and has been used several times since ((Pennec et al., 1999), (Stefanescu et al., 2004), (Vercauteren et al., 2007)). The claim is that the image matching is done with a rather tenuous analogy to Maxwell's demon [sic]. This is perhaps an unexpected consequence of thermodynamic and statistical physics.¹

Two uses of Thirion's daemonic effectors are presented in (Thirion, 1998): the author first applies effectors to an object boundary in the anchor image (for full non-rigid registration: effectors would be placed on a regular grid throughout the anchor image) and the ability of a corresponding float object to diffusively pass through this effector boundary according to some measure of increasing image similarity; the second considers the effectors as a *snake-like* contour on an object in the float image that can then be deformed to match a structure in the anchor image. The resulting method is extremely versatile but both these applications may require some object segmentation in both the anchor and float images.

¹Maxwell's demon was a concept devised to break the Second Law of Thermodynamics, stating that entropy always increases. For two adjacent boxes: one with some particles in and the other empty, the entropy (which can be imagined as disorder) of the system is quite low. If the partition between the boxes is removed, the particles will spread between the two boxes, therefore increasing the entropy. However, if a demon (and we will shortly see it is a *demon*) was to try to separate the particles by using the partition to only let particles into one side of the box, and never out, we would have a violation of the Second Law of Thermodynamics. The solution to this problem is along the lines that the demon is part of the system and he must receive information and do work on this information in order to separate the particles, hence for the *total* system the entropy must increase. In this sense, Maxwell's demon is certainly of the malevolent Medieval variety, seeking to disturb the laws of nature (and fortunately not succeeding). The demons of Thirion are not in the same sense as those of Maxwell, they are more *effectors* in the Classical sense, and act for neither good nor evil. In modern terminology this separation between Classical and Medieval has led to the terms for a malevolent *demon* and an effecting *daemon*, hence Thirion employs daemons, unrelated to Maxwell's experiment.

The method relies on a suitable definition of the daemons and their application. In this sense, simple brightness change could be used and the result is essentially an optical flow algorithm (Pennec et al., 1999).

B-Spline Registration

A different approach was proposed by Rueckert (Rueckert et al., 1999). Using a grid of regularly spaced *control points* across the float image, it is possible to move the control points and calculate the intermediate deformations according to a fitted spline. Basic splines (b-splines) are chosen which allow any deformation to be locally contained (mathematically they have limited support), this makes them very efficient to calculate as we need to consider only the few nearest neighbours to a control point. This method is also amenable to a hierarchical multi-resolution registration from coarse-to-fine scales.

The application of splines applied to biomechanical systems was first suggested by Bookstein (Bookstein, 1989), but Bookstein used them to produce smooth deformation fields under an applied force. The local deformation in the field T at a point x, y, z is given by Equation 2.18, relative to the nearest associated control point at $\Phi(i, j, k)$ and with respect to the distance of that point from that control point u, v, w . The appropriate spline is built up from a combination of cubic basis functions as shown in Equation 2.19. Additional regularisation is required to smooth the b-spline deformation and a bending energy regularisation in the spirit of the spline deformation is often applied.

$$T(x, y, z) = \sum_{l=0}^3 \sum_{m=0}^3 \sum_{n=0}^3 B_l(u) B_m(v) B_n(w) \Phi(i + l, j + m, k + n) \quad (2.18)$$

$$B_0 = \frac{1}{6}(1 - u)^3, B_1 = \frac{1}{6}(3u^3 - 6u^2 + 4), B_2 = \frac{1}{6}(-3u^3 + 3u^2 + 3u + 1), B_3 = \frac{1}{6}u^3 \quad (2.19)$$

The method has been widely used and the algorithm is often applied in the analysis of organ motion and deformation. McLeish applied the method to correct for motion of the heart during respiration, determining the extent of motion along each biological axis (McLeish et al., 2002). Tanner (Tanner et al., 2000), (Tanner et al., 2002) used the method in the registration of contrast-enhanced breast MRI. Potential mis-registration of enhancing features motivated the inclusion of a volume preserving constraint. Work on the validation of the B-spline method has also been carried out by analysing the B-spline registration results against a gold standard deformation generated from a biomechanical breast model (Schnabel et al., 2003). Similar work was carried out by Rohlfing (Rohlfing et al., 2003), analysing the Jacobian determinant of the deformation field and preventing any unrealistic volume change. Rohlfing also applied the B-spline method to analyse liver motion during the respiratory cycle (Rohlfing et al., 2004).

Fluid Registration

The modelling of image deformation as a fluid was first proposed by Christensen (Christensen et al., 1996). This method allows a more sophisticated regularisation than that used by a simple optical flow algorithm by coupling the component directions of the deformation field. For completeness, the full fluid equation is described in Equation 2.20 for a given pressure, \mathbf{P} , density, ρ , in a potential, ψ , with viscosity parameters μ and λ . The velocity of the flow, \mathbf{v} , is used to update the image transformation

over a given time-step. A well explained derivation can be found in the Feynman Lectures (Feynman et al., 1998) and is also included for the interested reader in Appendix A. Here we must add a force term associated with the similarity between our float and anchor images, $\mathbf{F}_B(\mathbf{A}, \mathbf{F})$. The range of phenomena described by the fluid equation is vast, hence a fluid is chosen that is isotropic, slow moving and viscous; physically this corresponds to a low Reynolds number (the number which represents the ratio of inertial to viscous forces). Since our viscous forces dominate and our inertial terms (those containing ρ) are correspondingly small (and we do not expect there to be pressure variations across the image), our equation reduces to Equation 2.21 for an isotropic medium. For a driving force calculated from our float image, \mathbf{F} , and anchor image, \mathbf{A} , our fluid will flow with velocity, \mathbf{v} , dragging the image it represents with it, which in turn alters the driving force. If we iterate forward in time, recalculating the new driving force after each time step, we hope to reach a situation in which the fluid stops flowing, corresponding to an image match: hence the force is zero and therefore the velocity field is zero.^{2 3}

$$\rho [(\mathbf{v} \cdot \nabla)\mathbf{v}] = -\nabla\mathbf{P} - \rho\nabla\psi + \mu\nabla^2\mathbf{v} + (\mu + \lambda)\nabla(\nabla \cdot \mathbf{v}) - \mathbf{F}_B(\mathbf{A}, \mathbf{F}) \quad (2.20)$$

$$\mu\nabla^2\mathbf{v} + (\mu + \lambda)\nabla(\nabla \cdot \mathbf{v}) = \mathbf{F}_B(\mathbf{A}, \mathbf{F}) \quad (2.21)$$

The solution of Equation 2.21 is a time consuming step. Christensen (Christensen et al., 1996) finds the solution using Successive Over Relaxation. By implementing the Successive Over Relaxation in a multi-grid solution, Crum (Crum et al., 2005) was successfully able to solve the fluid equation rapidly by propagating the solution at different resolutions between scales. This technique is well described in Numerical Recipes (Press et al., 2007). The solution may also be found using Fourier methods and this has been demonstrated by Cahill ((Cahill et al., 2007b), (Cahill et al., 2007a)) by carefully re-writing the fluid equation as a product of itself and its adjoint and recognising that the resulting solution of $\nabla^4 v$ can be expressed as a Sine, Cosine or Fourier Transform (depending on required boundary conditions) of the required velocity field. This solution is particularly desirable, since it is not only fast (requiring only a few Fast Fourier Transforms) but can be coded succinctly. Unfortunately Cahill *et al* do not provide a comparison of the the speed benefits of this solution with previous work on different size images, perhaps because of the difficulty of obtaining equivalent code.

Fluid registration has found greatest application in the brain. Crum (Crum et al., 2001) applied fluid registration to monitor hippocampal volume change in Alzheimers patients, suggesting the automated (and therefore labour saving) method was more accurate than manual-segmentation by 2%. d'Agostino (D'Agostino et al., 2003) applied the method for the analysis of multi-modal brain image registration using the Mutual Information cost function. Hecke (Hecke et al., 2007) used fluid registration to align Diffusion Tensor images of the brain, finding the alignment superior to affine image registration.

²When the velocity field \mathbf{v} in Equation 2.21 is substituted for the displacement field \mathbf{u} it becomes the solution of the linear elastic equation (it is also the most general solution of a function with only second-order derivatives). Registration using a linear elastic regularisation has been attempted ((Bajcsy & Kovacic, 1989), (Alexander et al., 1999)), but is limited since the overall displacement penalty term will grow with increasing displacement, putting a limit on registration success, which is not necessarily the case with a fluid deformation.

³When the viscosity coefficients have the condition $\mu = -\lambda$ the coupling term is removed and the equation becomes a (purely) diffusion equation, a generalisation of the daemon-based registration algorithm.

2.2.3 Recent Developments in Image Registration

A large number of articles have recently been published concerning developments in image registration. In particular there is a desire for diffeomorphic registration; deformation fields that can be inverted and therefore must provide a one-to-one mapping of one image to another (or from one space to another). In some circumstances, diffeomorphic registration is required from a biomechanical perspective, but it can only be true if the entirety of one object in one coordinate space is present in the second space. For monitoring the appearance of new features and possible changes in field-of-view the requirement provides little advantage. However, the desirability of an invertible solution from a mathematical perspective may be useful and this should be the predominant reason for use of a diffeomorphic transformation.

Recent developments have also investigated the inclusion of biomechanical models as transformation models in image registration. These might require a correspondence between the driving force and the deformation that is not really achievable using information theoretic cost-functions. Driving forces are likely to require optimisation over known motion models that describe biomechanically plausible deformations. Currently this type of image alignment is computationally expensive, but a growing interest in the use of graphical processing units for image processing is beginning to make implementation possible.

2.3 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) was developed through the 1970s and 80s from existing Nuclear Magnetic Resonance (NMR) chemical analysis into a full medical imaging technology (McRobbie et al., 2006).

The explanation of MRI is intriguing since it can only be explained theoretically using quantum mechanics, but for all practical purposes the description can be done in the classical sense. The reason for this is that we are dealing with large numbers of quantum objects which can then be considered classically. In MRI it is common to only look for a particular type of nucleus and the signals are very small. A large percentage of the body is made up of hydrogen, either bound as water or bound into compounds such as fat. Therefore the hydrogen atoms in water are a good choice for magnetic resonance. Hydrogen atoms consist of a proton orbited by a single electron: the proton nucleon has a directional intrinsic spin of $+\frac{1}{2}$ and therefore can be aligned (parallel or anti-parallel) with a magnetic field (as can any nucleus with odd numbers of protons or neutrons). At absolute zero, a sample of hydrogen would align all its spins parallel with that field. Unfortunately at body temperature $310K$ the difference in number of spins parallel to those anti-parallel is about $1 : 10^5$, however it is still possible to record a signal⁴.

Putting a hydrogen sample into a magnetic field splits the nuclei between two states by an energy $\hbar\gamma B$, where γ is the *gyromagnetic* ratio of a particle's charge to its mass $\frac{e}{m_p}$ which can be classically

⁴This number can be calculated from the (classical) Boltzmann Distribution having found the difference in energy between spin up and spin down states relative to body temperature T . $\frac{N_{aligned}}{N_{notaligned}} = \exp\left[\frac{g e \hbar B}{2 m_p k_B T}\right]$ where e is the electron charge, m_p the proton mass, \hbar Planck's constant divided by 2π , k_B the Boltzmann constant, B the field strength and g is the 'g-factor' calculated from quantum chromodynamics and having a value close to 2.

imagined as the relative strength of an orbiting particle's electromagnetic attraction to its desire to move in a straight line due to having mass. This should be modified slightly to account for small quantum effects using the g-factor mentioned in the footnote, $\gamma = \frac{g\mu_B}{\hbar}$.

Despite the small net number of atoms split by the magnetic field, there are sufficiently many that we can use a classical description of the net magnetisation of small regions. The most obvious feature of a MRI scanner is a large static magnetic field, in many systems this will be 1.5T or 3T, and in some cases more depending on application (for comparison, the Earth's magnetic field is between 30-60 μ T). This field runs parallel to a bore in the centre, in which the patient lies, hence the net magnetisation vector of the body is aligned along the bore. It is possible to alter the direction of the magnetisation vector by introducing electromagnetic radiation (in this case radio waves) at a particular frequency. This frequency is chosen so that it matches the natural frequency of the rotating spin of a particular substance in the body (here hydrogen) and is given by $f = \frac{\gamma}{2\pi}B$ Hz which for hydrogen is approximately $f = 42.58 \times 10^6 B$ Hz (a radio-wave) depending on how the hydrogen is bound to its surroundings this is known as chemical-shift. Adding radio-waves at this frequency allows us to alter the direction of the spin, this is a resonance effect, where we match our external force to the intrinsic frequency of a body proportional to $\frac{e}{m_p}$.

For a resonant frequency of $f = 42.58 \times 10^6 B$ with a single value of magnetic field, \mathbf{B} throughout the scanner, we can alter the direction of the net body magnetisation by using a radio-wave frequency that matches the resonant frequency of hydrogen nuclei. If we were to turn off the radio-waves, the net magnetisation vector would relax back to the direction of the large static B-field. The resonance matching frequency is dependent on the strength of the magnetic field, hence if we were to vary the B-field across the image, perturbing it with additional magnetic fields, we could apply a spectrum of radio-waves to match the resonance across the image. If we were to switch off this spectrum, the resulting signal would tell us the strength of the magnetisation at different points in the body, this would be an image. In practice, a large gradient is applied to the B-field so that only a slice of the body can be made to resonate using radio-waves close to a certain frequency (slice selection). A further gradient is applied along a second direction (often the largest body dimension in the plane of the image), so that a spectrum of radio-waves can be detected (frequency encoding). In the third direction we record each signal using another magnetic field to de-phase equivalent frequencies by a small amount (phase encoding). The signal acquisition relies on the intrinsic properties of the magnetised substance whose signal decays with rate constants T_1, T_2, T_2^* . In turn these describe: T_1 (longitudinal relaxation time); the decay of magnetisation of the spins with the 'lattice', or the surrounding environment, representing the loss of spin precession at the Larmor frequency due to spin-lattice relaxation; T_2 (transverse relaxation time); the decay of net magnetisation (by dephasing) of the spins with each other due to local spin-related magnetic field changes and T_2^* ; the observed T_2 decay of magnetisation of the spins including the effects of local field inhomogeneities.

The formation of the MR signal and the subsequent detection requires that the data be acquired in spatial frequency space (denoted k-space). The encoding of spatial position in the frequency and phase of the MRI signal requires that the complete signal map be built up in k-space before conversion to image

space. Collection of the signal in k-space may be done in any trajectory subject to hardware limitations; the ease at which the phase and frequency encoding gradients can be altered. Different acquisition schemes can be chosen to traverse k-space in a simple line-wise fashion from maximum to minimum phase and frequency encoding, with a radial line-wise profile or in a more complex fashion if there are advantages to the order of line acquisition. The conversion from spatial frequency to image space may be done by Fourier transform on completion of the acquisition. As a result the collection of the signal and the method of traversing k-space is an important factor in image speed and resolution. Motion artefacts may corrupt the acquisition of separate parts (or shots) of k-space leading to artefacts such as ghosting. A large body of work has been developed focusing on the correction of these intra-image motion artefacts. If motion occurs during acquisition of a single magnetic resonance image artefacts such as ghosting will occur. A method to autofocus individual images by correcting for phase-shifts due to simple motions was developed by Atkinson *et al* (Atkinson et al., 1997) and generalised by Batchelor *et al* (Batchelor et al., 2005) so that motion between imaging shots can be corrected arbitrarily. This method was adapted in order to correct for intra-image breathing motion by White *et al* (White et al., 2008) subject to the formation of a patient-specific breathing motion model. The correction of intra-image motion artefacts is not addressed in this thesis, we assume that intra-image motion artefacts are negligible.

From McRobbie (McRobbie et al., 2006), Equation 2.22 describes the signal strength for a particular tissue with intrinsic T_1 and T_2 in a spoiled gradient echo sequence with short echo time, TE and flip angle, α and can be calculated from the Bloch equations (which describe magnetisation changes with time). Gradient echo sequences are typically used in T1-weighted Gd-DTPA imaging as will be discussed in Section 2.3.1. The timing diagram for a spoiled gradient echo sequence is shown in Figure 2.1(top) the gradient strengths of G_{pe} and G_{fe} are stepped to allow both phase and frequency encoding; the sequence length can be shortened if we include spoiler gradients which increase spin dephasing, reducing the transverse magnetisation, as in Figure 2.1(bottom). Equation 2.23 represents the expected signal from a spin echo sequence and Equation 2.24 the expected signal from an inversion recovery sequence which could be used for determining intrinsic T_1 relaxation times (these equations typically assume that the echo time (TE) is much shorter than the repetition time (TR)); Figure 2.2 shows a typical spin echo timing diagram. In all cases the value of S_0 is used to absorb the effects of the signal detection apparatus and any additional tissue pathology effects additional to the intrinsic T_1 , T_2 and T_2^* values.

$$S_{SGE} = \frac{S_0 e^{-\frac{TE}{T_2^*}} (1 - e^{-\frac{TR}{T_1}}) \sin(\alpha)}{1 - \cos(\alpha) e^{-\frac{TR}{T_1}}} \quad (2.22)$$

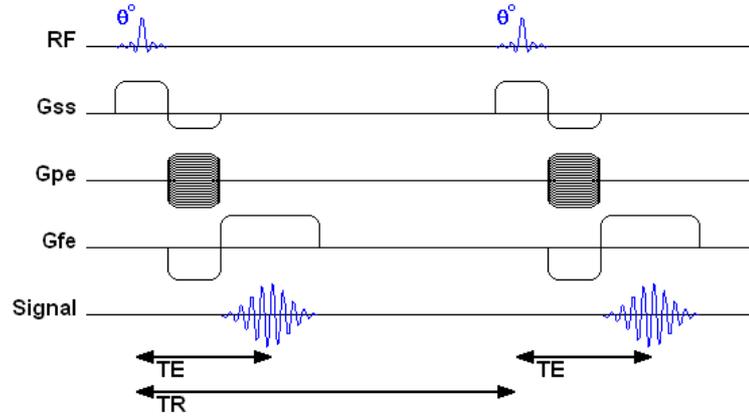
$$S_{SE} = S_0 e^{-\frac{TE}{T_2}} (1 - e^{-\frac{TR}{T_1}}) \quad (2.23)$$

$$S_{IR}^{180-90} = S_0 (1 - 2e^{-\frac{TI}{T_1}} + e^{-\frac{TR}{T_1}}) \quad (2.24)$$

2.3.1 Dynamic Contrast Enhanced MRI

Exogenous contrast agents such as gadolinium-DTPA (Gd-DTPA, gadolinium bound in a non-toxic chelate) increase the MRI signal by interacting with water to shorten its T_1 (and T_2 at large doses $> 1 \text{ mmol kg}^{-1}$). Water molecules passing close to the Gd-DTPA molecule are subject to a local field

Gradient Echo Timing Diagram



Spoiled Gradient Echo Timing Diagram

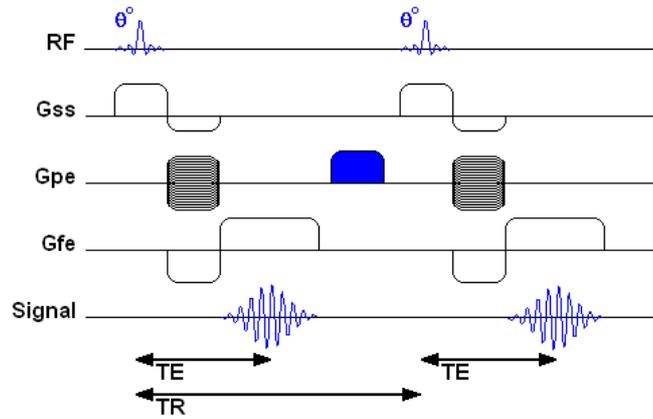


Figure 2.1: Illustrative example timing diagram for a Gradient Echo sequence. **Top**, Standard gradient echo protocol. **Bottom** Spoiled Gradient Echo including a spoiler gradient for more rapid acquisition, as used in Gd-DTPA contrast-enhancement imaging.

inhomogeneity and are more likely to move away from the Larmor frequency, contributing to the observed T_1 shortening (McRobbie et al., 2006). T_2 shortening is also due to increased dephasing due to water molecule interaction with Gd-DTPA. The modification of T_1 assumes a linear modification to the relaxation *rate* and is shown in Equation 2.25 for a given concentration, C , scaled by a substance specific relaxivity r (the method for the observed alteration to T_2 follows equivalent steps). The effect of shortening T_1 is seen to boost the signal for a given TR and flip angle α . This is shown in Equation 2.26 where the T_2^* term from Equation 2.22 is incorporated into S_0 since the effect of the contrast agent on T_2 may be ignored. The reason for this inclusion is that in human tissue T_2 is always far less than T_1 (Table 3.1), hence the effect of the linear correction to $\frac{1}{T_2}$ is normally very small (compare with Equation 2.25) subject to a relatively long TE .

$$T_1'(t) = \left(\frac{1}{T_1} + rC(t) \right)^{-1} \quad (2.25)$$

Spin Echo Timing Diagram

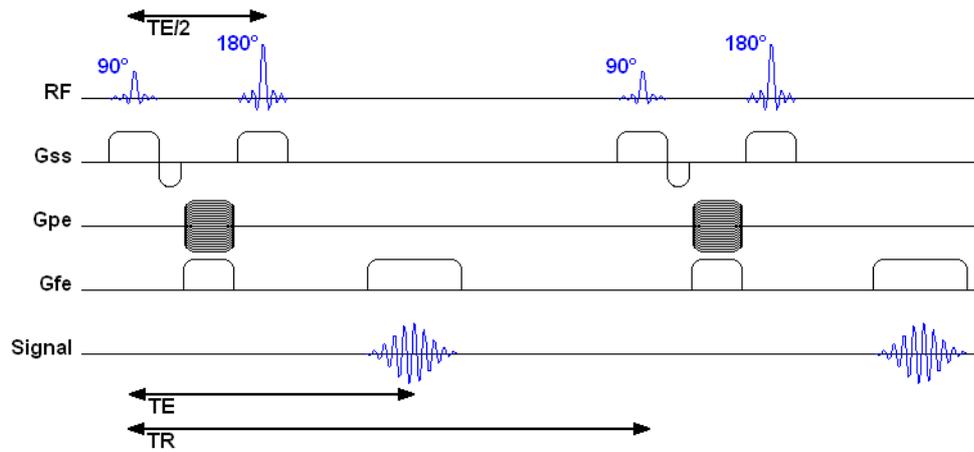


Figure 2.2: Illustrative example timing diagram for a Spin Echo acquisition sequence.

$$S(t) = S_0 \frac{\sin(\alpha)(1 - e^{-\frac{TR}{T_1(t)}})}{(1 - \cos(\alpha)e^{-\frac{TR}{T_1(t)}})} \quad (2.26)$$

Early work on the effects of Gadolinium as a contrast agent is found in (Weinmann et al., 1984) discussing the biological stability of the Gd-DTPA chelate and its effects on T_1 , T_2 at a range of doses. Donahue (Donahue et al., 1994) provide analysis of the relaxivity of Gd-DTPA, applying Equation 2.25 to find the T_1 relaxivities of blood plasma and cardiac tissue in a carpine and anural models. Work by Rinck (Rinck & Muller, 1999) analyses the magnetic field strength dependence of both T_1 and T_2 .

In the case of MRI contrast agents it is the effects of the contrast agent that we observe, rather than the contrast agent itself. This is important when considering that although, with a molecular mass of 500 nucleons, Gd-DTPA is able to leave capillaries and penetrate the extracellular-extravascular space, it cannot find its way inside cells. However, the water molecules it interacts with may cross the cell wall which might influence the observed T_1 values. Iron-oxides (coated in a carbohydrate shell, e.g. Feridex) are also used but since the magnetic field inhomogeneity is much larger than with Gd-DTPA, strongly reduce T_1 , T_2 and T_2^* over a large area even for small doses.

Endogenous contrast enhancement in tumours is also possible, for instance by the BOLD effect (Jiang et al., 2004), measuring blood oxygenation levels as used in functional MRI. This has been demonstrated using T_2^* measurements by both (Baudalet & Gallez, 2002) and (Taylor et al., 2001). The work by Taylor monitored tumour response on breathing carbogen (95% O_2 , 5% CO_2), concluding that it may be used to identify patients suitable for carbogen radiosensitisation pre-treatment. Contrast enhancement data may consist of only pre and post enhancement images which are then subtracted to show enhancing areas or it may be dynamic. Dynamic Contrast Enhancement monitors the progress and distribution of contrast agent through a particular organ by acquiring many images as a function of time. Dynamic information allows much more information to be obtained from the enhancement process as we will see below.

The Physiological Basis of Contrast Enhancement

Contrast agents provide a way of assessing treatments such as anti-angiogenic or anti-vascular therapies in oncology. Growing tumours require a blood-supply, and to achieve this, stimulate the growth of neo-vasculature in their surroundings. This is thought to occur by production of growth factors regulating Vascular Growth (Vascular Endothelial Growth Factor - VEGF) and Vascular Permeability (Vascular Permeability Factor - VPF) (Passe et al., 1997). A tumour's surrounding area will then consist of many tortuous, new and permeable blood vessels. Contrast agent reaching this area will transfer rapidly from blood plasma to extracellular-extravascular space, enhancing the MRI signal in the tumour boundary with a 'ring-shaped' enhancement. The relationship of DCE-MRI to histology has been made by many authors ((Buckley et al., 1999), (Knopp et al., 1999), (Harrer et al., 2004), (Patankar et al., 2005), (Cuenod et al., 2006)). Knopp *et al* analyse differences in enhancement due to tissue type by comparing enhancement with histology. They find there are significantly (with a statistical p-value < 0.001) faster exchange rates of contrast agent between vascular space and extra-cellular, extra-vascular space in malignant tissue compared to benign tissue. The authors suggest that contrast enhancement variations are mainly due to differences in vascular permeability manifest as a high expression of VEGF in histology. The successful application of DCE-MRI is discussed by Choyke (Choyke et al., 2003), providing an overview of practical DCE-MRI and its application in renal, cardiac and osteosarcoma applications. A report by Leach *et al* (Leach et al., 2005) provides recommendations on the required outcomes of DCE-MRI for the analysis of antiangiogenic and antivascular therapies.

Enhancement Curve Modelling

Early work on modelling the enhancement process of DCE-MRI was applied in the brain by Tofts and Kermode (Tofts & Kermode, 1991). They consider a two compartment model of contrast agent transferring between a vascular compartment and an Extracellular-Extravascular compartment (EES).

Distribution of contrast agent between the two compartments is governed by Equation 2.27 where the rate of change of contrast agent concentration in the extra-cellular extra-vascular space (EES), C_e given the fractional volume of the leakage space, v_e , is given by the difference between influx governed by the rate constant, K^{trans} as a function of the local permeability and surface area between compartments and the arterial input bolus, C_p (further discussed in the next section); and subsequent efflux from the EES leakage space. The integral solution of Equation 2.27 is given by Equation 2.28 which may be re-written as the convolution in Equation 2.29. Up to this point we have considered our tissue to consist purely of EES, a better representation is to consider that the total contrast agent contribution comes from a mixture of compartments with relative volumes representing the EES v_e , the blood pool v_p and the intra-cellular space v_i , hence $C_t = v_e C_e + v_p C_p + v_i C_i$ (Tofts, 1997). We assume that contrast agent (or the effects of the contrast agent) to not enter the intra-cellular space, hence $C_i = 0$. If we include the contribution of the intra-vascular space the result is the extended Kety model in Equation 2.30 which may be more appropriate for highly vascular regions of interest.

$$v_e \frac{dC_e}{dt} = K^{trans}(C_p - C_e) \quad (2.27)$$

$$C_t = K^{trans} \int_0^t C_p(t') \exp\left[-\frac{K^{trans}}{v_e}(t-t')\right] dt' \quad (2.28)$$

$$C_t = K^{trans} (C_p(t) \otimes e^{-\frac{K^{trans}}{v_e}t}) \quad (2.29)$$

$$C_t^{extended} = v_p C_p(t) + K^{trans} (C_p(t) \otimes e^{-\frac{K^{trans}}{v_e}t}) \quad (2.30)$$

Tofts (Tofts, 1997) assesses the inter-compatibility of other models ((Brix et al., 1990), (Larsson et al., 1990), (Buckley et al., 1994) and in the extensive collaboration of Tofts *et al* (Tofts et al., 1999) an effort is made to standardise the use of model parameters and their meaning under different conditions. The authors derive cases for both high permeability, in which contrast agent moves rapidly into the EES from the blood plasma, and low permeability models in which contrast agent transfer to the EES is slow and the model is dominated by vascular processes (the high-permeability model is known as the Kety model). Work on standardisation of parameters was also proposed by Armitage (Armitage et al., 2005), including a model by Hayton (Hayton et al., 1997) that is discussed further in Section 2.3.2. An important contribution by Armitage is the description of the non-linear relationship between contrast-enhancement and MR signal. Work by Buckley (Buckley, 2002) considers the uncertainty in parameter estimation, finding in particular that K^{trans} is systematically over estimated on model data. The suggested reasons are: ignoring the vascular contribution to the signal (over estimations of up to 54%), or non-uniqueness of the model fit. A good estimate of K^{trans} requires a good arterial input function and if this is to be obtained from the data, temporal resolution will need to be high enough to capture the signal.

Arterial Input Functions

Central to the pharmacokinetic model fitting process is the Arterial Input Function (AIF), $C_p(t)$ which is important in determining K^{trans} . In its most basic form the AIF can be calculated as a dual exponential decay from considerations of Equation 2.27 (Tofts & Kermode, 1991) where the parameter values are determined empirically but often taken from analysis by Weinmann (Weinmann et al., 1984) ($A = 3.99\text{kg}l^{-1}$, $B = 4.78\text{kg}l^{-1}$, $a = 0.144\text{min}^{-1}$, $b = 0.111\text{min}^{-1}$ and D is the injected dose), shown in Equation 2.31. The exponential decay is associated with initial mixing of contrast agent with tissue (and hence its loss from the blood pool) and alterations to the bolus shape through interaction with a 'body transfer function' (Orton et al., 2008). A further exponential could be used to relate the contrast agent removal by the kidneys with a biological half-life of Gd-DTPA. This is found to be about 20min under normal kidney function (Weinmann et al., 1984) and hence the contribution of this exponential decay is often ignored.

$$C_p(t) = D[Ae^{-at} + Be^{-bt}] \quad (2.31)$$

Although regularly discussed, the AIF used in Equation 2.31 is seldom used and the AIF is calculated directly from the acquired data. One early method (Andersen et al., 1996) injected a dose of $^{99m}\text{Tc} - \text{DTPA}$ in order to find the AIF by scintillation counting, assuming the equivalence of Gd-DTPA and Tc-DTPA pathways. Later, work by Port (Port et al., 2001) and Buckley (Buckley, 2002) discussed the importance of AIF estimation. The work by Port investigates differences in AIF between patients, finding not only that peak enhancement and time under the 10min curve vary by factors of 2.5 and 3.7 but that washout was more rapid with increased body mass. Using an AIF sampled directly from

the data has been used in work by Duhamel (Duhamel et al., 2006), Roberts (Roberts et al., 2006b) and Parker (Parker et al., 2006).

Work by Orton (Orton et al., 2007) presents examples of analytical arterial input functions built from exponential ($C_B(t) = a_B e^{-\mu_B t}$), gamma ($C_B(t) = a_B t e^{-\mu_B t}$) and cosine initial bolus shape functions ($C_B(t) = a_B(1 - \cos(\mu_B t))$) with a view to computationally efficient calculation of the convolution in Equation 2.29. The resulting blood plasma contrast-agent concentration $C_p(t)$ for the cosine bolus model $C_B(t)$ as it passes through the body, arriving at time t_B , is given by Equations 2.32 to 2.34 where it is convolved with a 'body transfer function' of the form $a_G e^{\mu_G t}$. The terms for μ_B and μ_G correspond to the rate constants associated with contrast agent mixing with the blood pool and whole body tissue respectively, with amplitude terms a_B (in $kg l^{-1}$) and a_G (in min^{-1}) describing the size of the bolus and the strength of its interaction with the body. The use of the cosine bolus function is empirical and presented in (Woolrich et al., 2004).

$$C_p(t) = a_B(1 - \cos(\mu_B t)) \otimes a_G e^{-\mu_G t} \quad (2.32)$$

$$C_p(t) = \begin{cases} a_B(1 - \cos(\mu_B t)) + a_B a_G f(t, \mu_G) & \text{for } 0 \leq t \leq t_B \\ a_B a_G f(t_B, \mu_G) e^{-\mu_G(t-t_B)} & \text{for } t > t_B \end{cases} \quad (2.33)$$

$$f(t, \mu) = \frac{1}{\mu}(1 - e^{-\mu t}) - \frac{1}{\mu^2 + \mu_B^2}(\mu \cos(\mu_B t) + \mu_B \sin(\mu_B t) - \mu e^{-\mu t}) \quad (2.34)$$

2.3.2 Registration of DCE-MRI

The increasing use of Dynamic Contrast-Enhanced MRI (DCE-MRI) in the assessment of therapy is discussed by Leach (Leach et al., 2005). However, the acquisition and further analysis of DCE-MRI is confounded by subject motion, due to the length of time needed to acquire a scan. Early results in functional SPECT showed that mis-registration by only 1/8 of a pixel can lead to count errors of 5%-10%, making the following pharmacokinetic analysis (Sychra et al., 1994) difficult. Similarly, pharmacokinetic analysis is subject to motion artefact errors in DCE-MRI. Early work by Zuo (Zuo et al., 1996) rigidly registered DCE-MRI volume pairs using a ratio-variance minimisation scheme, but the work was proposed as a method to automate manual registration of many DCE-MRI volumes and made no alteration for contrast-enhancement intensity profiles. Subsequent work can be divided loosely into two categories: *enhancement-cautious* approaches in which contrast-enhancement induced mis-registration artefacts are discarded as unrealistic motion behaviour in order to use a conventional registration; and *enhancement-driven* approaches in which enhancement profiles are used as additional information to guide the registration.

Early work on finding a cost-function for the registration of DCE datasets was produced by Acton (Acton et al., 1997). The work of Acton *et al* focuses on using principal components analysis to devise a cost-function that is robust to contrast-enhancement intensity changes. The work is applied to phantom cranial SPECT images of Dopamine receptors, in which images are corrupted from a gold-standard initial image by rigid body transformations. Contrast enhancement is modeled using patient data and corrupted by Poisson noise. Three cost-functions are compared: a (count) difference algorithm, a correlation algorithm, and the novel cost function which minimises the third order moment of the PCA eigenvalue

distribution. With the new cost function, the authors seek to maximise the variance that is contained in the early principal components, since it is assumed that motion-corruption leads to variance being shifted to the later, 'noisier' components. Effects of registration differences were measured using the χ^2 fit of the data with a bi-compartment kinetic analysis. The PCA cost-function was significantly better (with a p-value of < 0.001) at translational registration, but with no difference for rotational registration (it is suggested that the PCA method is sensitive to the interpolation method). The authors state that it is conceivable that the PCA-based cost function will fail in datasets in which the eigenvalue distribution represents the dynamic enhancement rather than the registration error, and this is likely to be the case for datasets with multiple enhancement patterns. The cost-function minimises the number of compartmental model fits by maximising the variance in the early components, therefore minimising the possible distribution of pixels.

An early attempt at enhancement-driven registration of DCE-MRI was made by Hayton (Hayton et al., 1997). The authors proposed a new model of contrast-enhancement uptake to allow the flexible monitoring of the effect of bolus injection. However, changes to the the bolus injection function led to only slight changes in contrast enhancement, a result that is no longer considered accurate (Roberts et al., 2006a). The registration operates on the assumption that if the images were perfectly registered, the residual of the model fits would be minimised: therefore a registration scheme (in this case optical flow adapted for brightness changes (Horn & Schunck, 1981)) can be driven by a cost-function that reduces the model-fit error. The work was tested on imposed translations in segmented 2D breast images. The method is unlikely to provide useful information for areas of insignificant enhancement where intensity noise dominates and the choice of model-fitting is inappropriate.

An enhancement-cautious approach was developed by Tanner (Tanner et al., 2000). Since the registration of contrast-enhancing features often results in unrealistic volume change (for instance, enhancing regions may be seen to shrink over the time-scale of the acquisition), the authors combined local rigid body constraints with a standard deformation to preserve shape and volume. A non-rigid registration algorithm (Rueckert et al., 1999) was used to parametrise the deformation, driven by an unaltered normalised mutual information cost-function (Studholme et al., 1999). It was suggested that this would be more appropriate since it does not require a linear relationship between intensities (although it does require that there is a consistent relationship). However, if the cost-function, and therefore the deformation, were truly appropriate, we might not expect any unphysical volume changes. Tanner *et al* also investigated volume changes by inspecting the volume change of a mask region and also investigated the use of multiple grid-point spacings. Coupling of grid-points within a free-form deformation were used to move local regions rigidly. Finer grids were demonstrated to result in larger volume change and the authors concluded that significant volume changes occur without correction (between -17% and 33% volume change). With correction by coupled control points, the volume change can be prevented. An attempt at validation was made by the same authors (Tanner et al., 2002), again in contrast-enhanced breast MRI. Biomechanical breast models were used to deform patient data selected with very little motion (introducing a whole breast volume change of 0.6%). More accurate registration was found over

the whole breast in the case of volume-preserving method (absolute volume change of 5.1%) than the comparison case of the standard unconstrained method (absolute volume change of 17.6%). This was an important validation step for the non-rigid registration of contrast-enhanced MRI, although the application of a biomechanical model is not an ideal gold-standard, particularly since the model deformations are not dissimilar to the returning registration deformations.

Further work on enhancement-cautious registration was proposed by Rohlfing (Rohlfing et al., 2003). The authors analysed the log of the Jacobian determinant during the progress of the registration and penalised any deformations from unity. Since the Jacobian describes the volume change associated with the change of co-ordinate system, a deviation greater than unity is an expansion and less than unity a contraction. As in Tanner (Tanner et al., 2000), a B-spline registration was used with a normalised mutual information cost-function. The authors compared a Jacobian-based volume change penalty term and a bending energy smoothness term in combination with the standard normalised mutual information cost function term. The results demonstrated significant volume decreases of between -1.3% to -78% for the standard registration method. The volume-change weighting factor proved robust and monotonic against volume change over a range of weightings, whereas the smoothness term did not. However, for both constraints there was a trade-off between volume preservation and motion correction, although less so in the case of the volume-preserving constraint. The authors suggested adaptively weighting the three cost-function factors as the registration proceeds. The use of constraining terms in non-rigid registration has shown success, however the methods do not suggest using contrast-enhancement information in the registration, which may provide much more information about the success of the registration.

Work by Hayton (Hayton et al., 1997) was extended by Xiaohua (Xiaohua et al., 2005) to allow combined image registration and segmentation of DCE-MRI. Overlap of the parameters from different tissue types in the model used by Hayton leads the authors to use a simple tissue attribute vector, consisting of the initial change in intensity after the addition of contrast agent and the slope of the late post-enhancement curve. Non-fat tissue is segmented into three classes and a Markov random field model is used to regularise and reduce noise. Optimum segmentation is given by the maximisation of the probability of pixels belonging to one tissue type. The assumption is made that optimum segmentation corresponds to optimum registration, since aligned pixels have less noisy model-fits and therefore can be better segmented. This is put into an iterative scheme and applied to DCE breast images. Segmentation results are improved after registration, but given that the process must be computationally intensive, and that there is no comparison with other registration methods, the true success of this process is questionable. It would be good to see the results of fitting different pharmacokinetic models.

Work on registration of myocardial perfusion images using active appearance models was proposed by Stegmann (Stegmann et al., 2005). The method uses a training set, that can be computed off-line, from analysis of the variance of data from previous perfusion study patients. Image registration of new data can then proceed using perfusion specific shape models. The method works well under enhancement and despite the computationally intense model-building, can register rapidly.

Another enhancement-driven registration method was proposed by Buonaccorsi (Buonaccorsi et al.,

2006). The authors realise that conventional cost-functions require that source and target image intensities maintain the same relationship and it is this relationship which is violated with contrast enhancement in DCE-MRI. In order to use conventional cost-functions, you should provide target images that resemble the intensity profile of the source. This is done by generating target images from model-fits of the unregistered data. The standard and extended Kety models are compared, using an arterial input function from the literature (Tofts & Kermode, 1991). This process can then be iterated, registering the original sources to the model target images and then re-fitting the model to the registered data and repeating the process. Unfortunately this process is limited by the appropriateness of model choice, so may only be done over a small region of interest. The size of the region of interest is also limited by the computational time of the pixel-wise model fitting procedure. The method is applied to abdominal tumours, which are considered rigid, and only translations are considered. The model parameter estimates vary as the registration proceeds, generally increasing K^{trans} and v_e , which may be expected as pixels come into better alignment relative to their pre-registration positions. The results of the model-fitting registration algorithm are compared to results produced by registration to the time-series mean: registration to the mean image reveals significant distortion to the subsequent model-fitting parameters. Residual model-fit errors are reduced after registration by the iterative model-fitting method, suggesting improved final model-fitting and therefore a more successful registration. Although results for the extended Kety model are considered more appropriate for the data used in this study, the success of the method relies on the choice of a good model which is difficult to determine. The process is currently only applied on small regions of interest, but if extended, would require the consideration of a non-rigid registration algorithm. An increased region of interest would require accurate controls on the model-fitting to ensure fitting of the correct model to different enhancing features. The increased computational time might also be prohibitive.

A combination of the work by Buonaccorsi (Buonaccorsi et al., 2006) and Hayton (Hayton et al., 1997) was proposed by Adluru (Adluru et al., 2006) for use in cardiac perfusion imaging. The authors use the extended Kety model applied in a form that is computationally efficient to fit. The data is registered by generating synthetic target images from model-fit data, and the registration is driven using the residual model-fit error as a cost-function (although again, all motion is assumed to be described by translations only). The results show an improvement in the estimation of kinetic parameters of 83% using the iterative registration scheme when compared to 68% with registration to a single image in the dataset. However, these results are calculated using the before and after model-fitting from the model used in the registration. Again, this work relies on the application of an appropriate model and it would be interesting to see it validated with a gold-standard simulation.

Work on the registration of DCE-MRI has also been proposed by Milles *et al* (Milles et al., 2008). The method can be compared to the work described in Chapter 5 (Melbourne et al., 2007b) but here an Independent Component Analysis (ICA) is substituted for a Principal Components Analysis (PCA). Milles *et al* find three independent components from dynamic contrast-enhanced cardiac MRI and optimise the 2D translation parameters between the original images and images generated by a combination

of the independent components. This work is currently limited to translation; non-rigid registration using this method would be an interesting extension. Using ICA as a direct substitution for the PCA in the work of Chapter 5 is limited by the independent components having no preferred order, PCA contains an implicit ordering of the principal components so that images can be generated consistently with principal components with higher variance being used earlier.

Work by Martel (Martel et al., 2007) applied work by Barber (Barber & Hose, 2005) who adapted the optical flow image transformation method of Horn (Horn & Schunck, 1981) to allow for contrast enhancement. DCE-MRI was simulated from patient data using the first two components of a principal component analysis to simulate enhancement profiles and a biomechanical finite element model to generate deformations, giving a gold standard. The optical flow method was compared with the results of an affine registration and a registration using a B-spline algorithm (Rueckert et al., 1999). The optical flow method, applied over a control-point distribution to reduce the degrees of freedom, was found to perform poorly at full resolution, but implementing a multi-resolution approach led to reduction in registration error. Control point spacing was required to be less than 16mm in order for successful registration. The optical flow method outperforms the affine registration and the results are comparable to the B-spline registration; however, the algorithm is only compared to the B-spline algorithm at 16mm control point spacing, so comparing results to the optical flow algorithm at a finer control-point spacing may be inappropriate. The optical flow algorithm is extremely fast when compared to the B-spline registration, but implicitly struggles in areas in which contrast enhancement is more significant than motion changes because by implication it assumes that intensity changes are due to motion.

There is a growing body of work concerning the registration of DCE-MRI. Many methods have been applied to the problem of registration although none are in widespread use. It is clear that conventional, general registration methods cannot be used without modification, either to the cost function or to the transformation method. Many recent methods require an iterative registration scheme in which a standard registration is used multiple times to allow for the updating of an external measure of success. The extraction of reliable and reproducible pharmacokinetic parameters may be improved using these image registration methods, this thesis presents work that might enable further improvements to pharmacokinetic parameter extraction, allowing improved and accurate diagnosis and assessment of therapy.

2.3.3 Diffusion Weighted MRI

The recent development and use of Diffusion Weighted MRI to infer structure by analysing restrictions to isotropic diffusion has yielded a large body of research work. The concept was devised with the addition, by Stejskal and Tanner, of extra diffusion sensitising gradients to the spin echo acquisition sequence (Stejskal & Tanner, 1965) as illustrated in Figure 2.3.

Diffusion Weighted MRI can be used to analyse the strength of restricted diffusion in a particular direction, under the diffusion imaging equation shown in Equation 2.35 representing the signal for a given b -value (in units of $s.mm^{-2}$). The b -value is the imaging parameter used to weight the diffusion signal and encompasses the effects of the gradient amplitude g , gradient duration δ and temporal separation

Spin Echo Diffusion Weighted Timing Diagram

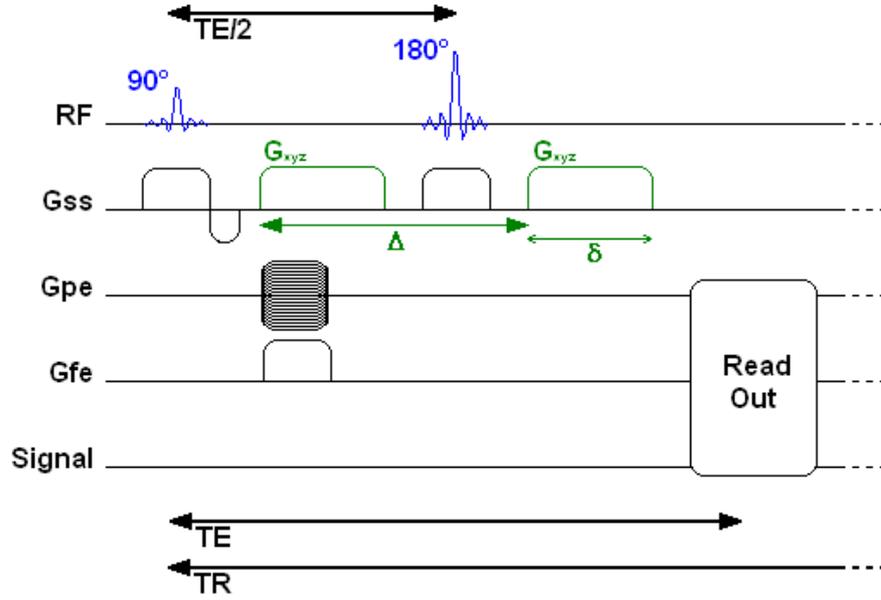


Figure 2.3: Example schematic timing diagram for a Diffusion Weighted Spin Echo acquisition sequence. Diffusion gradients (along any desired direction, but shown here along G_{ss}) are added either side of the 180° RF pulse, causing spin phase shifts that are refocussed dependent on the position and motion of the spins. For acquisition time reasons, the read-out block is likely to consist of an EPI sequence.

between the twin gradient echoes Δ used for diffusion imaging (Equation 2.36).

$$S(\mathbf{g}) = S_0 e^{-b\mathbf{g}^T \mathbf{D} \mathbf{g}} \quad (2.35)$$

$$b = \gamma^2 \delta^2 g^2 \left(\Delta - \frac{1}{3} \delta \right) \quad (2.36)$$

By observing the diffusion in multiple directions by varying the gradient direction \mathbf{g} , it is possible to calculate a second-order diffusion tensor \mathbf{D} as demonstrated by Basser *et al* (Basser *et al.*, 1994) (Basser & Pierpaoli, 1996) and also discussed by Batchelor (Batchelor *et al.*, 2003). In three dimensions we have a 3×3 Tensor but it is symmetric (i.e. $D_{xy} \equiv D_{yx}$), hence we need a minimum of 6 gradient directions to determine the tensor, plus a b_0 map (the image with zero b -value); however it is common to use many more gradient directions. Finding the eigenvectors of the diffusion tensor allows the inference of the dominant diffusion directions, which, in the brain, can be used to represent nerve-fibre orientation. Following dominant nerve-fibre bundle orientation between pixels has resulted in many groups publishing work on DT-MRI tractography using different methods (Melhem *et al.*, 2002), (Bammer *et al.*, 2003), (Behrens *et al.*, 2003), (Parker & Alexander, 2003).

2.3.4 Registration of Diffusion Weighted MRI

The length of the acquisition of a diffusion weighted MRI dataset exposes it to motion artefacts. Strong, long-duration diffusion gradients induce eddy currents to which the EPI read-out is sensitive, leading to image distortions. Patient movement may produce diffusion weighted images along different gradient directions that are misaligned. For further analysis or good tensor estimation, these images may need to be realigned. Attempts have been made to register scalar direction images ((Leemans et al., 2005), (Tao & Miller, 2006) and more similar to the method presented in Chapter 7 that of Bai and Alexander (Bai & Alexander, 2008)) or to remove the influence of outliers in the tensor estimation procedure (Chang et al., 2005). More commonly, registration is done on the post-analysis diffusion tensor images and a large amount of work has been produced ((Alexander et al., 2001), (Guimond et al., 2002), (Hecke et al., 2007) and (Chiang et al., 2008)). Registration of the component direction images is made difficult by the varying local contrast as a function of fibre direction. This local contrast variation may invalidate the assumptions of the registration similarity measures discussed in Section 2.2.1. This is analogous to the problems of registration of DCE-MRI and again, the problem is addressed in this thesis by developing a full field-of-view non-rigid registration method.

Chapter 3

Creation of Simulated Dynamic Contrast

Enhanced MRI Data

3.1 Introduction

In this chapter we develop a model to simulate abdominal dynamic contrast enhanced MRI (DCE-MRI) data. The chapter will discuss the development of simulated abdominal data incorporating both a deformation and an enhancement model for use later in the thesis. The model is developed primarily to provide an extensive basis for the testing of novel registration algorithms. With regards to the deformation model, an elastic deformation is used that will allow a coherent non-rigid deformation combined with a volume preservation modification to model stiff tissue regions. The simulated data includes an enhancement model to allow estimation of the recovery of parameter values after motion corruption.

A simulated deformation model of a breathing liver in an individual who is free breathing is developed in order to evaluate registration success when varying motion and enhancement parameters in DCE-MRI. We use this method to better understand the conditions of success for different registration methods and to better understand where they fail and to gain insight on the reasons for this failure. This will enable us to design better algorithms for these applications. A DCE-MRI scan often takes minutes in order that the contrast-agent washout can be observed, therefore the patient must breath and organs of interest such as the liver will move over time. Image registration can be used to re-align organs within the images, allowing further analysis for use in diagnosis and therapy evaluation. However, conventional registration methods require that images being registered have the same information and structure, but this requirement is not met in DCE-MRI, since the enhancement introduces new information into the images.

The liver is subject to motion due to subject breathing motion, the adjacency of the superior liver to the lungs and diaphragm exposes the organ to large superior-inferior deformations with the breathing cycle. To an extent the superior liver is protected by the ribcage and therefore may be expected to move predictably with the breathing cycle. This is unlikely to be the case for the inferior liver which is physically closer, and influenced, by the orientation and contractions of both gastro-intestinal and automotive abdominal muscles. The connection of the liver by five ligaments to the moving diaphragm and abdominal walls (Gray, 1918) exposes the organ to non-rigid deformations and complicated forms

of motion not dealt with by current image registration methods. Motion such as a relative movement of the liver sliding over fixed abdominal walls is particularly problematic and not dealt with by current registration algorithms. Deformation of the superior liver will be dominated by breathing type and depth. Breathing type is controlled by both the autonomic and conscious nervous system and may be affected by positioning of the subject, so to some extent the deformation may be trained or restricted. Breathing motion is likely to impart a cyclic deformation on the liver, causing a repetitive superior-inferior displacement according to lung filling. Depending on position, this motion may have important anterior-posterior and, to a lesser extent, medial-lateral components (Rohlfing et al., 2004). The motion itself is unlikely to be regular due to the competing influence of a large-tolerance feedback system trying to achieve blood-gas homeostasis (or at least clearance of fluctuating carbon dioxide levels) and irregular additional commands from the central nervous system. Additional intra-cycle variability arises due to the physics of breathing; it is more difficult to breath in than out because of the pressure gradients: i.e. the relaxed state of the respiratory system is gentle exhale and breathing air into the lungs requires the subject to do work. As a result, modelling of the breathing cycle is difficult and is compounded by unpredictable variations in phase and depth over both short and long-term periods.

The use of a comprehensive *in silico* simulation of DCE-MRI data allows complete control of the deformation and enhancement parameters. The creation of a gold-standard allows an assessment of the comparative success of image processing algorithms subject to how well the model represents the real situation it is approximating.

3.1.1 Finding an intrinsic T1 map

The intrinsic imaging parameter in DCE-MRI is the T1 value of tissue. As previously discussed, the effect of the popular Gadolinium based contrast agents is to reduce the observed T1 of a particular region. The influence of Gadolinium contrast agents on the T2 value is assumed to be negligible in the following analysis. Using a spoiled gradient echo sequence, the observed signal for a given T1 value is given by Equation 3.1 for flip angle α and repetition time TR . The influence of TE and other scanner parameters are included in S_0 . The change in T1 under the presence of contrast agent will be discussed shortly. Hence we can find a T1 map of intrinsic relaxation times by comparing the signal under varying flip angles. Equation 3.2 shows how to find the observed T1 from two images of different flip angle. Using multiple flip angle values allows a better estimation of the T1 value, for instance, by least-squares fitting.

$$S = S_0 \frac{\sin(\alpha)(1 - e^{-\frac{TR}{T_1}})}{(1 - \cos(\alpha)e^{-\frac{TR}{T_1}})} \quad (3.1)$$

For two different flip angles α_1 and α_2 , the T1 map can be found using Equation 3.2. For a spoiled gradient echo sequence, typical parameters might be $\alpha_1 = 2^\circ$, $\alpha_2 = 24^\circ$ with a TR of $4ms$.

$$T1 = TR \left[\ln \frac{\frac{S_1 \sin(\alpha_2)}{S_2 \sin(\alpha_1)} \cos(\alpha_1) - \cos(\alpha_2)}{\frac{S_1 \sin(\alpha_2)}{S_2 \sin(\alpha_1)} - 1} \right]^{-1} \quad (3.2)$$

Some typical T1 values for different organs are presented in Table 3.1 recorded from (de Bazelaire et al., 2004) for 1.5T and 3T.

Table 3.1: Typical T1 values for organs at 1.5T and 3T from (de Bazelaire et al., 2004). All values in *ms*.

Organ	1.5T				3T			
	T1	±	T2	±	T1	±	T2	±
Liver	586	39	46	6	809	71	34	4
Kidney Cortex	966	58	87	4	1142	154	76	7
Kidney Medulla	1412	58	85	11	1545	142	81	8
Spleen	1057	42	79	15	1328	31	61	9
Pancreas	584	14	46	6	725	71	43	7
Paravertebral Muscle	856	61	27	8	898	33	29	4
Subcutaneous Fat	343	37	58	4	382	13	68	4
Prostate	1317	85	88	-	1597	42	74	9

For the neuro-endocrine DCE-MRI data used in this thesis, the underlying T1 maps can be calculated from two different flip angle images. Flip angles of 2° and 24° were acquired with a repetition time TR of $4ms$. Via Equation 3.2 it is possible to calculate T1 values. It is also possible to produce a truly synthetic T1 map from Table 3.1, but this makes the generation of realistic images considerably more difficult and is left as future work. For the synthetic data generated in this chapter, we estimate T1 values from individual input images using Equation 3.1.

3.2 Developing A Liver Model

3.2.1 Liver Deformation

We aim to generate a deformation model of the liver that will allow it to appear to deform realistically and reversibly. Breathing motion is cyclic and undergoes a hysteretic motion, although the cycle may not be closed and its end point may drift over time (Blackall et al., 2006). Approximately 70% of motion is in the superior-inferior direction, with motion of a smaller extent in both the anterior-posterior (24%) and medial-lateral (7%) directions (Rohlfing et al., 2004). We model our image as an isotropic elastic medium, this is a reasonable approximation for non-rigid objects which resist an applied force and return to their initial configuration on removal of the force (an anisotropic modification is discussed below). The method will not be appropriate near objects such as bone. Initial results are shown for 2D motion, neglecting small medial-lateral deformations. To ensure a model that is both general and that has good deformation properties, we model a global image deformation, in which organ specific motion is induced by careful placement of forces. Deforming forces are placed in order to mimic breathing motion; they are strongest, resulting in largest displacement, in the region of the diaphragm and weighted toward a deformation in the superior-inferior direction. Forces may also be placed in the lower abdomen to mimic peristalsis. We require our forces to be time-varying, allowing the generation of a cyclic model meeting the properties discussed above. Random variation of the force magnitude allows us to simulate repeated breath-hold imaging conditions.

The forces in each direction are described here as Gaussian point forces. A location is chosen at a

point $[x_0, y_0]$ and a force applied symmetrically around this point for all $[x, y]$, the force has time-varying magnitude $A(t)$, and spatial extent b (Equation 3.3).

$$\mathbf{F}(x, y, t) = \frac{\mathbf{A}(t)}{b\sqrt{2\pi}} \exp \frac{-((x_0 - x)^2 + (y_0 - y)^2)}{2b^2} \quad (3.3)$$

The breathing model applies many forces of the above style; for instance, superior-inferior forces located in the lung region will drive the dominant breathing motion, greatest in magnitude in the diaphragm region. Additional anterior-posterior or medial lateral forces will impart perturbations to the breathing cycle. Modulating the forces in magnitude and direction with a sine-wave ($\mathbf{A}_y(t) = \sin(t)$, $\mathbf{A}_{xz}(t) = 0$) or a spline-based model will generate images across the breathing cycle. The solution of this force-field on the image is found by the solution of the isotropic linear elastic equation (Equation 3.4), allowing a displacement field to be calculated across the image.

An example is shown in Figure 3.1 for three superior-inferior force centres selected using a graphical user interface. This model allows more advanced force models to that described above. Each force point can have its magnitude modulated by a raised sinusoid or linear ramp over a period of time in order to model breathing or other types of force (see (George et al., 2005)). A raised sinusoid is often used, but it does not address the fundamental issue that a single breathing cycle is not symmetric - breathing in is more difficult than breathing out due to the pressure differences. More advanced work has been produced by McClelland (McClelland et al., 2006), modelling a single breathing cycle with a spline. The incorporation of a spline model is a desirable step, particularly when incorporating a more natural model of variations in breathing phase and depth. An example of using the spline model to describe breathing variation is shown in Figure 3.2 for six consecutive breath-holds. We define an initial spline (red), shown here as slightly saw-tooth (breathing in takes longer) and with a magnitude that will correlate with breath-depth. The spline nodes (green/yellow) are allowed to vary from these locations in subsequent breaths with a Gaussian distribution.

The solution for the displacement induced by the map of forces as described above is given by Equation 3.4. The solution of the displacement \mathbf{u} from the force \mathbf{F} (the parameters for μ and λ are set to 1 and 0 respectively as the elastic medium is both isotropic and we have no information to guide the choice for these parameters) is found here using a method developed by Cahill (Cahill et al., 2007b) for fixed boundary conditions. By writing Equation 3.4 as the product of itself with its adjoint (Equations 3.5-3.7), we can then consider the eigenvalues of ∇^4 to find an analytical solution using a few fast Fourier transforms (if the boundary conditions are not periodic (i.e. Neumann or Dirichlet) we may use an equivalent fast sine or fast cosine transform).

$$\mu \nabla^2 \mathbf{u} + (\mu + \lambda) \nabla (\nabla \cdot \mathbf{u}) = \mathbf{F} \quad (3.4)$$

$$L(\mathbf{u}) = \mathbf{F} \quad (3.5)$$

$$L^\dagger L(\mathbf{u}) = L^\dagger \mathbf{F} \quad (3.6)$$

$$\mu(\lambda + 2\mu) \nabla^4 \mathbf{u} = L^\dagger \mathbf{F} \quad (3.7)$$

Here we choose a zero boundary condition, requiring the fast sine transform to be used. This is

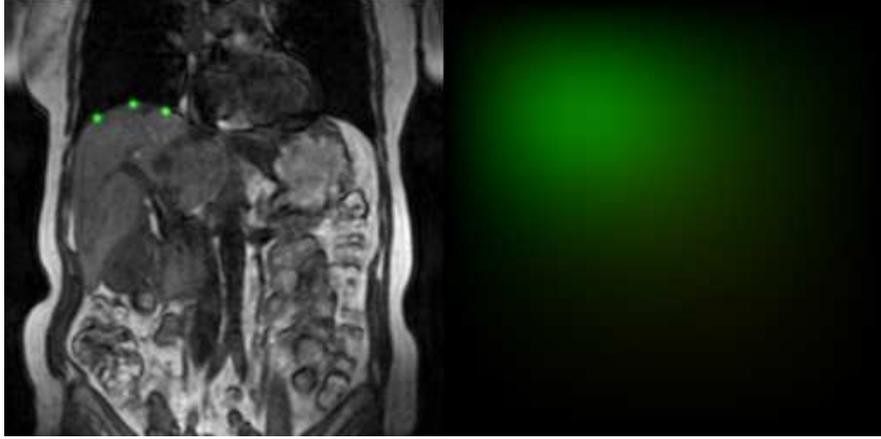


Figure 3.1: Example force seeds added to an underlying base image (left) and resulting deformation field (right). The green points represent centres of superior-inferior driving force from which the image displacement is calculated using an elastic equation to give a corresponding deformation field (right hand image). Breathing is modelled by varying the magnitude of the force centres by a raised sinusoid.

done to prevent objects moving from the field of view, if we want large displacements, the field of view can be made larger). Writing the discrete sine transform as Ψ (in 3 dimensions with sizes in each dimension denoted by M, N, P), it can be shown that the solution of \mathbf{u} for a given force is found by Equation 3.8 where the division by β is an element-wise divide - corrected for the undefined point at β_{000} . Formulations for the sine-transform and for β are shown in Equations 3.9-3.10. In this case the inverse of the sine transform is the same as the forward sine transform $\Psi\Psi(u) = u$. The fast sine transform is coded using the method described in Numerical Recipes (Press et al., 2007) based on a single fast fourier transform in each dimension.

$$\mathbf{u} = \Psi\left(\frac{\Psi(L^\dagger\mathbf{F})}{\beta}\right) \quad (3.8)$$

$$\text{where, } \Psi(\mathbf{u}_{ijk}) = \frac{8}{MNP} \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} \sum_{k=0}^{P-1} \mathbf{u}_{ijk} \sin\left(\frac{\pi mi}{M-1}\right) \sin\left(\frac{\pi nj}{N-1}\right) \sin\left(\frac{\pi pk}{P-1}\right) \quad (3.9)$$

$$\beta_{ijk} = 8\mu(\lambda + 2\mu)\left(\cos\left(\frac{\pi i}{M-1}\right) + \cos\left(\frac{\pi j}{N-1}\right) + \cos\left(\frac{\pi k}{P-1}\right) - 3\right)^2 \quad (3.10)$$

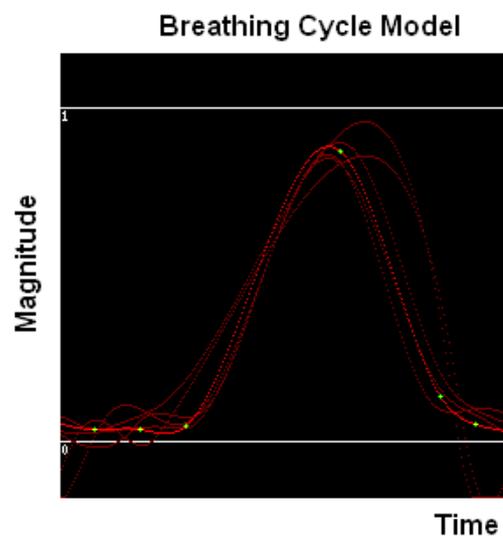


Figure 3.2: Using a spline model to describe breathing variation for six consecutive breath-holds. An initial spline (red) (shown slightly saw-tooth so that breathing in takes longer) is defined. The initial spline nodes (green/yellow) are allowed to vary from their defined locations in subsequent breaths with a Gaussian distribution in both time and magnitude.

3.2.2 A Contrast Enhancement Model

We now add a contrast enhancement model; for this we require a segmentation, different organs have different overall enhancement characteristics depending on blood requirements and vascular distance from the heart. Figure 3.3 demonstrates a hand segmentation of gross abdominal features overlaid on anatomical reference images. Each image is one of twenty slices of a contiguous abdominal volume of pixel size $1.37 \times 1.37 \times 5mm^3$, segmented by hand into liver, kidney, aorta & vascular features, heart (left & right side). Pathology may also be marked but additional pathology will be included as discussed below. The enhancement process in each of these can be modelled according to observed physical properties. Vascular features including blood vessels and tumour boundary angiogenesis enhance rapidly. Enhancement also fades rapidly from these regions. The highly vascular bulk liver enhances brightly, whilst the bulk tumour will have delayed enhancement, depending on the tissue status of its interior. A physiological description of this process and the total contrast agent, $C_t(t)$, at a given time is described by the widely-used Kety model (see Section 2.3.1). The parameters K^{trans} , v_p , v_e correspond to the volume transfer coefficient of contrast agent between blood plasma and extracellular-extravascular space (EES), and the fractional volumes of blood plasma and EES, respectively. $C_p(t)$ is the 'Arterial Input Function' describing the injection of contrast agent into the organ of interest. Equation 3.11 describes the total tissue concentration of contrast agent using the extended Kety model (see Equation 2.30). Since the liver has a dual blood supply we include a model of the contribution of both the contrast agent arriving via the hepatic artery $C_t^{arterial}(t)$ (Equation 3.11) and portal vein $C_t^{portal}(t)$ (Equation 3.12). The weighting of each contribution is given by γ representing the hepatic perfusion index (HPI) describing the observed ratio of arterial to total liver perfusion, for instance the contribution of the hepatic blood supply is about 25% from the aorta and 75% from the gastro-intestinal (portal) system, hence for Equation 3.13, $\gamma \approx 0.25$. However, this number will vary between individuals and due to pathology.

$$C_t^{arterial}(t) = v_p C_p^{arterial}(t) + K^{trans} \int_0^t C_p^{arterial}(t) \exp\left[\frac{-K^{trans}}{v_e}(t-t')\right] dt' \quad (3.11)$$

$$C_t^{portal}(t) = v_p C_p^{portal}(t) + K^{trans} \int_0^t C_p^{portal}(t) \exp\left[\frac{-K^{trans}}{v_e}(t-t')\right] dt' \quad (3.12)$$

$$C_t^{total}(t) = \gamma C_t^{arterial}(t) + (1-\gamma) C_t^{portal}(t) \quad (3.13)$$

A correct Arterial Input Function (AIF) is often difficult to determine, so an empirical model may be used or may be determined from the data (Buonaccorsi et al., 2006) (e.g. by tracking contrast enhancement in a segmented region of the aorta). Here we use a dual input model based on a cosine input function as developed by Woolrich (Woolrich et al., 2004) and discussed by Orton (Orton et al., 2008) given by Equations 2.33 and 2.34. Separate cosine arterial input models are applied within $C_p^{arterial}(t)$ and $C_p^{portal}(t)$ with values given in Table 3.2. The function coefficient values for μ_B and μ_G correspond to the rate constants associated with the cosinusoidal contrast agent bolus arrival and its temporal shape modulation under recirculation by a 'body transfer function' given by $a_G e^{-\mu_G t}$ (amplitudes are often expressed as $kg l^{-1}$ and rate constants in min^{-1}), hence this 'body transfer function' is given different coefficients for the arterial and portal input functions. A single input system such as the aorta can be modelled with a value of $\gamma = 1$.

Table 3.2: Modelled Arterial Input Function Parameters for Dual Cosine Input Model (see (Orton et al., 2008) and (Parker et al., 2006) for source.)

	a_B	μ_B	a_G	μ_G	t_0
	mM	min^{-1}	mM	min^{-1}	min
$C_p^{\text{arterial}}(t)$	4.90	22.8	1.36	0.171	0
$C_p^{\text{portal}}(t)$	1.69	11.8	2.33	0.145	0.1

Table 3.3: Modelled Pharmacokinetic Parameters for given organs from consideration of vascular properties (see text for description) (Parker et al., 2006)

Organ	K^{trans}	v_p	v_e	HPI	Onset
	(min^{-1})				(min^{-1})
Right Heart	-	1	0	1	0
Aorta	-	0.8	0	1	0.12
Kidney	0.33	0.2	0.2	1	0.2
Liver (A)	0.27	0	0.25	0.3	0.23
Liver (P)	0.27	0	0.25	0.3	0.33

A further important factor when considering contrast enhancement of multiple organs with multiple blood supplies is the bolus onset time. In the case of the liver, we alter the relative onset times of $C_t^{\text{arterial}}(t)$ and $C_t^{\text{portal}}(t)$ by adjusting the bolus arrival time. We alter the parameters K^{trans} , v_e , v_p and HPI to give enhancement profiles with the behaviour we wish each region to display. Model parameters are shown in Table 3.3 for comparison with literature values (Parker et al., 2006). We currently use hepatic values of $v_p = 0$ for simplicity, hence the enhancement curve modelling does not necessarily correspond to the *in vivo* biological situation. For the purposes of registration testing in this thesis it is the contrast variation that is important, however, the model may be refined in future. As illustrated by Parker *et al* (Parker et al., 2006), the range of pharmacokinetic parameter values is quite variable, so the relative values chosen are important. In the liver, the delay between the arterial and portal enhancements is shown and the HPI is given a value of 0.3 (Totman et al., 2005). The kidney is given a large v_p to represent a large blood supply and corresponding large plasma fraction. Large blood volumes are modelled with v_p for the heart and aorta. It would be possible to mimic some dispersion to the bolus in the left ventricle and aorta by using a large K^{trans} to model the disruption to the bolus passing between cardiac chambers, although this result would be difficult to interpret biologically.

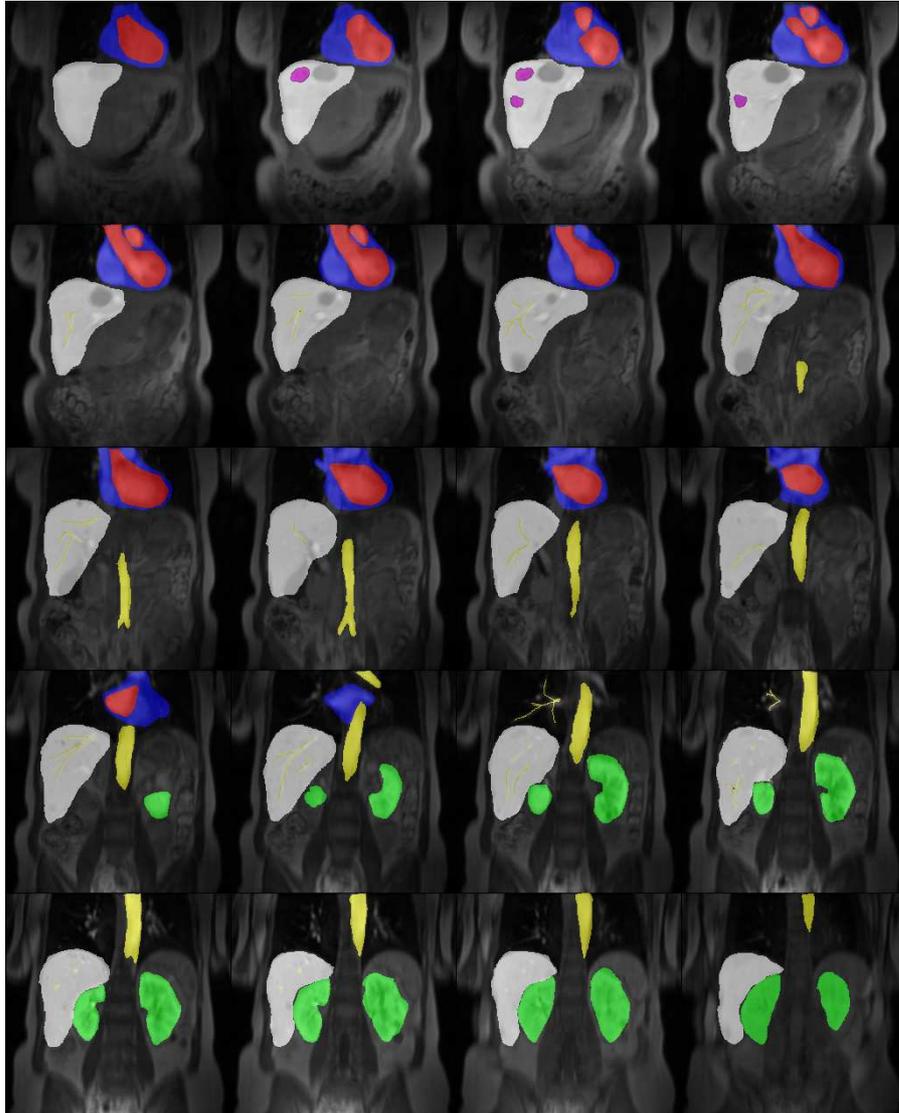


Figure 3.3: Example of 3D Gross Abdominal Segmentation overlaid on anatomical reference images. Segmented by hand into liver (white), kidney (green), aorta & vascular features (yellow), heart (left & right side (red & blue respectively)). Pathology may also be marked (magenta).

3.2.3 A Tumour Model

In addition to the gross organ segmentation described above, it is useful to add particular pathology. This will be particularly important for future work investigating the success of registration in areas of complex enhancement. The model in this region is a crucial inclusion in the development of simulated DCE-MRI since it is the pathology that motivates the acquisition of DCE-MRI data. Here we introduce models of tumours with different enhancement artefacts in different locations. After selecting a tumour location, a roughly circular (in 3D, spherical) boundary is drawn around this point. This is done by setting two parameters, one governing the radius and one allowing the radius to deviate away from a circle. The radius is defined at a set of spoke locations in turn and may be allowed to vary with a Gaussian distribution. The gaps between the spokes are interpolated with a cubic-spline to ensure a smooth boundary. The circle is filled to a given radius to give different pharmacokinetic properties between the boundary and tumour core (see Figure 3.4).

Pharmacokinetic parameters are chosen to mimic particular types of tumour: 1) isotropic tumour enhancement; 2) filling tumours where the rim enhances quickly and later the tumour core enhances; 3) necrotic core tumour where the rim enhances quickly and the core does not enhance; 4) poorly defined tumours with a *mosaic* appearance of different enhancement characteristics. An illustration of the application of this tumour model is shown in Figure 3.5. More vascular areas such as angiogenic regions are given larger values for the transfer constant K^{trans} and blood volume v_p . The *HPI* may also be increased to reflect increased arterial vascularity (Totman et al., 2005).

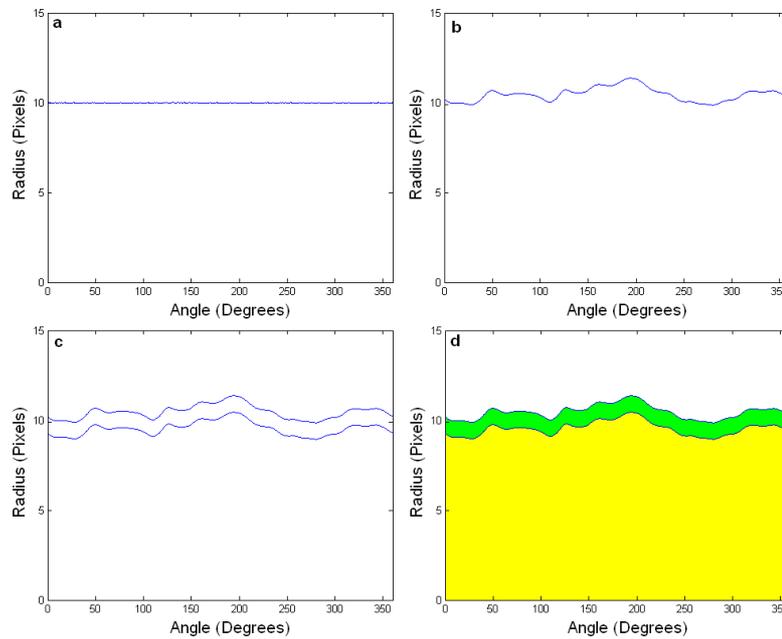


Figure 3.4: 2D pathology construction. **a** Circular boundary. **b** distorted boundary, smoothed with splines. **c** boundary thickness definition. **d** labelling pharmacokinetic features (colour coding interior and exterior).

Parameters for varying pathology are shown in Table 3.4. Tumour interiors are given lower values

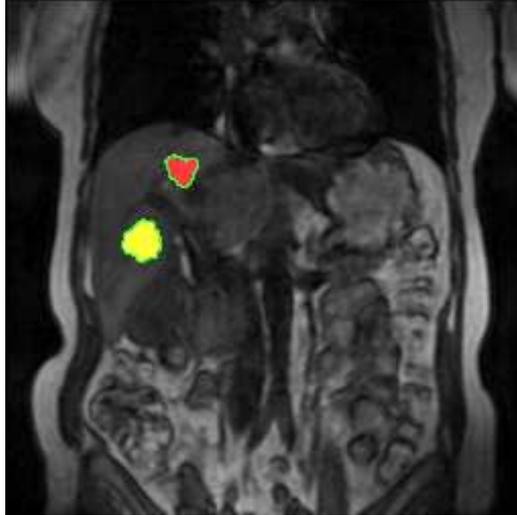


Figure 3.5: Example tumours added to underlying (pre-enhancement) image. Segmentation colours correspond to different pharmacokinetic parameters for tumour rim and tumour core. For tumour generation process see text.

Table 3.4: Modelled Pharmacokinetic Parameters for Simulated Pathology (see text for description and compare with Table 3.3). We increase the K^{trans} of the tumour boundary to mimic the expected rate-constant increase due to angiogenesis and add some delay in the onset of enhancement in a filling region.

Organ	K^{trans} (min^{-1})	v_p	v_e	HPI	Onset (min^{-1})
Tumour Boundary	0.33	0.25	0.2	1	0.23
Filling Region	0.17	0	0.18	0	1.3
Necrotic Region	-	0	0	0	0
Normal Liver	0.27	0	0.25	0.3	0.23

for K^{trans} , a necrotic region would not enhance. Values for v_p are increased to represent increased vascularity in angiogenic regions and are kept at zero in other regions, for comparison with Table 3.3.

Figure 3.6 shows the corresponding parameters for three types of tumour corresponding to the first three types discussed above. The colour segmentation is converted into contrast-enhancement uptake curves via table 3.4, which are then converted to signal via the spoiled gradient echo equation as in Equation 2.26. Figure 3.6 presents an example showing contrast-enhancement as a function of time for ten time-points. Figure 3.7 demonstrates the intensity time-curves generated for a specific flip angle, T1 and TR for different organs, two sub-figures are shown demonstrating signal intensity curves for both gross organ segmentation and for pathology with values illustrated in Table 3.4. Intrinsic T1 values are drawn from Table 3.1.

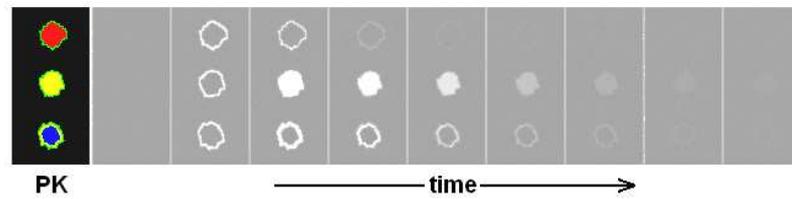


Figure 3.6: 2D pathology examples over an approximately 3 minute time period for top) peripheral tumour enhancement; middle) filling tumours where the rim enhances quickly and later the tumour core enhances; bottom) necrotic core tumour where the rim enhances quickly and the core does not enhance.

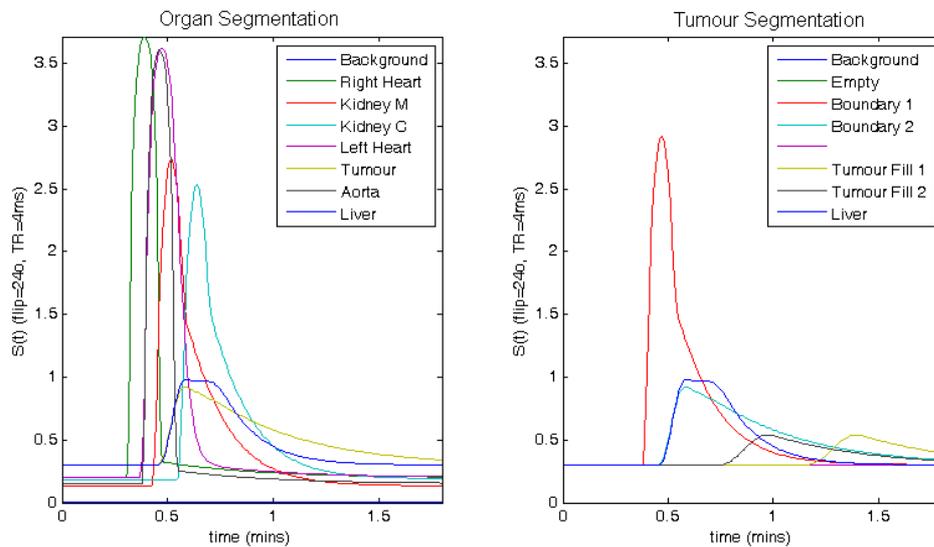


Figure 3.7: Example intensity-time curves used to simulate the uptake profile of major organs (left) and curves used in Section 3.2.3 to model pathology. Signal is generated for $\alpha = 24^\circ$, $TR = 4ms$ and for intrinsic tissue $T1$ found from Table 3.1. Note that these curves have an unrealistically rapid wash-out due to an implementation error that should be resolved prior to additional work regarding biological pharmacokinetic parameter extraction. See text for parameter choice and further clarification.

3.2.4 Proposed Volume Preservation Modification

Our deformation model is intrinsically isotropic, a situation which is implausible where tissues have different stiffness. As an example, in the breast, tumour tissue is found to be up to 15 times stiffer than normal breast tissue (Sarvazyan et al., 1994). A full discussion of the character of the liver, one that would be useful for further development of a biomechanical model, is given by (Liu & Bilston, 2000). To accommodate stiffness variation, we modify our deformation fields retrospectively to ensure that tumours move rigidly. The elastic equation in Equation 3.4 is isotropic, the parameters relating resistance to shear forces (μ) and resistance to internal expansion and compression forces (λ) are fixed throughout the medium. If we wish to vary them locally, we must expand Equation 3.4 to Equation 3.16 (see (Lester et al., 1998) and (Little et al., 1997)). The motivation for Equation 3.16 is given by Equations 3.14 and 3.15 describing the force as a function of the stress (see also Appendix A).

$$\sigma_{ij} = \mu \left[\frac{\delta \mathbf{u}_i}{\delta x_j} + \frac{\delta \mathbf{u}_j}{\delta x_i} \right] + \lambda \delta_{ij} (\nabla \cdot \mathbf{u}) \quad (3.14)$$

$$\mathbf{F}_{visc} = \sum_{j=1}^3 \frac{\delta \sigma_{ij}}{\delta x_j} \quad (3.15)$$

$$\mathbf{F}_{visc} = \mu \nabla^2 \mathbf{u} + (\mu + \lambda) \nabla (\nabla \cdot \mathbf{u}) + (\nabla \mathbf{u}^T + (\nabla \mathbf{u}^T)^T) \nabla \mu + (\nabla \cdot \mathbf{u}) \nabla \lambda \quad (3.16)$$

The use of Equation 3.16 may be appropriate for non-rigid image registration if we wish to prevent the deformation of objects we know to be well-registered. To some extent this is incorporated into the paper by Lester (Lester et al., 1998). An explicit alteration in the case of DCE-MRI would be to monitor the success of the model-fitting. If pixels are well-fitted (E.g. if they have relatively low residuals) then the viscosity may be locally increased. The success of the registration can then be governed by the fraction of pixels considered to be well-fitted. This concept is left as future work.

It is also possible to ensure rigidity if we segment the tumours and give every pixel in the tumour values corresponding to an approximation of the best-fitting 4×4 affine matrix. This can be calculated in a least-squares fashion to obtain parameters for rotations, scales and shears. Depending on the properties of the desired resultant deformation, we can remove the influence of particular parameters (for instance the scaling). This process does not remove the problem of smoothing the deformations together. This can be done as by using a smoothing filter at the region boundary or by using a spline to interpolate from the affine block into the fluid deformation field. Alternative formulations of locally rigid (or locally affine) registrations have been developed by Narayanan (Narayanan et al., 2005) and Commowick (Commowick et al., 2008).

Figure 3.8 demonstrates the requirement for the generation of a good intermediate deformation field. For more complex objects, the definition of the deformation across the boundary is crucial to avoid discontinuities in the image. Figure 3.9 demonstrates the incorporation of a spline-based interpolation of the deformation fields between the affine and elastic deformation blocks. The correction algorithm is described in Table 3.5. In practice and in the results generated in Chapter 5 we implement a Gaussian filter at the boundary to smooth the non-rigid and affine deformations together.

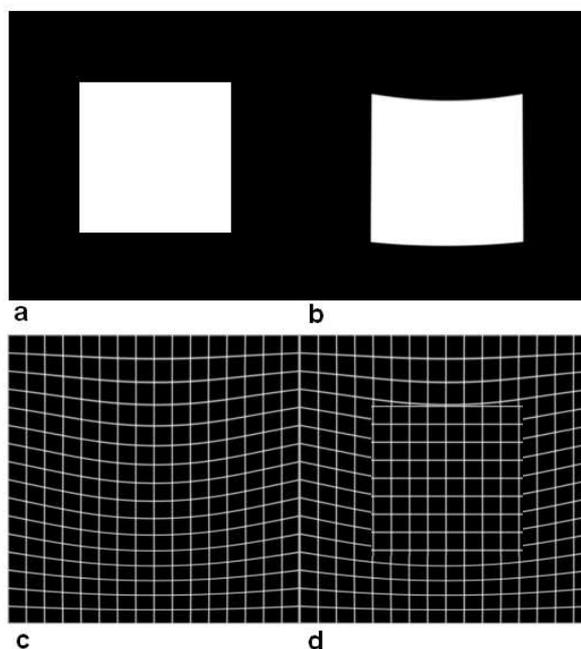


Figure 3.8: Maintaining rigid object shape in an elastic deformation: **a** Original Square. **b** Square moved downwards using elastic deformation. **c** Unmodified elastic deformation. **d** affine deformation patched into elastic deformation, note the unacceptable deformation discontinuity at the boundary.

Table 3.5: Algorithm for patching affine transformation into a global elastic deformation (see text for discussion)

For each affine object:

- 1) Segment non-rigid part of elastic deformation field that will be made rigid.
- 2) Approximate segmentation as affine using least-squares fitting.
- 3) Re-insert affine deformation into elastic deformation.
- 4) For each point on the boundary of the affine deformation:
 - i) find the B-spline that extends a depth $\pm 2xD$ into both the elastic and affine deformations.
 - ii) replace pixels within $\pm D$ by the interpolated value.

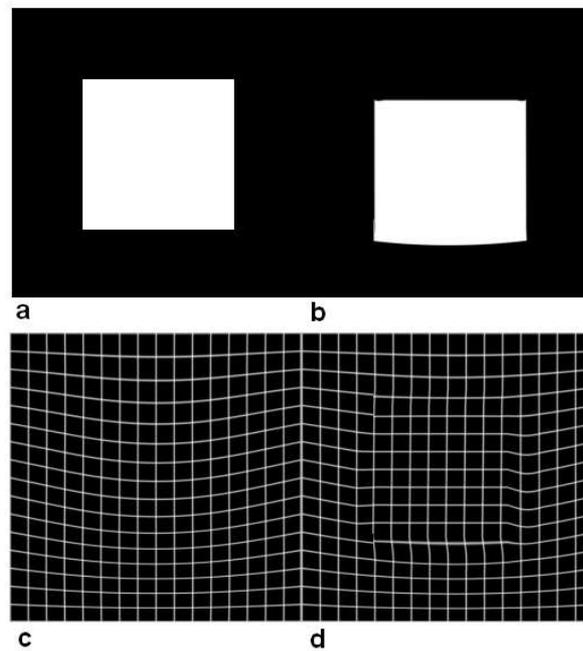


Figure 3.9: Maintaining rigid object shape in an elastic deformation: **a** Original Square. **b** Square moved downwards using elastic deformation. **c** Unmodified elastic deformation. **d** affine deformation patched into elastic deformation, now interpolated between affine and elastic deformations for comparison with Figure 3.8 using the method described in Table 3.5

3.3 Motion Model Examples

Figure 3.10 demonstrates images subject to the motion model described in Section 3.2.1. A force model is used to deform the liver in a superior-inferior direction with sinusoidal (breathing-like) amplitude with a period of 10 images. The green overlay outlines the liver position in the first image. An example movie is included on the supplementary CD (See Appendix E).

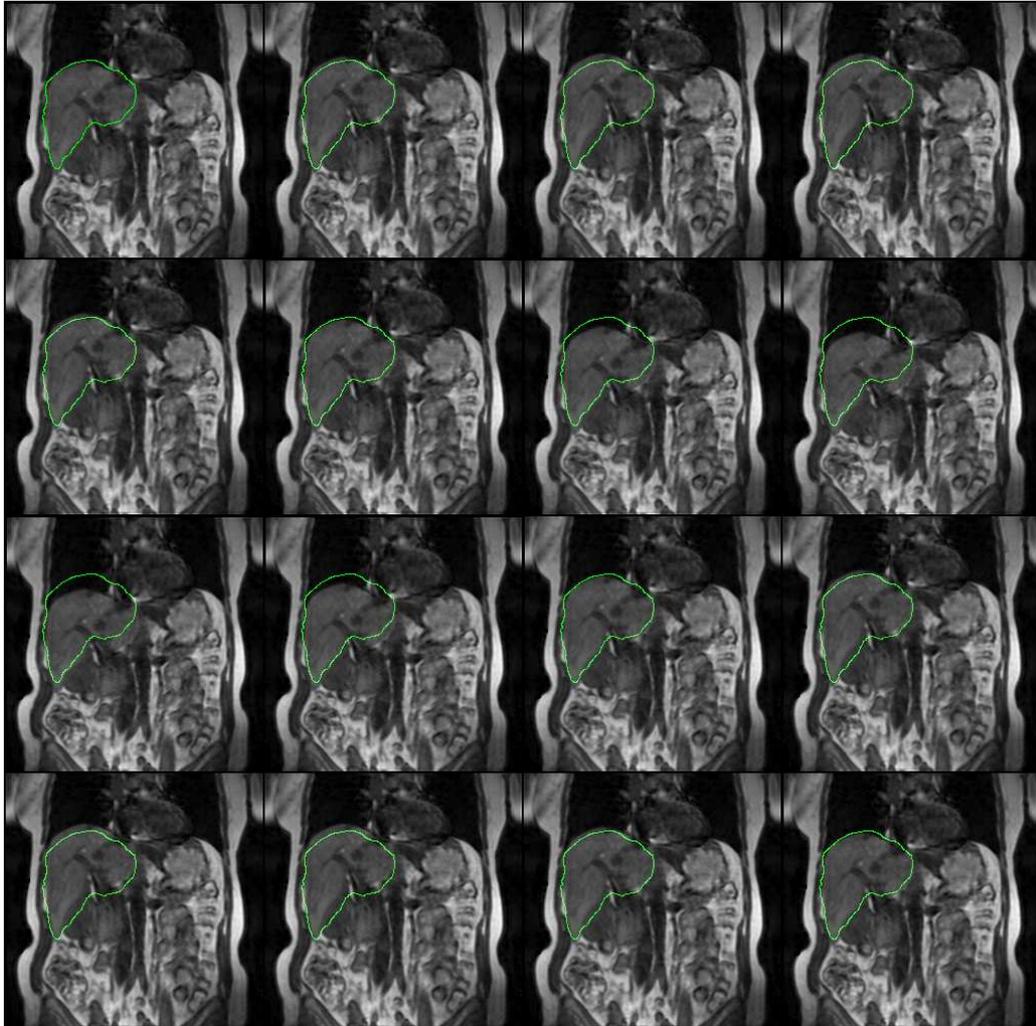


Figure 3.10: Example deformation for superior-inferior *breathing motion* with 10 image cycle length with additional random medial-lateral deformations.

Figure 3.11 shows the (normalised) difference images of each frame in Figure 3.10 with the first image.

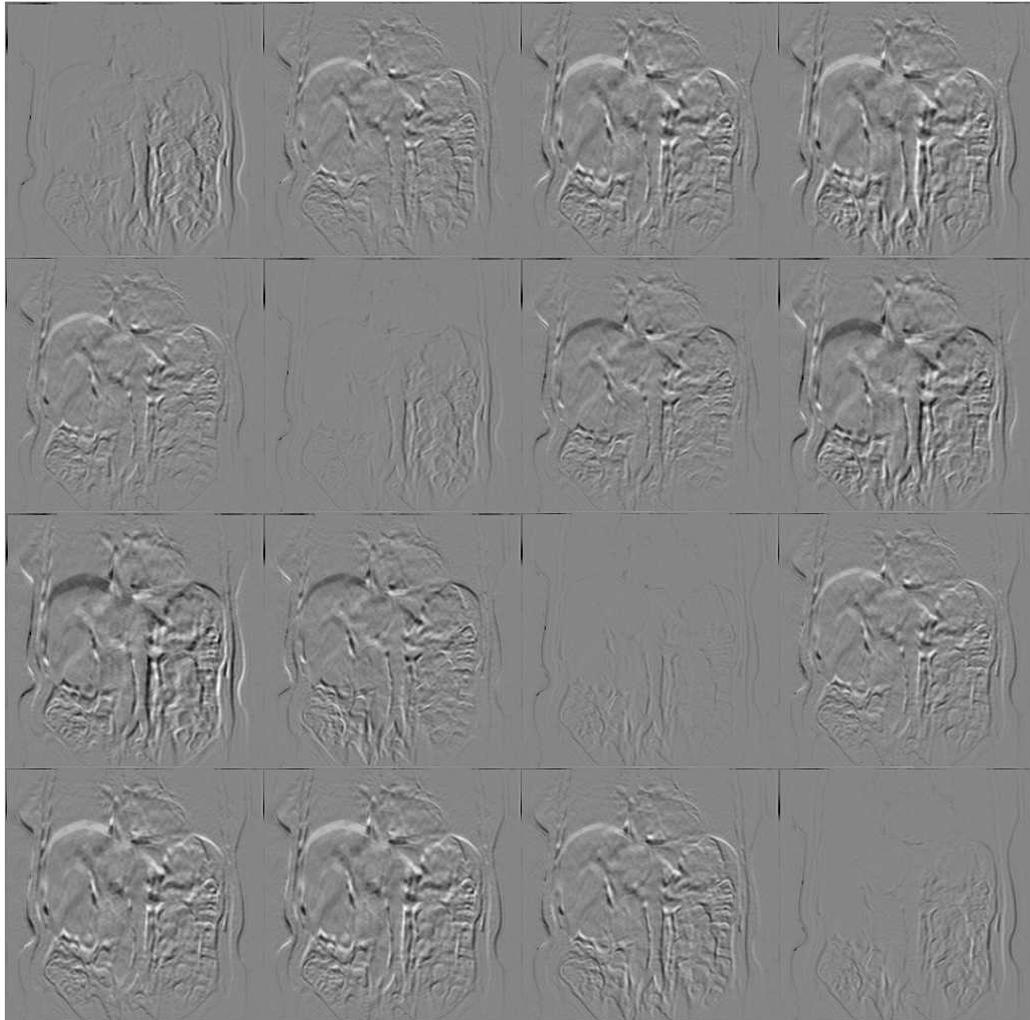


Figure 3.11: Example difference images (with original un-deformed image) for superior-inferior *breathing motion* with 10 image cycle length for comparison with 3.10.

3.4 Enhancement Model Examples

Figure 3.12 demonstrate the contrast enhancement model given the segmentation illustrated in Figure 3.3. Early enhancement of the heart and aorta is followed by enhancement of the liver and an embedded tumour. In this example, enhancement occurs over a short period of time and wash-out is unrealistically rapid. Figure 3.3 shows a further example of the enhancement model from a sagittal perspective. In this case the enhancing kidney is shown and two large hepatic masses are modelled. Example movies are included on the supplementary CD (See Appendix E).

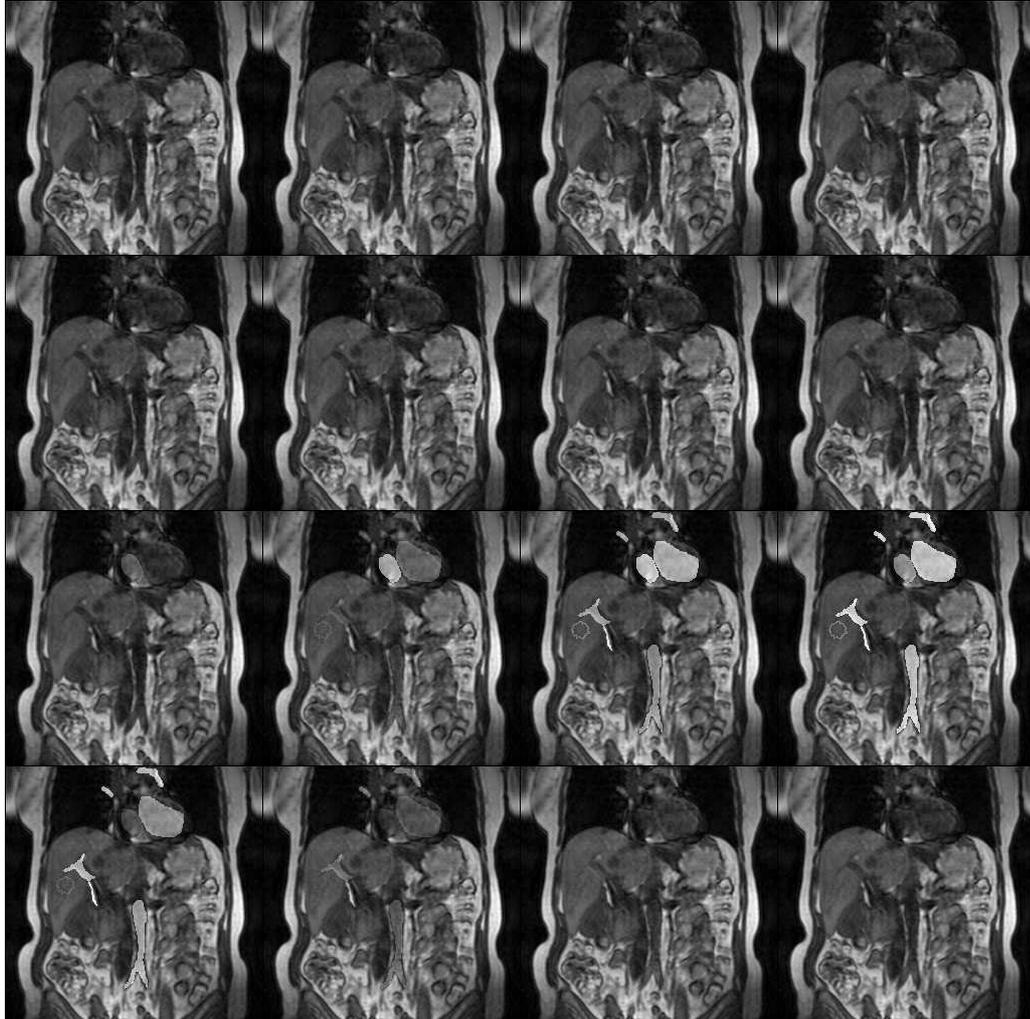


Figure 3.12: Example contrast enhancement time course for coronal images.

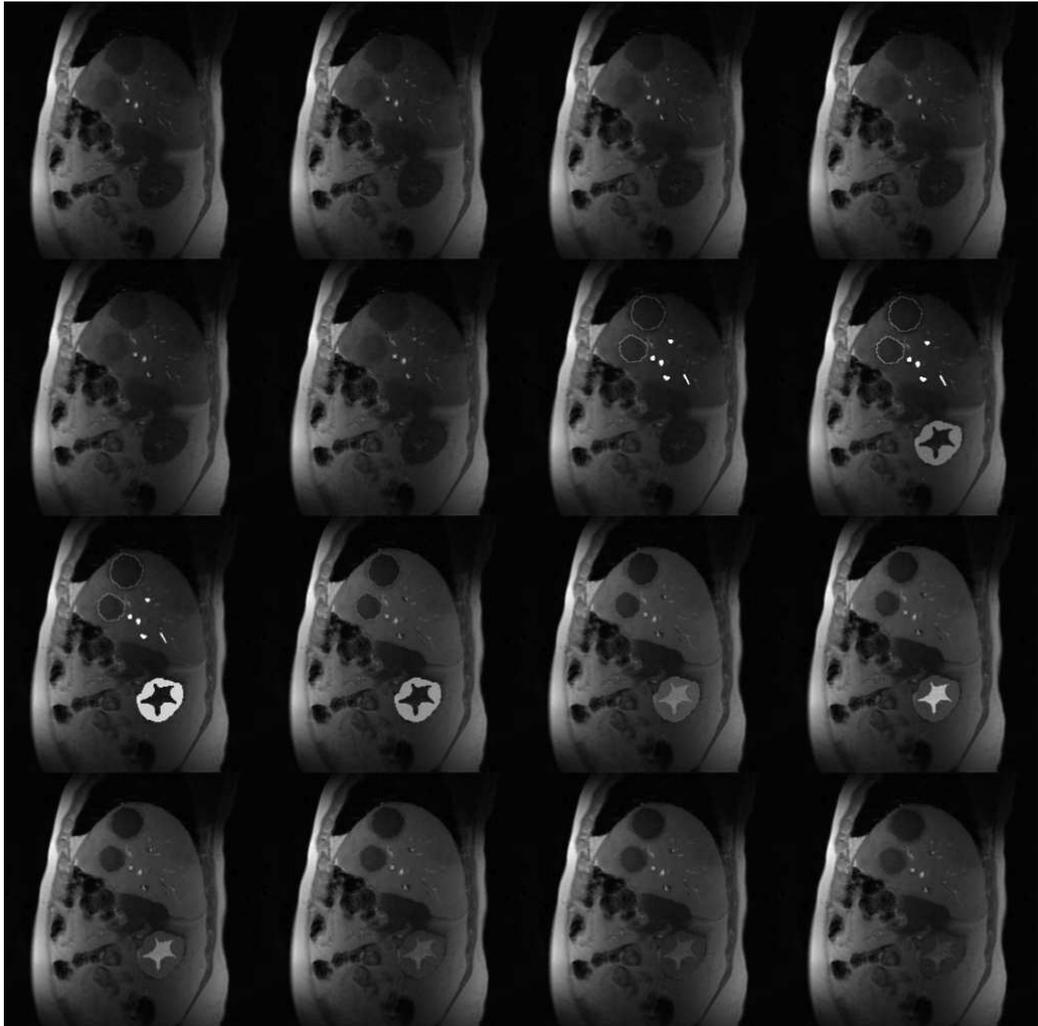


Figure 3.13: Example contrast enhancement time course for sagittal images.

3.5 Conclusion

The model presented above has been developed to provide an extensive basis for the testing of novel registration algorithms. With regards to the deformation model, the elastic deformation allows a coherent non-rigid deformation to be used. The deformation is one that should be correctable by a registration algorithm, provided the choice of cost-function is appropriate. The deformation is appropriate for algorithm testing. However, the global elastic deformation lacks the realism required for inferring registration success in real-world applications.

The inclusion of an organ specific contrast-enhancement model is an important step. If registration accuracy was to be tested by software designed for pharmacokinetic model-fitting, the simulated data has included a well-developed enhancement model that could allow an estimation of real-world parameter extraction accuracy. The importance of testing registration algorithms on known pharmacokinetic parameters is required for validation purposes: registration may be visually accurate but the acid test remains the ability to extract accurate pharmacokinetic properties from the entire DCE-MRI dataset. The parameter curves shown in Figure 3.7 do not appear realistic, having a rapid wash-out phase in all cases. This is due to an error in the implementation of Equations 3.11 and 3.12. Although this makes the generated enhancement curves unrepresentative of the *in vivo* situation, they may still be used for registration validation within the scope of this thesis.

The inclusion of organ specific deformations, perhaps using finite-element methods, would be a necessary development if the model was to be used outside of its purpose of testing the success of registration algorithms. If realistic biological deformations were used, the method may be used to generate synthetic data to match an existing dataset, predicting both accurate biological deformations and enhancement parameters.

As the model stands, it may be used for analysing registration success and in particular the failure of registration under contrast enhancement due to inappropriate cost-function selection. The choice of registration cost-function is discussed in the next chapter.

Chapter 4

Cost Functions and Contrast Enhancement

4.1 Introduction

Registration of Dynamic Contrast-Enhanced Magnetic Resonance Images (DCE-MRI) of soft tissue is difficult. Conventional registration cost-functions that depend on information content are compromised by the changing intensity profile, leading to mis-registration. This chapter will outline the requirement for a registration method that accommodates contrast enhancement by discussing the failures of common cost-functions. A method is also introduced to enable the assessment of the registration status of a group of images to one another. This method is then used to provide some assessment of registration accuracy when attempting to determine the best choice of anchor image within a group of images.

4.2 Conventional Cost-Functions

Image registration cost functions do not distinguish between differences due to motion artefacts and differences due to contrast enhancement, therefore when calculating forces that minimise a cost function, contrast enhancement can induce mis-registration. This often leads to distortion at enhancement boundaries, compromising registration success. Figure 4.1 demonstrates enhancement of features seen in a dynamic contrast enhancement sequence. Figure 4.1a is a pre-enhancement image. Figures 4.1b and 4.1c are images acquired during the passage of the bolus through the heart from right side to left side and into the aorta. This is an extreme case of the changes in images under enhancement. The rapid enhancement of the heart and the transitional appearance of vascular features (such as those seen in the liver) combine to confound registration cost-functions. The reason for this failure is now discussed. The derivation of the force gradients for each cost-function are included in Appendix B.

4.2.1 Method

For the cost-functions discussed in Chapter 2 we will discuss the formation of image registration force gradients when registering a float image \mathbf{F} to an anchor image \mathbf{A} . The images that will be analysed are the real images in Figure 4.1 where we register the two post-enhancement images, \mathbf{b} and \mathbf{c} , to the pre-enhancement image, \mathbf{a} . These are images from the central slice of patient three in Table 1.2 and contain both small amounts of motion and contrast enhancement. The dense force gradient images are suitable for implementation directly into a fluid or diffusion based registration algorithm. Analysis of

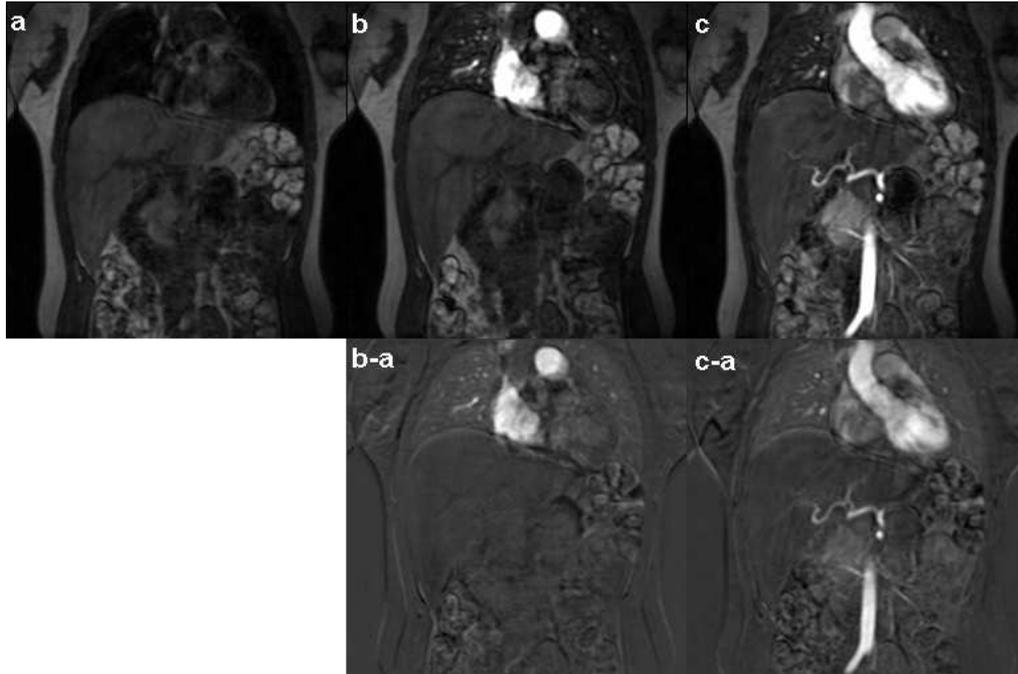


Figure 4.1: Example real contrast-enhanced images **b** and **c** and pre-enhancement image **a** demonstrating passage of bolus through right and left sides of the heart and aorta. Difference images for **b-a** and **c-a** are also shown to illustrate differences in position of abdominal wall and superior liver.

the gradient images for both x and y force directions will be presented as evidence for failure of the cost-function under contrast-enhancing features. Arrow plots of the resulting deformations after five iterations of a fluid based registration algorithm are shown as additional evidence on the fourth row of Figure 4.2 to 4.6 (see Appendix C).

4.2.2 Results

Statistical Alignment

A simple least squares alignment is given by Equation 4.1 as found in (Christensen et al., 1996). Its derivation is found in Appendix B. Differences in intensity between images are penalised and the force reduces to zero for $A = F$. Figure 4.2 demonstrates the dominance of contrast enhancement on the cost-function gradient images that would then be used in a fluid or b-spline registration algorithm. This dominance reduces the chances of correct registration in areas outside the dominant areas. Since gradients are largely the result of contrast-enhancement, we do not expect (and do not achieve, when used) a correct registration. The gradients in Figure 4.2 will result in shrinking of the enhanced heart and aorta as implied by the deformation field arrow plots.

$$\frac{dLS}{dx} = (A - F)\nabla F \quad (4.1)$$

An alternative is the cross-correlation force (or its variants using different normalisations), found by considering the change to the overall cross-correlation value for individual pixel displacements. The

gradient of this (unnormalised) cost-function is represented by Equation 4.2. If we were to properly normalise this measure we would have to include additional terms that are derived in Appendix B. This mediates against displacement fields dominated by a few regions of large force values such as those found using least-squares. However, the fact that there is often not a linear relationship between pixels (due to one-to-many intensity relationships between anchor and float) theoretically limits the cost-function's applicability. The gradient images in Figure 4.3 show detail not visible in Figure 4.2 and there are correcting gradients in the medial-lateral direction on the abdominal walls. Gradients in the enhancing heart and aorta are less well-defined, but the increased noise might result in unpredictable registration results that introduce distortions in the resulting deformation field.

$$\frac{dCC}{dx} = A\nabla F \quad (4.2)$$

Information-Based Alignment

Starting with joint entropy, the calculation of forces proceeds by analysing the change to the total entropy by moving one pixel between two intensity bins. The derivation is found in Appendix B, but for large numbers of pixels in each intensity bin, the change in joint entropy is given by Equation 4.3 (Crum et al., 2005) where $P(F_{ij}, A_{i-1j})$ is the joint histogram entry for the intensity values at location ij in the anchor \mathbf{A} and float \mathbf{F} images (N is the total number of pixels). Figure 4.4 uses 64 bins to give a good bin population. Superficially the gradients do not seem to act to disrupt the enhancing features, however, as discussed in Chapter 2, joint entropy on its own is not a good cost-function to minimise theoretically. Chapter 5 (see Figure 5.8) will also show some evidence of difficulties in finding a smooth minimisation of joint entropy.

$$\frac{dJE}{dx} = -\frac{1}{N} \log \frac{P(F_{ij}, A_{i-1j})}{P(F_{ij}, A_{i+1j})} \quad (4.3)$$

Similarly for Mutual Information, we make adjustment to include the effect of pixel movement on the marginal entropy (Appendix B) to achieve Equation 4.4 where we include the entry from the respective single (marginal) image histogram for the intensity value at location ij in the float \mathbf{F} image (N is the total number of pixels). The resulting gradient images are shown in Figure 4.5. The correction of the abdominal wall displacement is visible, but there remains some evidence of mis-correction in the enhancing heart and hepatic artery. This structure again, by inspection of the gradient images, will result in shrinking of enhancing regions.

$$\frac{dMI}{dx} = -\frac{1}{N} \log \frac{P(F_{ij}, A_{i-1j})P(F_{i+1j})}{P(F_{ij}, A_{i+1j})P(F_{i-1j})} \quad (4.4)$$

Gradients for Normalised Mutual Information can be developed by considering the two previous results for joint entropy (H_{AF}) and mutual information ($H_A + H_F - H_{AF}$) (Crum et al., 2005). The resulting gradient images are shown in Figure 4.6. Although theoretically the result is a general image similarity measure (as discussed in Chapter 2), gradients are seen that will shrink enhancing features.

$$NMI = \frac{H_A + H_F}{H_{AF}} = \frac{MI + JE}{JE} \quad (4.5)$$

$$\frac{dNMI}{dx} = -\frac{1}{JE^2} \left[JE \frac{dMI}{dx} - MI \frac{dJE}{dx} \right] \quad (4.6)$$

4.2.3 Conclusion

In this section we have demonstrated that both statistical and information based cost-function gradients are affected by contrast-enhancement. The extent of enhancement shown here is likely to be unrepresentative of an entire dataset since the examples show the first passage of contrast agent through the heart. However, over the course of a dynamic series, the effects shown above will be manifest at different levels. One solution is to separate the motion artefacts from contrast-enhancement artefacts. In the absence of a cost-function that implicitly does this, it is necessary to focus on the formation of images that are contrast-matched; where some effort is made to match enhancing features between float and anchor images. Therefore the effect of contrast-enhancement can be removed from the resulting cost function gradients. The discussion in the next chapter of the Progressive Principal Component Registration method discusses one such way in which contrast-enhancement matched images can be generated.

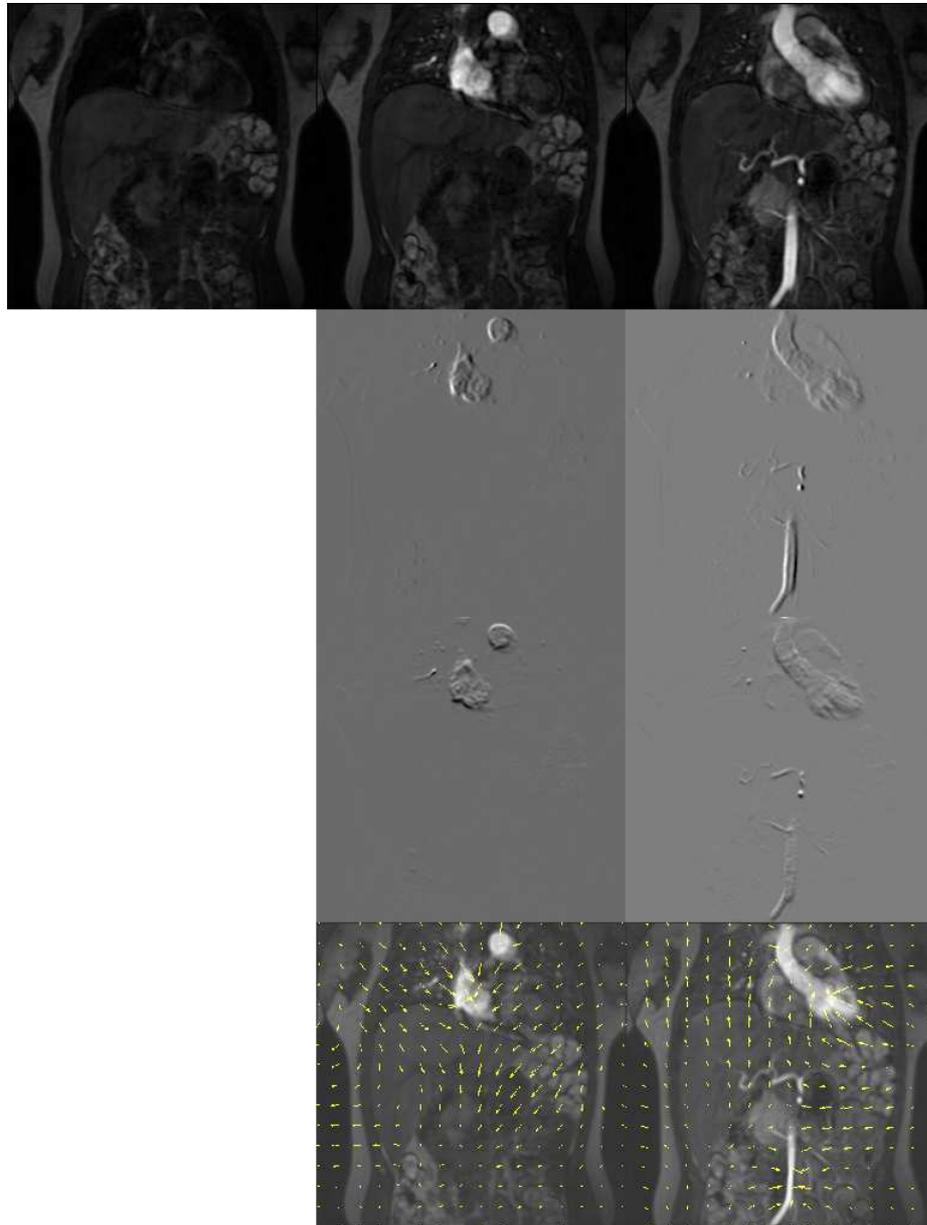


Figure 4.2: Derived image-similarity local gradient images for Least Squares Cost Function. **Top row** pre-enhancement image and two post-enhancement images. **Second row** corresponding (normalised) force gradients in x -direction. **Third row** corresponding (normalised) force gradients in y -direction. **Bottom row** corresponding displacement vector fields.

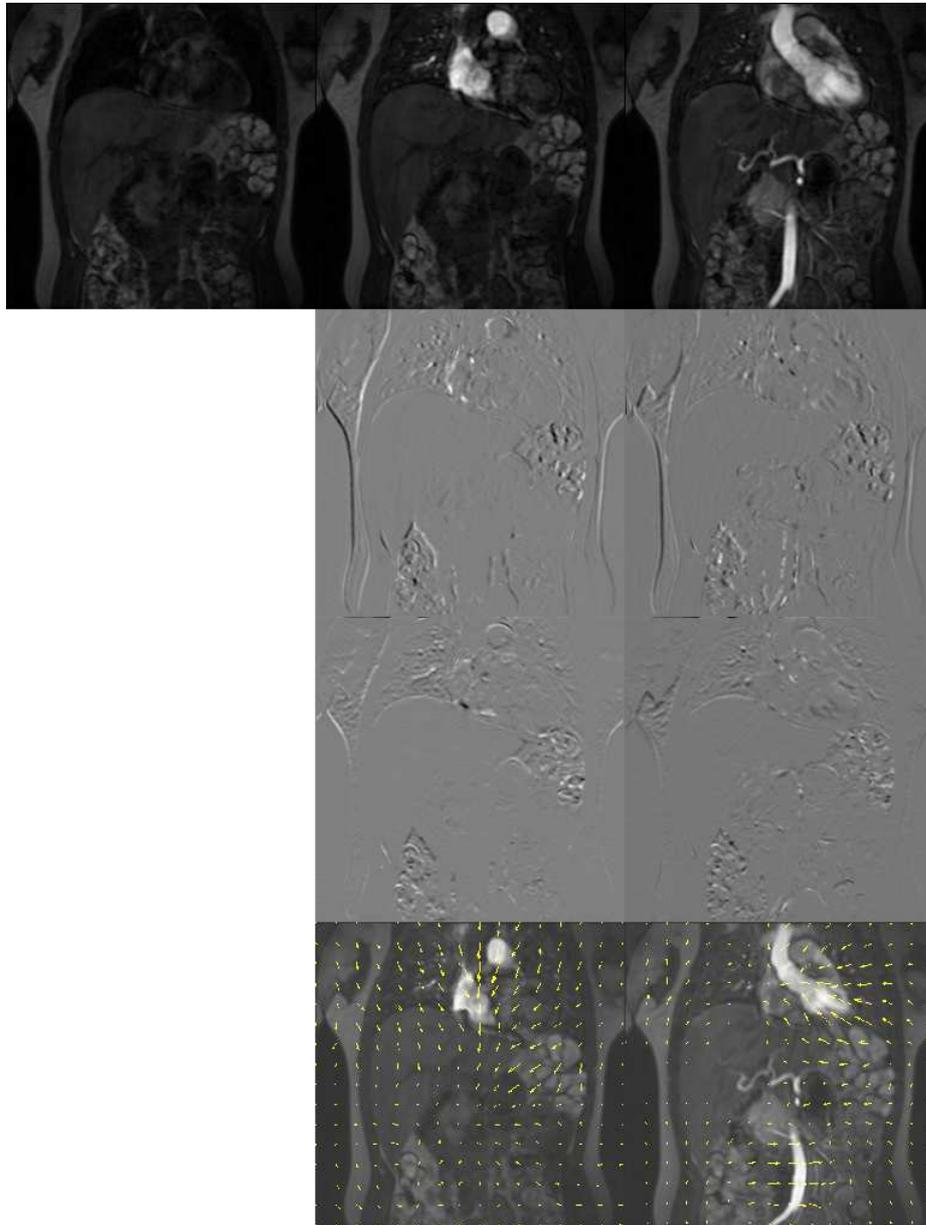


Figure 4.3: As Figure 4.2 but for Cross-Correlation Cost Function.

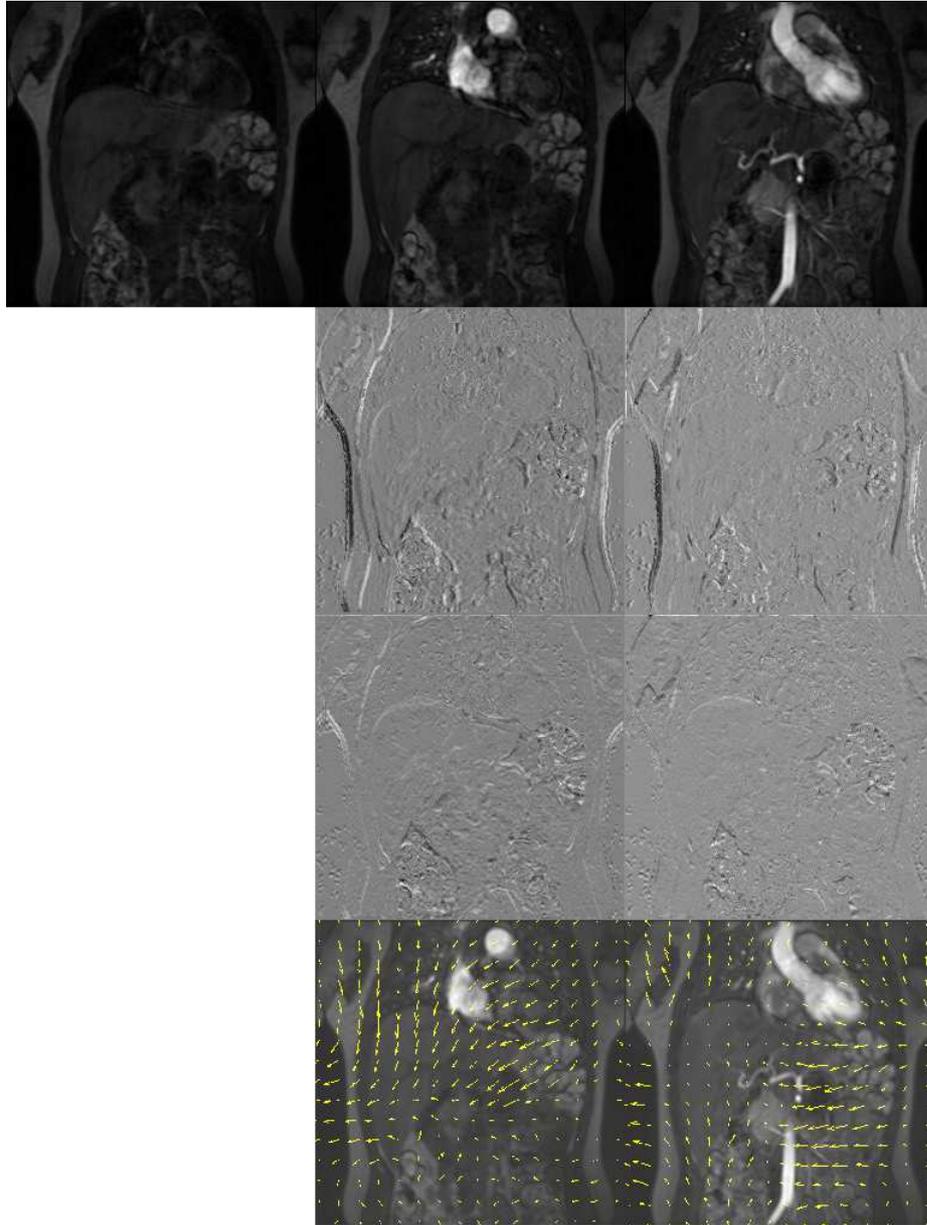


Figure 4.4: As Figure 4.2 but for Joint Entropy Cost Function.

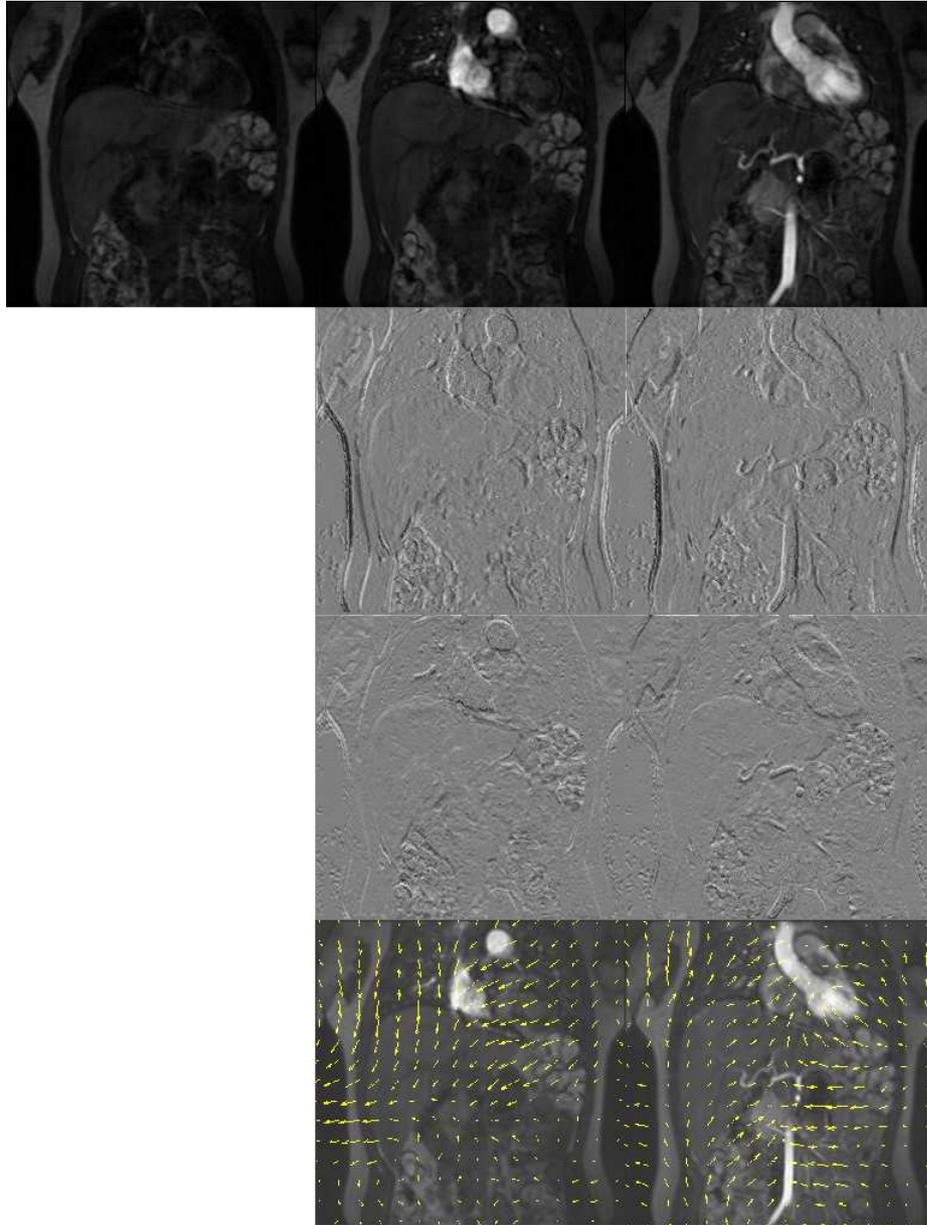


Figure 4.5: As Figure 4.2 but for Mutual Information Cost Function.

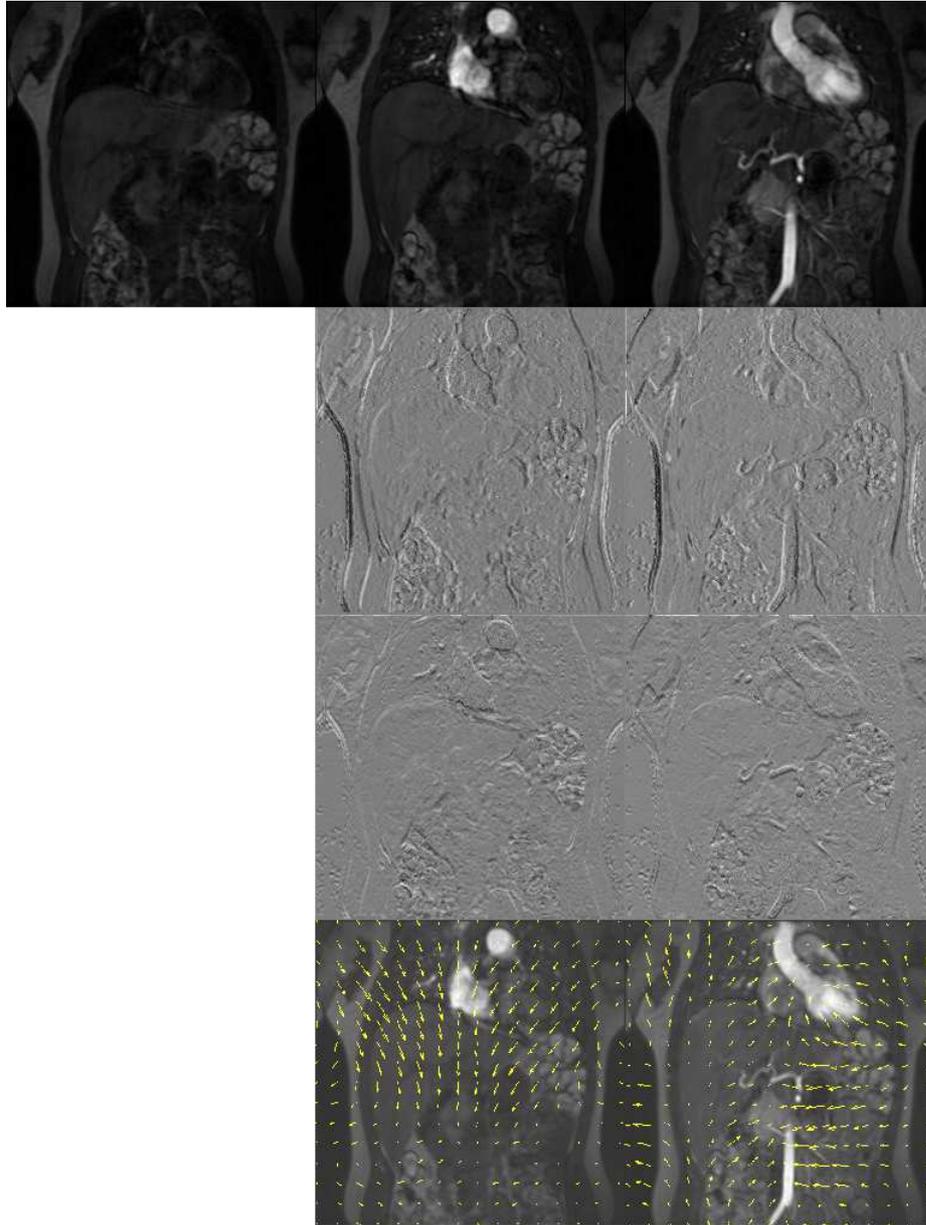


Figure 4.6: As Figure 4.2 but for Normalised Mutual Information Cost Function.

4.3 Cost Function Minimisation

We can test the suitability of different cost-functions on different types of image by progressively applying a known force with the simulation discussed in Chapter 3 and comparing the original and deformed images. In this way we can assess the ability of a given cost-function to retrieve the true deformation parameters which in this case will be a force with a magnitude and (2D) direction. If the cost-function is appropriate and therefore produces a monotonically decreasing path towards the true solution, we would expect to be able to take any (gradient-descent) trajectory through the cost-function space to find the parameters that were used to produce the second image from the original image. The method is discussed in this section as a prelude to further use in Chapter 5.

4.3.1 Method

The experiment proceeds as follows. Two images (which should be in good feature alignment) are considered, for instance two identical images. One image is reserved as the anchor whilst the other image, the float image, is deformed by a known force. In the case shown here a known force is applied to the centre of the image (Figure 4.7). The force is varied in magnitude and direction to deform the float image. The deformation of the float image is found by solving the linear elastic equation for the displacement from the force as discussed in Chapter 3. The cost-function value between anchor image and deformed float image is recorded and plotted in a space corresponding to force magnitude in each direction. The intensity in the images of Figure 4.8 and 4.10 represent the cost function value for a given value of force magnitude and direction corresponding to the x and y axes. The centre of each cost-function space corresponds to zero force, which for identical images corresponds to a perfect image alignment. If we were to use the cost-function space information in a registration algorithm, we would follow the gradient of the cost-function to its minimum and therefore register the images.

The cost-function spaces provide evidence for how well registration would proceed. If the cost-function space has a well defined minimum (or maximum) we would expect an appropriate registration algorithm (in this case perhaps an elastic registration algorithm) to deform the image to achieve this minimum. On the other hand, if the cost-function space has no clear minimum, the registration might be expected to proceed poorly.

In the cases shown here, the images are deformed with a single force causing an elastic deformation in the centre of the image. An elastic registration algorithm with this prior knowledge could apply a single force to the centre of the image and optimise the value accordingly. However, in realistic image registration examples we do not know the type of deformation that brought about the changes to the image, nor do we expect the deformation to have a simple form; we can only assume that using a fluid or b-spline registration (or a good regularisation) will result in finding a good approximation to the true deformation parameters.

4.3.2 Results

Two identical images

Figure 4.8 demonstrates the minimisation space of the respective cost-functions for varying the magnitude and direction of a force applied to an image. For two identical images (Figure 4.7), the search space is particularly well-defined for all cost-functions. For ease of visualisation the negative log values for mutual information, normalised mutual information and cross-correlation are shown, hence the best value for the cost function is bright in all images. It appears that all cost-functions may be used in this case, with the exception of differences in the speed at which they are optimised under gradient-descent.

Contrast-enhanced images

If we inspect the cost-function space of a pre and a post contrast enhanced image (e.g. one in which there are valid one-to-many pixel relationships between anchor and float images (Figure 4.9)), we get a much less well-defined minimisation (Figure 4.10). Again, for ease of visualisation the negative log values for mutual information, normalised mutual information and cross-correlation are shown.

4.3.3 Conclusion

The results in this section provide the groundwork for Section 5.4. The results presented here use a particularly simple form of deformation, applying a single force to the centre of the image. Figure 4.10 provides evidence that registration using sum-of-squared differences or cross-correlation will lead to erroneous registration when enhancement is present - in this case a compression of the enhancing region (this is due to the compression optimising the cost-function by removing enhancing pixels). For the very simple deformation model described in this section, the information based cost-function minimisation spaces appear to suggest that they are able to correct the deforming displacement. However, with regards to full non-rigid registration, the cost-function force-field is calculated locally across the entire image and we may not expect registration to minimise so well. The global form of the deformation is perhaps oversimplistic when compared to the inverse registration problem. Non-rigid image registration algorithms are local in application, so local changes are likely to impinge on registration success using information-theoretic cost-functions, resulting in the effects seen in Figures 4.4 to 4.6.



Figure 4.7: Two identical images **a** and **b**, float image is deformed by a known force and then the cost-function value is found between deformed float and anchor. **c**: difference image between **a** and **b**. **d**: difference image between **a** and **b** at maximum deformation.

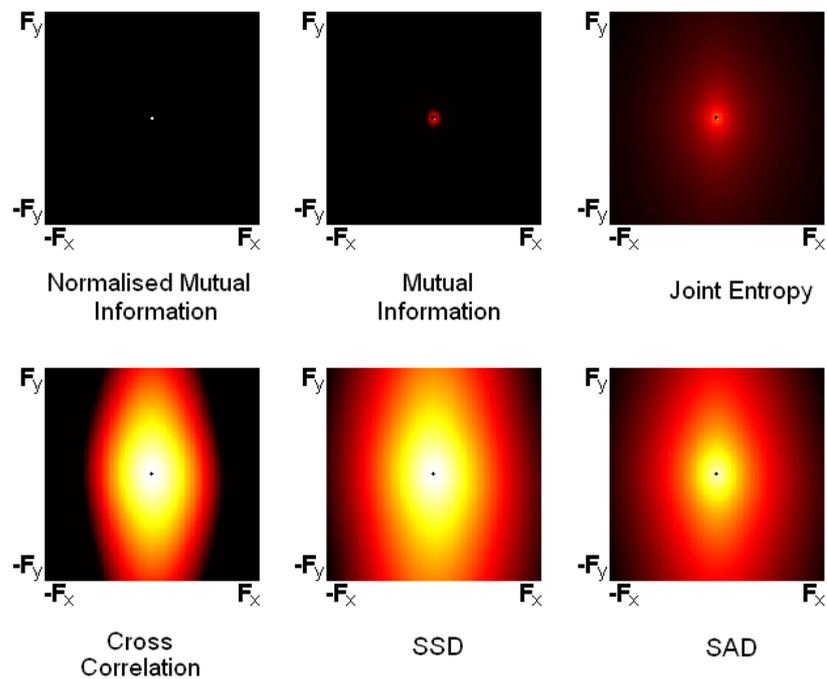


Figure 4.8: Cost-function minimisation space for labelled similarity measure for images corresponding to Figure 4.7, identical images where the float is deformed by a known force of varying magnitude and direction. Distance from centre on X and Y axes represents force strength in that direction (zero force at centre).

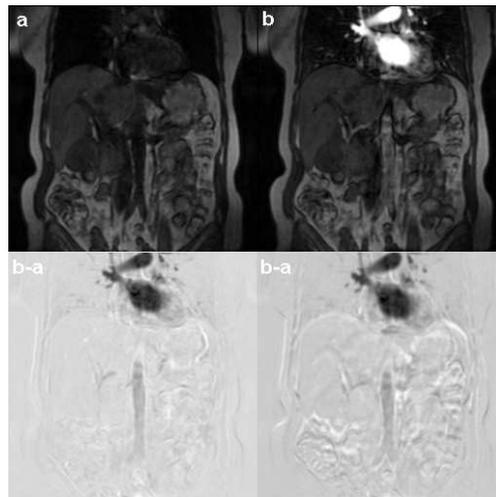


Figure 4.9: Pre and Post Contrast Enhancement images **a** and **b**, float image is deformed by a known force and then the cost-function value is found between deformed float and anchor. **c**: difference image between **a** and **b** with no deformation (there is small existing misalignment). **d**: difference image between **a** and **b** at maximum deformation.

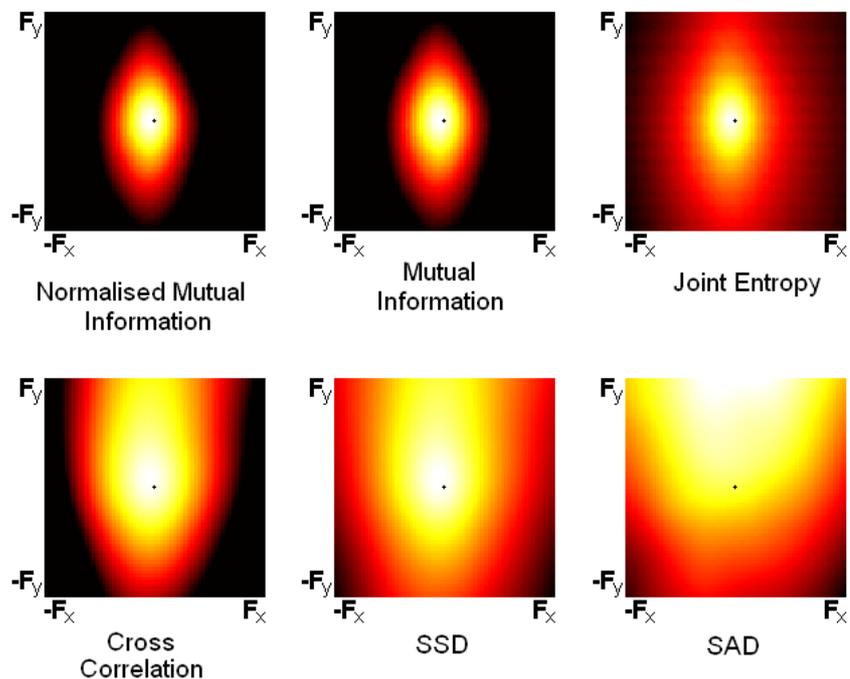


Figure 4.10: Cost-function minimisation space for labeled similarity measure for images corresponding to Figure 4.9, pre and post enhancement images where the float is deformed by a known force of varying magnitude and direction. Distance from centre on X and Y axes represents force strength in that direction (zero force at centre).

4.4 The Cost Function Matrix Mean (CFMM)

Assessment of the performance of the registration can be done comparatively between two images before and after registration. This is also true when registering groups of images. When registering a group of images we desire that they are all registered to the same coordinate system, but we not only need to assess how well they are registered to one image, but how well they are registered to every other image in the group. Therefore, we need to assess the relative improvement of a *cost-function matrix* of cost-function values of each image in a group to every other. The formulation of the mean value of this matrix is shown in Equation 4.7 for a symmetric similarity measure for T images (this fills half the matrix; a non-symmetric cost-function would fill every entry in the matrix). An explicit example is shown in Equation 4.8, again for T images for the Normalised Mutual Information cost-function.

$$\widehat{S} = \frac{2}{T(T+1)} \sum_{i=1}^T \sum_{j=1}^i \text{cost}(A(i), F(j)) \quad (4.7)$$

$$\widehat{NMI} = \frac{2}{T(T+1)} \sum_{i=1}^T \sum_{j=1}^i \frac{H_{A(i)} + H_{F(j)}}{H_{A(i)F(j)}} \quad (4.8)$$

$$(4.9)$$

If our group of images is well-registered relative to an unregistered set of images, we would expect an improvement throughout the cost-function matrix since every image should be better aligned to every other. Therefore we might expect the total gain in image similarity from the image registration can be analysed by assessing the mean of the cost-function matrix (The cost-function matrix mean (CFMM)). This method also allows a comparison of the relative merits of different registration algorithms. Other measures of the change to the cost-function matrix after image registration might also be proposed but for this work, we consider the mean value of the matrix elements only.

4.4.1 Using Simulated DCE-MRI to investigate the CFMM

The behaviour of the Cost-Function Matrix Mean (CFMM) under varying influence of motion and contrast enhancement may be investigated with simulated DCE-MRI data produced from the method in Chapter 3. By setting a standard deformation and segmentation, we may then vary the force magnitude and contrast-agent 'dose' parameters used in this standard deformation and investigate the stability of the CFMM. We take one image and generate a dataset of 20 images with varying motion strength and contrast agent dose parameters. Motion strength is varied in 9 steps from zero motion to an average pixel displacement of 1.13 pixels and a corresponding maximum displacement of 33.6 pixels. Enhancement strength is also varied in 9 steps from zero enhancement to an increase of 150% in the region of greatest enhancement. Simulated deformation and enhancement are applied to a coronal liver image; a dominant force direction moves the liver in a superior-inferior direction with a sinusoidal motion. Datasets are generated, each of twenty images, varying the motion and enhancement parameters in nine steps between zero displacement (and enhancement) and twice the standard displacement (or contrast agent dose). The NMI-CFMM is then calculated for each of the 81 (9 motion \times 9 enhancement steps) datasets.

Figure 4.11 shows the results of the NMI-CFMM for each of the nine motion and nine enhancement levels described above. The influence of motion level dominates the curve position and the influence of

contrast-enhancement introduces a small dose-dependent perturbation to each motion-level curve, seen as a decrease in the similarity measure with increasing dose. The top curve representing zero-motion contains the effects of contrast-agent only and hence decreases slowly from the maximum value of 2 with increasing enhancement level.

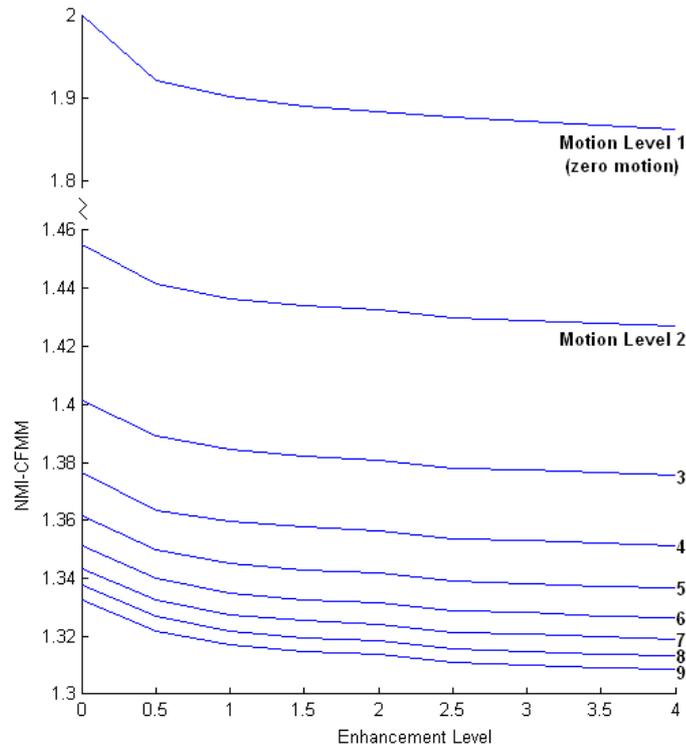


Figure 4.11: Plot of NMI-CFMM values for varying sinusoidal-motion magnitude in linear steps between minimum and maximum force strength and contrast agent 'dose' varied in linear steps before conversion to signal by the spoiled gradient echo equation (Equation 2.26).

The results above suggest that the CFMM measure is a suitable measure for determining the performance of registration on groups of images, dependent on the appropriateness of the measure used on each pair of images. The value of the NMI-CFMM is determined predominantly by the motion parameter, causing the large jumps between curves of NMI-CFMM value for varying enhancement, with additional perturbations due to contrast-enhancement. The method may provide some robustness to contrast-enhancement, particularly when large numbers of images are in the wash-out phase. When comparing the alignment of separate groups of images, care must be taken when using NMI, since the measure is non-linear between NMI value and probabilistic image similarity as seen in Figure 4.11. The cost-function matrix may also be analysed by the standard deviation of its values. In this sense, a reduction in variability of the matrix values corresponds to good overall registration. In future work, this statistic could be used to reveal cases where the cost-function matrix mean is biased by a few very good or bad registrations.

4.5 Choice of Anchor Image in Conventional Registration

An important consideration in conventional registration of DCE-MRI is the choice of anchor image. This provides the co-ordinate system for all images. The cost function matrix mean (CFMM) assessment discussed in Section 4.4 allows a comparison of the results of selecting different anchor images in conventional registration. It represents a measure of internal registration consistency - perfect registration of identical images would result in a maximum (or minimum) value of the cost-function matrix mean. Figure 4.12 shows the final NMI-CFMM for real datasets from Table 1.1 registered using each image in turn as the anchor: a higher value of the NMI-CFMM represents a better registration of all images within the dataset to one another. Results are shown for individual registrations using both cross-correlation and NMI as the registration cost-function. An interesting point for each entry in this graph, is that for individual registrations using NMI as the image similarity measure, the final NMI-CFMM is lower than (not as good as) that found using cross-correlation for the individual registrations. Visual inspection of the individual registered images reveals that cross-correlation is more likely to give a better result, improving the overlap of image features such as the diaphragm. This result is used as a justification for using cross-correlation in later work.

As a further note, it is also possible to register all images to the mean image. However, the reduction in image resolution by the summation of the original image intensities may result in either under-registration, due to features being aligned to the same position in the mean-image, or mis-registration of features. Under-registration is likely due to the formation of the cost function image forces: the loss of definition in the pixel intensity mappings may cause the driving forces to be weaker; hence it should be preferable to register to real features in a carefully selected anchor image.

The results of Figure 4.12 demonstrate the difficulty of selecting the best anchor image from the dataset. In three cases, the anchor image resulting in the highest (best) NMI matrix mean is one of the pre-enhancement images. However, selection of this optimum pre-enhancement anchor image requires proceeding with the registrations for all other anchor images. For this reason, subsequent conventional registrations presented here use the first image in the time-series as the anchor image. Outlier values of the NMI-CFMM plots are likely due to registrations carried out toward an anchor image that is itself an outlier (i.e. contains large motion deformation relative to the other images in the dataset).

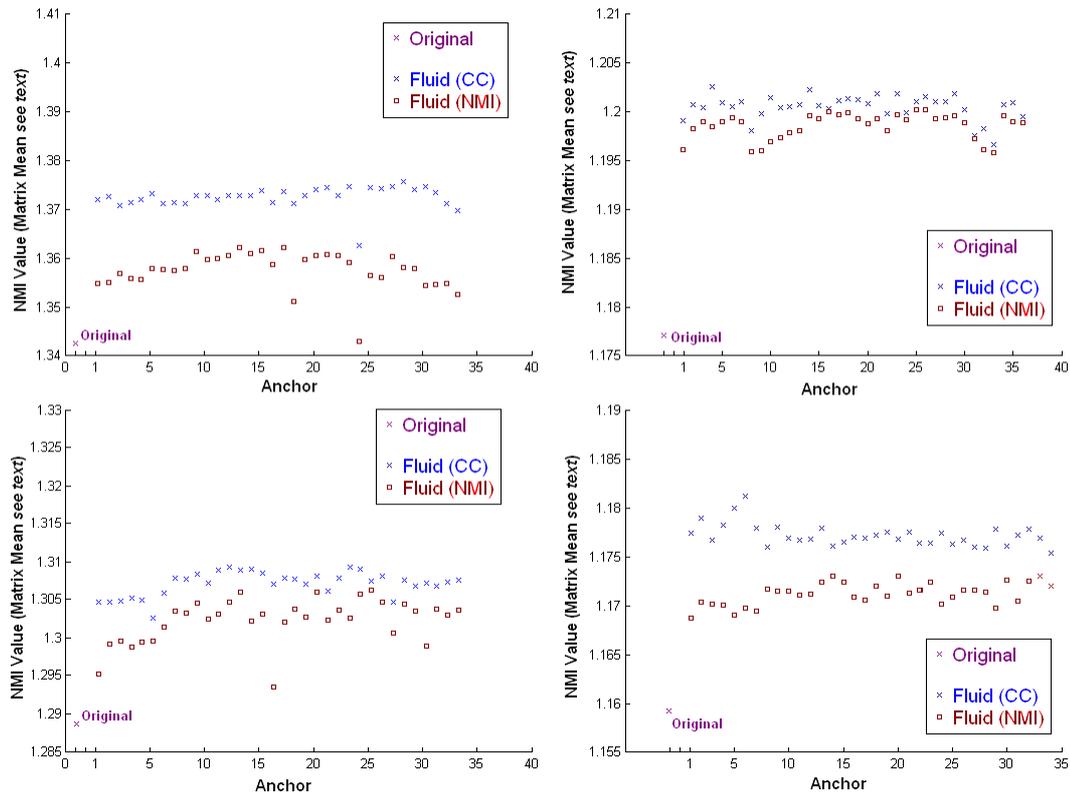


Figure 4.12: Assessment of registration result of variations in target image selection using Cost Function Matrix Mean (see Section 4.4). For four separate datasets from Table 1.1, fluid registration proceeds using the n th image as the anchor image. The NMI Matrix Mean is shown for the result of registration using Cross-Correlation (blue) and Normalised Mutual Information (red). Also shown is the original NMI Matrix Mean before registration (magenta).

4.6 Conclusion

This chapter has provided examples and discussed the reasons for poor image registration of contrast-enhanced images. The reason for poor image registration is the cost-function; each of the cost-functions demonstrated here prefer a one-to-one relationship between image intensities in order to operate successfully. The effect of contrast-enhancement, particularly the large intensity changes induced by bolus arrival, on image registration force-gradients is likely to cause mis-registration as demonstrated by the compression of the enhancing features in Figures 4.2 to 4.6.

Section 4.3 investigated the disruption to the cost-function minimisation space for a range of cost-functions brought about by contrast-enhancement. Despite a relatively simple elastic deformation being used, the effect on the cost-function space is shown to affect the potential for good minimisation.

The formation of the cost-function matrix mean (CFMM) is presented as a method of inspecting the overall registration status of a group of images. Despite the problems of image registration in the presence of contrast enhancement, the use of this measure as a post-registration measure of registration success is still possible. Providing the images have been registered by an algorithm robust to contrast-enhancement intensity changes, it is possible to cautiously (given the relationship of the cost-function to the actual image deformation) compare results of the CFMM between registration methods and this analysis will be used in future chapters. The method is first used in Section 4.5 to investigate the choice of an optimal anchor image for the registration of a group of images. The results show no preference for anchor choice, although registration to any anchor image is likely to provide some benefit over the unregistered data. In future chapters, when testing against algorithms robust to contrast-enhancement, unmodified direct fluid registration will always proceed using cross-correlation to align each image to the first image in the dataset.

Chapter 5

Progressive Principal Component Registration (PPCR)

In this chapter we present the development of a progressive, temporal principal-component based registration algorithm (PPCR). The model developed in Chapter 3 is used to explore the interplay between motion type, the extent of organ motion and contrast-enhancement on PPCR performance. Further testing is carried out to evaluate the performance of the PPCR algorithm on real Dynamic Contrast Enhanced MRI data (DCE-MRI). The model of DCE MRI of the liver from Chapter 3, incorporates an isotropic elastic non-rigid deformation to simulate both breathing and breath-hold data, a volume-preserving modification for tumour regions is also included. Contrast enhancement is simulated by applying a pharmacokinetic model. In this chapter, for each simulated dataset, a direct fluid registration of each image to the first in the dataset is compared to the contrast-enhancement guided Progressive Principal Component Registration (PPCR). Analysis of the correction to the deformation fields, tumour volume change and dispersion of joint image histograms are used to show the importance of motion type on PPCR performance and of enhancement level on direct fluid registration performance. For breathing motion, we will see that PPCR registers groups of images in different phases of the breathing cycle to separate final positions, but maintains enhancing tumour volume. This is not the case for direct registration where volume changes of up to 7% are observed. For cases in which the patient holds their breath at different levels of expiration, PPCR out-performs direct registration, particularly for large enhancement levels. Analysis of the joint image histograms suggests that the generation of target images using PPCR reduces histogram dispersion due to contrast enhancement. Since this distinction is not made using direct registration, it is unable to register images when large enhancements are present. On the other hand, under cyclic breathing motion, PPCR target images are ill-defined, increasing dispersion in the joint image histograms, leading to failure or separation of the images into clusters driven by breathing phase. Also analysed are the effect of more careful choice of anchor image in conventional image registration and an investigation of the progress of PPCR with each iteration. Analysis of the formation of images in the first PPCR iteration is also carried out on real DCE-MRI data from Section 1.3 using the method developed in Section 4.3.

5.1 Principal Components Analysis

Principal Component Analysis (PCA) is a method of representing data in a coordinate system so that the maximum data variance occurs along the first axis and the second variance component along the second axis and so on. The ordering of variance allows components that represent a large fraction of the variance of the data to be considered separately from those that contain little variance. This makes it useful both for eliciting trends from data and in compressing data. Both approaches may be useful for the registration of DCE-MRI but it is the data compression, applied in the time domain, that is used here.

To calculate principal components we must reinterpret our data so that every pixel in an image is described by a function detailing its change in intensity through the dynamic sequence. We define the dataset as \mathbf{A} , each individual time-frame must have the mean pixel value of that time-frame subtracted from it, as required by the PCA. An individual pixel-function located at spatial index i, j is denoted \mathbf{A}_{ij} and since it consists of T time points, is a vector of dimension $T \times 1$. We now compare all pixels in the dataset, obtaining a covariance matrix of size $T \times T$ encapsulating information from every pixel in the dataset. Finding the eigenvectors of this matrix and ordering them by eigenvalue magnitude, we obtain the PCA result.

Calculating the Principal Components Analysis of a DCE-MRI dataset in the manner described above can be used to generate registration anchor images that are contrast-matched to their respective float images. Crudely, the first principal component will resemble the general intensity profile of the images with respect to one another. Hence, all pixels can be weighted with how much of this principal component they contain. Differences due to organ motion (in particular organ motion due to inconsistent breath-hold depth) are not strongly represented in the first principal component because this motion is represented by local intensity fluctuations in relatively few pixels; the variance that this motion represents is likely to be small and hence is likely to appear in later principal components. Hence the large scale intensity changes are dominant in the generated anchor images; the anchor images are contrast-matched to their float images and so registration by conventional cost-functions becomes feasible. In subsequent iterations, previous registrations have hopefully removed some of the organ motion, and so principal components increasingly contain information about changing intensity profile in preference to residual motion artefacts.

5.1.1 PCA for Functional Analysis

It is often claimed that combinations of Principal Components can be used to represent physical or biological information within a dataset. Principal Components Analysis is often used in statistical shape analysis to describe the principal axis along which shapes vary - for instance in our laboratory, the femur's principal direction of variation is along its length with subsequent important variations in femoral head size and axial twist (Chan et al., 2004). Extensive work on statistical shape models incorporating principal components analysis has also been produced by Cootes (Cootes et al., 2008).

Figure 5.1 shows the first four principal component eigenvectors calculated from a Dynamic Contrast Enhanced MRI dataset of the abdomen. As described above, each pixel is treated as a data entry for the principal components calculation. In this case we have 40 timepoints and so we have 40 principal

component eigenvectors each of 40 timepoints in length, the first four of which are shown here. The first principal component corresponds to the most representative pixel intensity fluctuations (most variance will be contained in a principal component that represents the overall increase in intensity with bolus arrival and dispersion) around which further fluctuations are modulated by later principal components. To try to extract physiological parameters relating to contrast uptake and to try to infer K^{trans} , v_e or v_p would be extremely difficult due to the averaging of pixel intensity information through the covariance matrix; the principal components are also orthogonal whereas any physical parameters may not be. The PCA provides an efficient re-parameterisation of the data, but there is no immediate reason for this parameterisation to be better at yielding pharmacokinetic parameters.

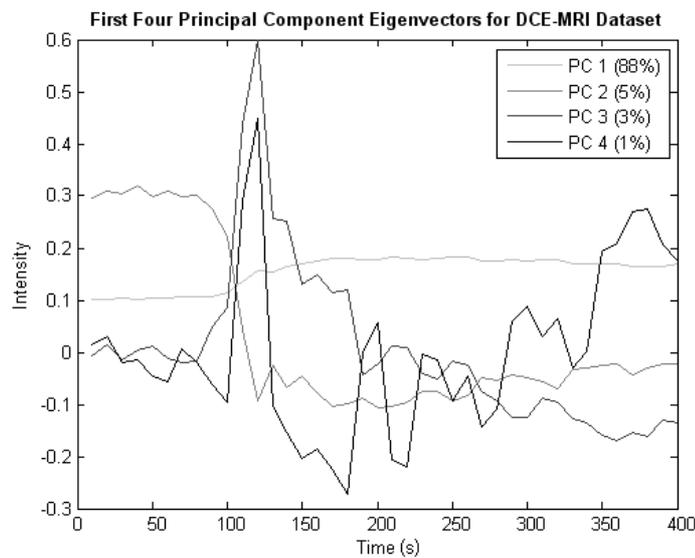


Figure 5.1: First four normalised principal components for DCE-MRI dataset. PC1 is a general enhancement profile incorporating the mean intensity change over a time-scale of a few minutes. PC2 appears to act to correct those pixels that are not enhancing. PC3 and PC4 appear to enhance PC1 in areas of rapid initial enhancement, further describing differences between pixels in the wash out phase. Extracting pharmacokinetic parameters from these components is likely to be difficult despite the fact that these four components contain 97% of the dataset variance.

5.1.2 PCA Used for Data Compression

Principal Components Analysis can be used for lossy data compression, instead of transmitting an entire DCE-MRI dataset we could in principle transmit only a few early principal components and their weightings. Figure 5.2 illustrates the resulting data-compressed images. Slices are shown for pre-enhancement, bolus arrival in the left heart, bolus arrival in the liver and late post-enhancement. The top row contains the original images and subsequent rows contain the images rebuilt from 1,2,3,4 principal components respectively (see Figure 5.1). It is clear that in this case, for a relatively small number of time-points (20), that the early principal components are dominated by explaining the early enhancement of the heart and aorta. As a result the pre-enhancement heart is not well represented until four principal components

are used.

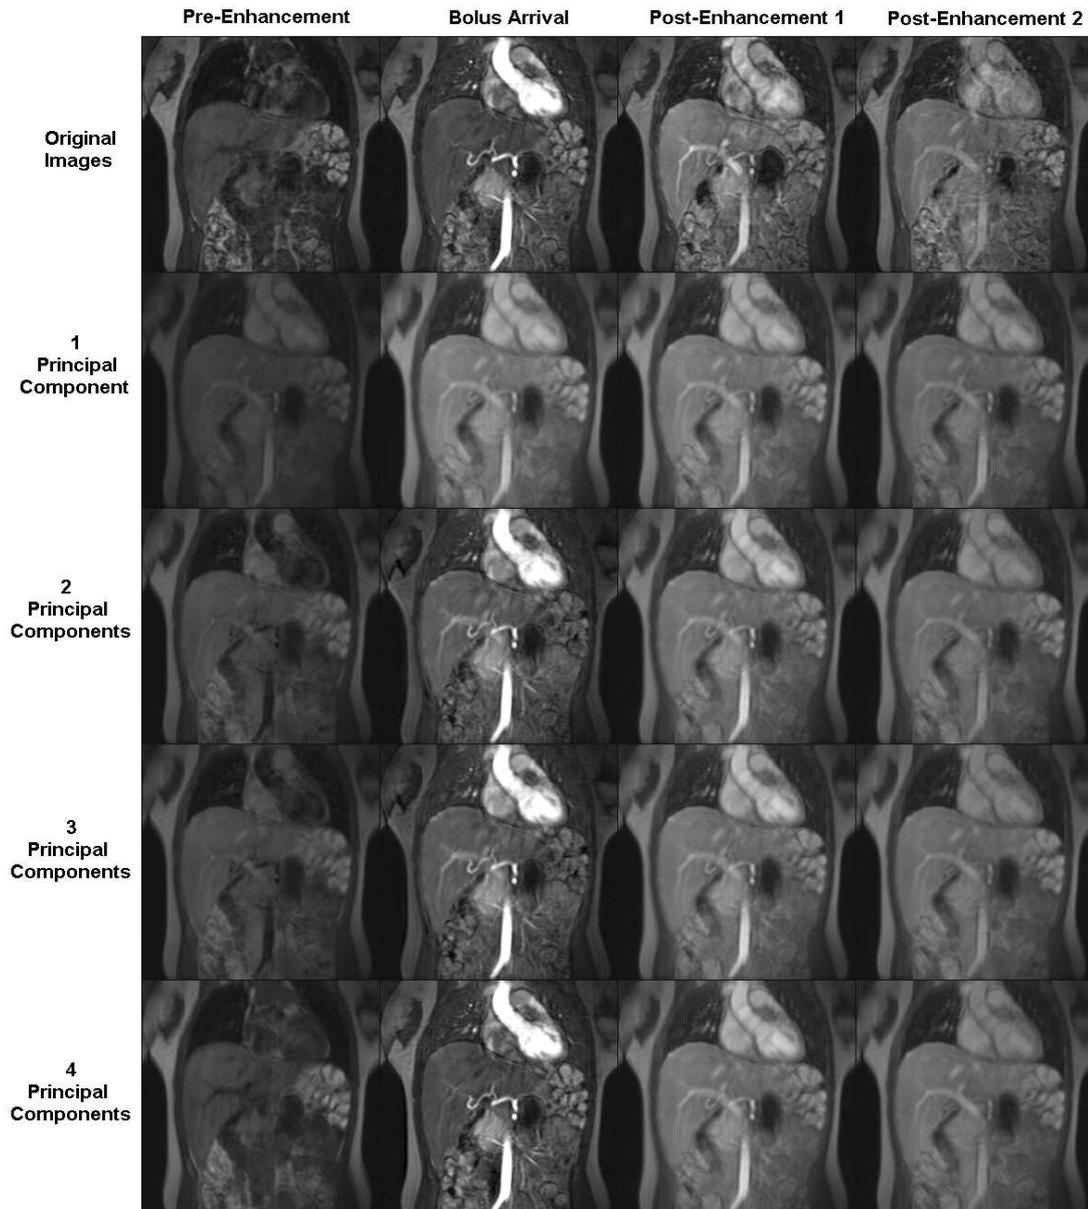


Figure 5.2: **Top row:** Images from a DCE-MRI dataset for pre-enhancement, bolus arrival in the left heart, bolus arrival in the liver and late post-enhancement. **Subsequent Rows:** Images rebuilt using, on each row, 1,2,3,4 components respectively.

5.1.3 Formation of the Covariance Matrix

The covariance matrix from which principal components are calculated is governed by the strength of temporal relationships between pixels. The mean intensity across the image will vary, but this will not contribute to the covariance values since it is subtracted. Regions that enhance will contribute, due to changes in pixel variance. Random noise is not influenced by time-point and will appear uniformly over the covariance matrix, but time-dependent pixel trends bias the covariance matrix. This bias influences

the resulting eigenvectors and eigenvalues of the covariance matrix. As a result, strong pixel time-point trends result in large eigenvalues combined with an eigenvector reflecting this trend. Ordering the eigenvectors by eigenvalue magnitude sifts long-temporal pixel trends from short-temporal random noise. If motion appears random and of similar magnitude through all time-points, it is conceivable that image motion artefacts will appear in principal components with small magnitude eigenvalues. Conversely, pixels undergoing different enhancement profiles will have those profiles encapsulated in combinations of the first few principal components although we do not expect to be able to differentiate between different underlying physiological profiles.

5.2 The PPCR Algorithm

This section develops the Progressive Principal Component Registration (PPCR) process. The method applies registration repeatedly to an artificial time-series of target images generated using the principal components of the current best-registered time-series data. The aim is to produce a dataset that has had random motion artefacts removed but long-term contrast-enhancement implicitly preserved. The method requires neither segmentation nor a pharmacokinetic uptake model and can allow successful registration in the presence of contrast-enhancement.

The PCA produces T eigenvectors \mathbf{U}_m , each $T \times 1$ in size where $m = [1, 2, \dots, T]$. We can write our data in terms of these eigenvectors. Each pixel has an amount $C_{ij,m}$ of eigenvector \mathbf{U}_m . $C_{ij,m}$ is the dot product of \mathbf{U}_m and the original pixel \mathbf{A}_{ij} . Hence the original pixel, \mathbf{A}_{ij} , located at index i, j can be re-built as in equations (5.1) and (5.2).

$$\mathbf{A}_{ij} = \sum_{m=1}^T C_{ij,m} \mathbf{U}_m \quad (5.1)$$

$$\mathbf{A}_{ij} = \sum_{m=1}^T (\mathbf{U}_m \cdot \mathbf{A}_{ij}) \mathbf{U}_m \quad (5.2)$$

Our assumption is that most signal associated with enhancement is contained in the earlier eigenvectors of the PCA. Hence we can approximate a pixel's intensity values \mathbf{A}_{ij} by leaving out the less significant eigenvectors.

$$\mathbf{A}_{ij} \approx \sum_{m=1}^{\eta < T} (\mathbf{U}_m \cdot \mathbf{A}_{ij}) \mathbf{U}_m \quad (5.3)$$

PCA extracts trends from the time-series data in order of significance. As a result long-term contrast-agent uptake trends should appear in the earlier principal component eigenvectors, whereas any short-term random motion would be represented in later principal components. If data are rebuilt from only the first few eigenvectors, we should be able to register to essentially motion-free data. Having done this we are free to repeat the process. The initial synthetic dataset is given by (5.4), where the star-superscript represents PCA generated data and we include a numerical superscript to denote iteration number: we write $\mathbf{A}_{ij}^1 \equiv \mathbf{A}_{ij}$, since this will be our first iteration. Eigenvectors are also given a superscript denoting the iteration in which they were calculated (e.g. \mathbf{U}_m^n for iteration n).

$$\mathbf{A}_{ij}^{1*} = (\mathbf{U}_1^1 \cdot \mathbf{A}_{ij}^1) \mathbf{U}_1^1 \quad (5.4)$$

$$\mathbf{A}^2 = \mathbf{A}^1 \mapsto \mathbf{A}^{1*} \quad (5.5)$$

We register each image in the original data, \mathbf{A}^1 , to its corresponding image in the PCA generated, data, \mathbf{A}^{1*} (see (5.5) where \mapsto represents image registration), using a suitable registration algorithm. This will result in a dataset for the second iteration, \mathbf{A}^2 , that is coarsely registered. Repeating these steps we can re-calculate the PCA on this coarse-registered data, \mathbf{A}^2 , giving us a new set of principal component eigenvectors \mathbf{U}_k^2 (where the superscript denotes that this is the second iteration). Data are rebuilt from both the first and second eigenvectors from the new PCA (5.6). We are free to do this because the coarse registered data has less motion than the original data, so motion artefacts should appear even later in

the principal component eigenvectors. We now register our coarsely registered data, \mathbf{A}^2 , to the dataset produced by (5.6), \mathbf{A}^{2*} , to find \mathbf{A}^3 (5.7).

$$\mathbf{A}_{ij}^{2*} = (\mathbf{U}_1^2 \cdot \mathbf{A}_{ij}^2) \mathbf{U}_1^2 + (\mathbf{U}_2^2 \cdot \mathbf{A}_{ij}^2) \mathbf{U}_2^2 \quad (5.6)$$

$$\mathbf{A}^3 = \mathbf{A}^2 \mapsto \mathbf{A}^{2*} \quad (5.7)$$

By repeating this process from $n = [1, 2, \dots, T - 1]$ we aim to achieve well-registered data (5.8). The registration for $n = T$ would be a registration of the registered dataset to itself, as in (5.2), and so is omitted. This whole process we term Progressive Principal Component Registration (PPCR).

$$\mathbf{A}_{ij}^{n+1} = \mathbf{A}_{ij}^n \mapsto \sum_{m=1}^{n < T} (\mathbf{U}_m^n \cdot \mathbf{A}_{ij}^n) \mathbf{U}_m^n \quad (5.8)$$

At each iteration, image registrations may be implemented by any method that permits quantitative analysis of the intensity values on the images. Here, it is the transformation embedded in the registration algorithm which determines the preservation of intensity rather than the PPCR algorithm and implicit PCA. Due to the approach used here where the covariance includes whole image data, we are free to use a large region of interest. A non-rigid registration method allows the accommodation of differences in type and extent of motion.

With each PPCR iteration it is also in principle possible to vary the number and choice of principal components. The PPCR algorithm throughout this thesis proceeds by starting with the inclusion of a single principal component and adding that with the next highest variance at each iteration. If methods were developed for the inspection of principal components at each iteration, it may be possible to adaptively select groups of principal components in order to better guide the image registration procedure. An adaptive technique might also allow computational benefits. Registration to the time-series mean as compared in the work of Buonaccorsi (Buonaccorsi et al., 2005) may be imagined as a nullth PPCR prior to the addition of principal components. The PPCR algorithm will always skip this step, beginning with the inclusion of the first principal component. This step is crucial as it allows the first level of contrast enhancement matching between the current set of float and anchor images, preventing mis-registration that might otherwise occur due to contrast enhancement.

Figure 5.3 is an illustration of the PPCR algorithm applied to five images (for conciseness). The PCA is applied as described above to produce five eigenvectors and five principal component weighting maps representing the amount of each principal component needed to recreate each pixel. In the first iteration only the first principal component and weighting map are used to generate the first set of anchor images (denoted with an asterisk). After the first set of registrations, the PCA is recalculated and a new set of anchor images produced this time including both the first and second principal components and their respective weighting maps.

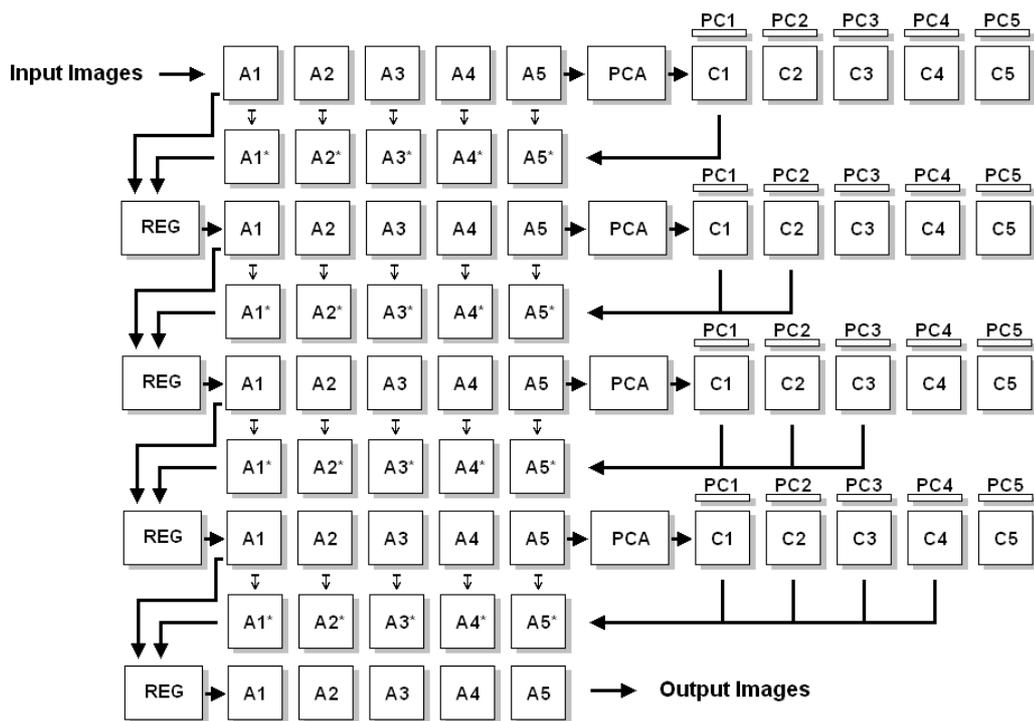


Figure 5.3: Illustration of the PPCR Algorithm for a dataset with five images. Principal Components Analysis of 5 images produces 5 Principal Components and 5 Weighting Maps. The algorithm increments the number of Principal Components used to generate target images at each iteration, recalculating the PCA after each iteration. The last (fifth) Principal Component is not used as this will result in five anchor images that are identical to their float images.

5.3 Registration of Simulated Data using PPCR

This section investigates the performance of PPCR when applied to simulated data over a range of motion types and contrast-enhancement levels. The correction of the deformation fields applied using the model in Chapter 3 are assessed alongside discussion and investigation of the formation of the joint image histograms used in registration using information theoretic similarity measures.

5.3.1 Method

Data Generation

DCE-MRI datasets are simulated for a range of motions and enhancements. Two types of motion are considered, cyclic motion due to breathing-like motion and breath-hold depth inconsistency in which liver position is determined by a Gaussian distribution around zero displacement (this simulation is consistent with the data presented in (Melbourne et al., 2007b)). A dataset is chosen with a sagittal-oblique perspective. Motion is added as a large superior-inferior force positioned in the superior liver combined with a smaller anterior-posterior force in the superior-posterior liver. This combination of deformation produces a force that changes with an elliptical pattern. Different motion levels 1-4 (level 4 corresponds to a maximum of 20 pixels (35mm) displacement) are considered. Different levels of enhancement are included, with scaling levels 0-3 (with level 3 corresponding to a maximum increase of 50% pixel intensity). Two registration types are compared, the first is a direct (fluid) registration of each image in the simulated dataset to the first image in the dataset, the second is the PPCR algorithm.

Registration Methods

Direct Image Registration Image registration proceeds by registering every image in the DCE-MRI dataset to the first (pre-enhancement) image in the dataset as summarised by Equation 5.9. A Normalised Mutual Information based cost-function may be supposed to be most able to cope with changing intensity patterns and so is used here.

$$\mathbf{A}(t) = \mathbf{A}(t) \mapsto \mathbf{A}(t_0) \quad (5.9)$$

We choose three different analysis criteria: the residual motion in the deformation fields, this is found by taking the gold-standard deforming transformation and adding to it the correcting registration displacement field. Also the tumour volume, both to assess the volume preserving constraint implemented in Chapter 3 and to observe any additional tumour volume changes due to the registration process. Finally, the joint image histograms of the unregistered and the first PPCR target images to analyse dispersion due to contrast enhancement and motion.

PPCR The registration method is shown in Equation 5.10 where the result at the next iteration, $n + 1$, is given by the registration of the best registered data, \mathbf{A} , from the previous step n , registered to artificial images generated from a temporal principal components analysis of the registered data from the previous step n , rebuilt using n principal components. Registration of source images to artificial target images uses a fluid registration algorithm (Crum et al., 2005) with a cross-correlation cost function, appropriate

for the images being registered.

$$\mathbf{A}^{n+1} = \mathbf{A}^n \mapsto \sum_{m=1}^{n < T} (\mathbf{U}_m^n \cdot \mathbf{A}^n) \mathbf{U}_m^n \quad (5.10)$$

5.3.2 Results

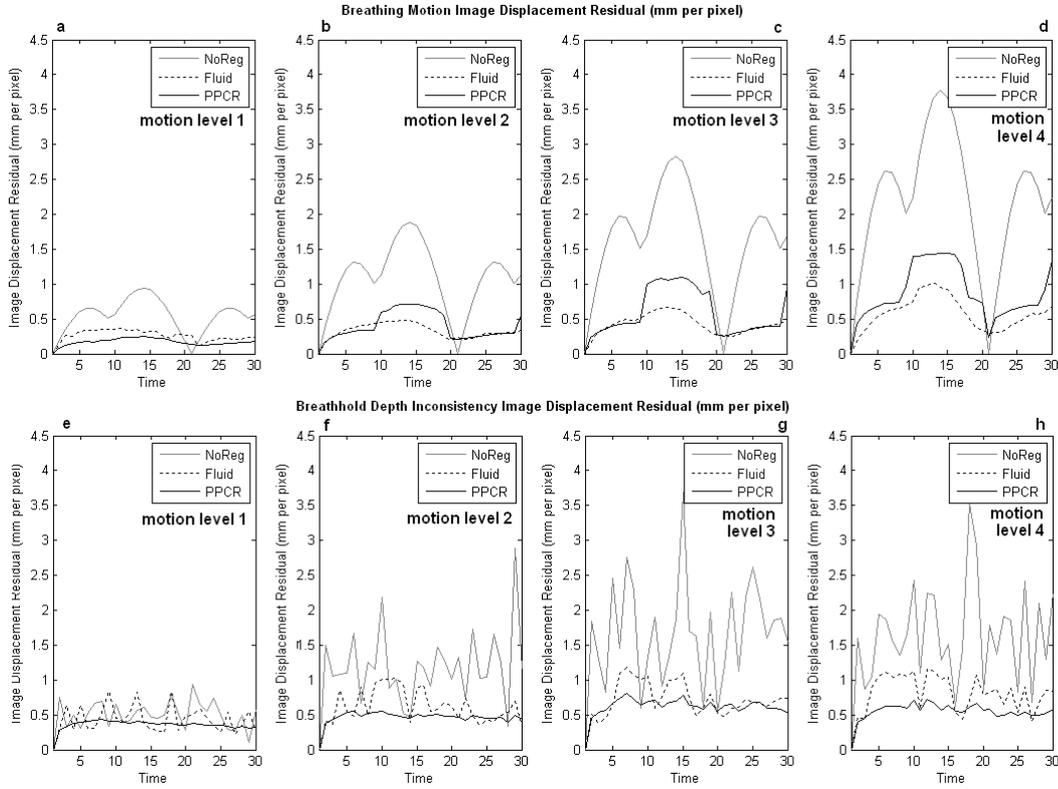


Figure 5.4: Graphs of absolute image residual displacement with time for varying motion (levels 1-4 corresponding to $9mm$, $18mm$, $26mm$, $35mm$ maximum displacements). **a-d** Cyclic breathing motion for contrast-enhancement level 3 (step-like curve for PPCR) shows separation of final registration position between two locations *see text*. **e-h** Breath-hold depth inconsistency for contrast-enhancement level 3 showing artefacts for direct fluid registration under increasing enhancement.

Figures 5.4a-d show the residual deformation after registration, demonstrating PPCR registration errors increasing with the amplitude of periodic motion. The profile of the unregistered deformation is a consequence of elliptical motion, taking different paths during inhale and exhale. Registration by PPCR for large periodic motion separates the final images into clusters at two locations. This is due to the periodic motion influencing early principal components, reducing the variance in early principal components and generating target images that are ill-defined (Figure 5.5b), containing a spread of image positions. Images in different phases of breathing are successively driven towards separate locations during subsequent iterations, since these are reinforced in the second calculation of the PCA, producing clusters of well-registered images.

If the final PPCR clusters are distinct, it may be possible to manually correct this effect by finding the transform between clusters, thus bringing the PPCR result to an equivalent correction to direct reg-

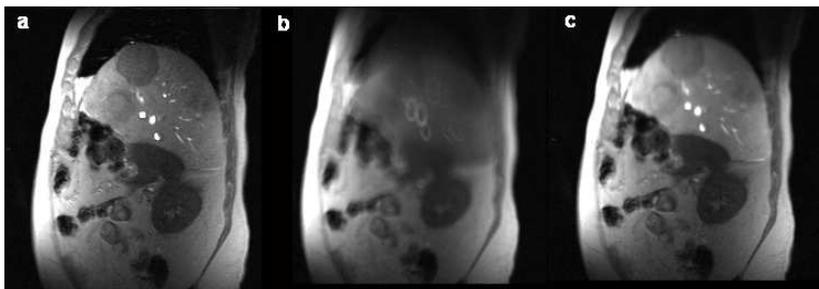


Figure 5.5: **a**) Source image from dataset with motion level 4 (up to 35mm displacement) and enhancement level 3. **b**) PPCR generated target image from first iteration for model with cyclic breathing motion for comparison with **c**) PPCR generated target image from first iteration for model with breath-hold depth inconsistency. The ill-defined nature of **b** results in separation of final registration position *see text*.

istration. This is left as future work but it is possible that since images within the separate locations are well registered to one another, we require only a single deformation between clusters. The transformation might be determined by the registration between images either side of the first jump between clusters and applied to all images within the cluster. This correction should be implemented in the first iteration of the PPCR method; inspecting the inter-image residual deformation allows clusters to be determined and then removed as above.

Figures 5.4e-h demonstrate the difference in registration success for varying breath-hold consistency level and fixed contrast enhancement. With no enhancement, both fluid and PPCR demonstrate successful registration, illustrated by a reduction in the residual motion. With increasing enhancement, fluid registration begins to mis-register enhancing regions, particularly visible in the images as distortions to the rigid tumours. This effect is shown as a failure to reduce the residual deformation level. For level 3 enhancement, the fluid registration is actively mis-registering a large proportion of the images. The eventual success of the PPCR method is implied in the first target image shown in Figure 5.5c in which features are given a well-defined average position.

Figure 5.6a plots fluctuations in tumour size. It is clear that the tumour volume-preserving modification in Chapter 3 is not entirely successful, compressing the tumour up to 1% with increasing motion levels (see the NoReg data in Figure 5.6a). The failure of the modification is less important when considering volume change due to the direct fluid registration. Volume change is visually correlated with the strength of contrast enhancement and volume changes of up to 7% are observed, visible in the actual images. Tumour volume change is never more than 1% using the PPCR method. Since tumours are the very objects we are likely to be interested in, the success of registration in this region is crucial. This is particularly true when monitoring the response to therapy with longitudinal scans.

The joint image histograms between pre and post enhancement images contain dispersion from both motion and contrast enhancement processes. PPCR separates these processes but is more effective between inconsistent breath-hold depth than in cyclic breathing motion. Histogram dispersion is a result of one-to-one pixel intensity relationships becoming one-to-many due to spatially dependent intensity

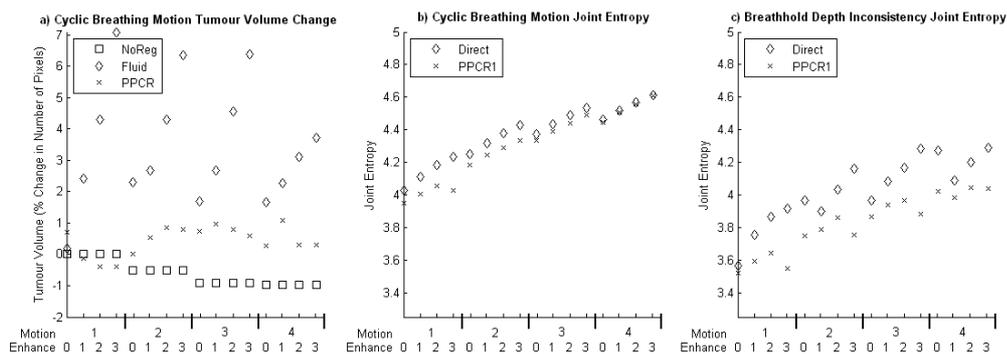


Figure 5.6: Graphs plotting statistics for variations in **a)** tumour volume change for breathing-motion extent and enhancement level, note the trend for decreasing tumour size in the gold standard. **b)** average joint entropy of target and source images per simulation for direct registration (to pre-contrast image) and to first set of target images generated by PPCR for breathing motion. **c)** for breath-hold depth inconsistency. Note increasing joint entropy with motion *and* enhancement.

variations. Crudely, this is seen as lobe-like arms in Figure 5.7 for a real DCE-MRI dataset. Information based cost-functions aim to minimise this dispersion but do not distinguish between the two sources of dispersion. For PPCR generated target images, the contrast enhancement level is more closely matched since gross intensity changes are encoded in the early principal components, therefore dispersion due to contrast enhancement is reduced and the remaining dispersion is more strongly associated with motion. This can be seen in Figure 5.6c as an enhancement dependent reduction in the joint entropy between source images and PPCR generated target images at each motion level. If the PPCR generated target images contain ambiguous boundaries or poor contrast-enhancement matching relative to the source images, then there may be an increase in joint image histogram dispersion. Poor target representation in the breathing motion case can be compared to the breath-hold case in Figures 5.6b and 5.6c; improved target representation in the breath-hold case allows a reduction in joint entropy with reduced dependence on enhancement level than in the breathing motion case. Additional dispersion in the joint histograms is always likely to occur under PPCR (this is best shown in Figure 5.7 Row 1). This is because PPCR generates target images in early iterations that contain averaged representations of enhancement, but if the reduction in dispersion due to contrast enhancement is greater than this additional blurring, registration by PPCR may proceed.

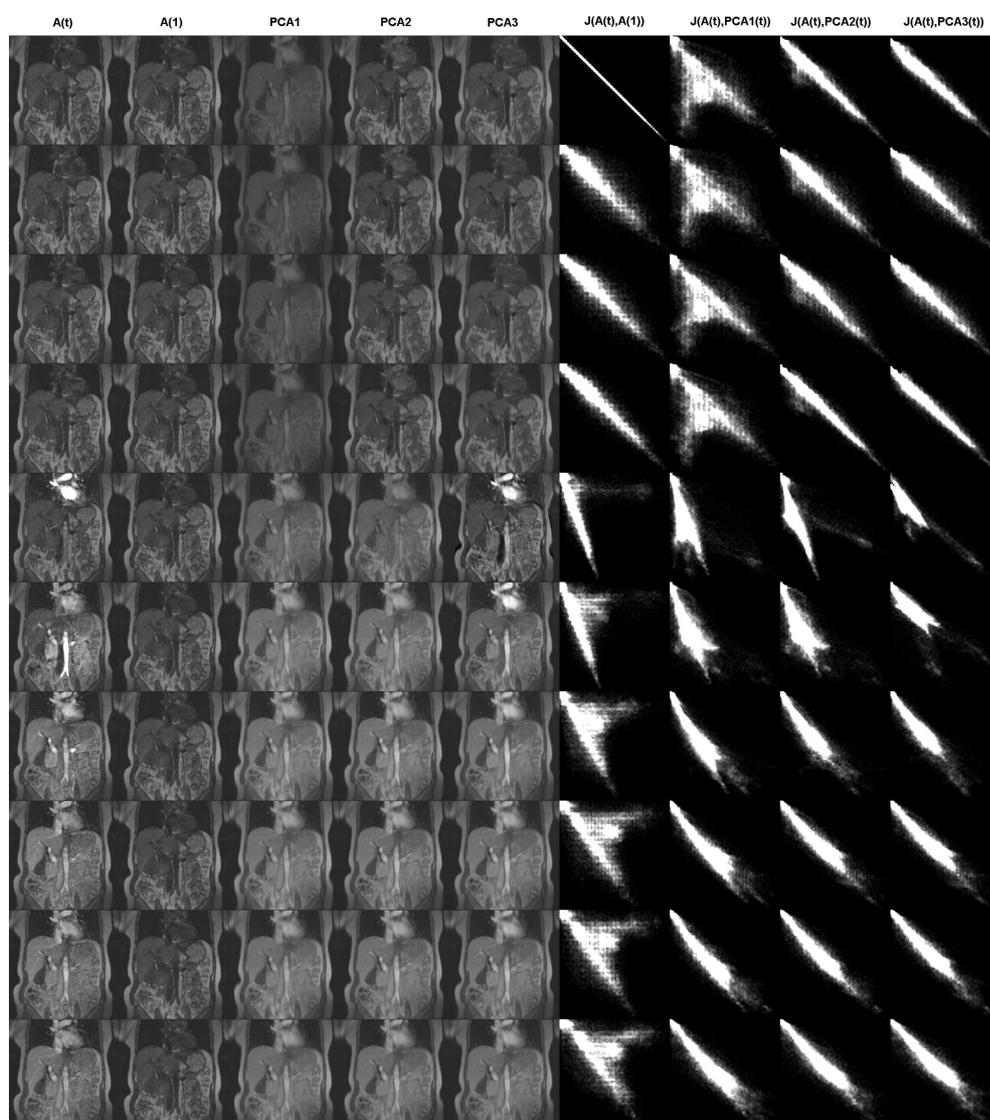


Figure 5.7: Effect of PCA on joint image histogram formation from a real dataset of 20 2D images. **Column 1** original (real) dataset (images 1-10), **Column 2** first image in dataset, **Columns 3-5** images 1-10 rebuilt using 1-3 principal components. **Column 6** joint image histograms (x-axis float, y-axis anchor image intensities) of images in column 1 with those in column 2, **Columns 7-9** joint image histograms of Column 1 with Columns 3-5. Dispersion in Column 6 is the result of both motion and contrast enhancement, using PCA allows some removal of enhancement (lobe-like) dispersion, although the effect is reversed for the pre-enhancement images in Column 7. Inclusion of further principal components in Columns 8-9 removes.

5.3.3 Conclusion

We have shown the applicability of two image registration methods under different levels of motion and enhancement. Under cyclic breathing motion, PPCR finds it difficult to generate a representative set of target images (Figure 5.4a-d), but this type of motion produces a predictable result, requiring a modification or diagnosis that could be included in the algorithm. Such a method could inspect early principal components or resulting anchor images in order to suggest or predict the performance of the PPCR algorithm. The clustering shown in Figure 5.5 is not observed in any real data, and subsequent anchor image formation, later in the thesis. The PPCR method is able to preserve the volume of enhancing regions unlike direct registration which begins to fail under increasing contrast enhancement. This is a failure of the cost-function to account for the appearance of new structure. Using information based cost-functions, there will be an increase in joint histogram dispersion which is not distinguished from mis-alignment dispersion, making the cost function inappropriate. By encapsulating intensity variations in early principal components, PPCR generates enhancement matched target images, reducing dispersion in the joint histogram due to contrast enhancement. This allows registration to proceed, but only in cases where target images are well-matched to their source images, which is not the case under cyclic breathing motion. In this case, reductions in contrast enhancement dispersion are offset by an increase in dispersion due to poor target matching and PPCR will break down. Future work will develop the breathing-model to allow for more realistic unpredictability in breathing-depth and phase, improve the enhancement model to make it organ specific, extend the work to full 3D and include medial-lateral deformations.

5.4 The Effect of PPCR on Cost-Function Minimisation Space

We inspect the effect of PPCR on the formation of cost-function space as presented in Section 4.3. By applying an elastic deformation to the centre of our image with varying magnitude and direction, we can compute the cost-function space associated with minimisation of a particular cost-function (see Section 4.3). The centre of the following cost-function spaces represents the value of the cost-function between these image pairs.

5.4.1 Results

Figure 5.8 demonstrates the effect of contrast enhancement on cost-function optimisation. The first column represents the anchor image(s), which in this case is just the first pre-enhancement image from the first 10 images of a DCE-MRI dataset. The second column contains all ten images and we inspect the result of calculating different cost-functions between the image-pairs.

The cost-function spaces of Figure 5.8 reveal a large amount of information about the image alignment process. Inspection of the pre-enhancement images reveals that we should expect to be able to align the float image if it is deformed as described above for all cost-functions. Discrepancies in the shape of the cost-function value are largely the result of the correction of minor mis-alignment due to breath-hold depth consistency in these real images: for instance a minor superior→inferior displacement of the liver is partially corrected by an inferior→superior force of the type described above. The appearance of contrast enhancement in the heart and aorta disturbs the cost-function space, particularly for the statistical cost-functions: cross-correlation, sum of squared differences and sum of absolute differences. In the case of the fifth image pair, the cost-function spaces show that it is preferable to distort the images using an inferior→superior force. This would have the effect of compressing the enhancing heart, removing the intensity discrepancy and therefore maximising the image similarity. Similarly for the sixth image pair with enhancement of the aorta, both cross-correlation and sum of squared difference cost-functions suggest that a positive medial→lateral force will improve the image similarity. This can be seen to be the case, since the aorta will begin to overlap the brighter regions adjacent to the dark pre-enhancement aorta. After the first passage of the bolus the cost function extremum, in all cases, becomes disperse suggesting that finding the correct image alignment by registration will become difficult. In particular, when enhancement is present Figure 5.8 suggests that cross-correlation and sum of squared differences should not be used. Minimisation of the joint entropy cost-function also suggests that there might be some problems finding a smooth gradient descent through the cost-function space. The effect of the first anchor images generated by the PPCR algorithm on the cost function minimisation space is shown in Figure 5.9. The first column represents the anchor images, generated from a PCA of the ten float images using the PPCR method. The second column contains all ten images and we inspect the result of different cost-functions between the image-pairs. This value corresponds to the centre of the following cost-function spaces in which the float image is not deformed. We now apply a deforming force to the centre of the image and solve the linear elastic equation to find the resulting deformation over the image. By varying this force (in two-dimensions: medial-lateral and superior-inferior) and looking at the resulting cost-function value between the anchor image and the deformed float image, we can investigate

how easily we would expect to be able to recover the deformation by registration using that particular cost-function.

The cost-function spaces of Figure 5.9 reveal the effect of the first PPCR iteration on the cost-function minimisation space. Inspection of the pre-enhancement images reveals that we should expect to be able to align the float image if it is deformed as described above for all cost-functions. However, the minimisation does not appear to be as well-defined as in the registration to the first image case. This is due to information about other structures being present in the anchor images for each image pair. As before, discrepancies in the shape of the cost-function value, asymmetry around the centre of the cost-function minimisation space, are largely the result of the correction of minor mis-alignment due to breath-hold depth consistency in these real images. The appearance of contrast enhancement in the heart and aorta still disturbs the cost-function space despite using PPCR, particularly for the sum of squared differences. In the case of the sixth image pair with enhancement of the aorta, the problems seen in the corresponding image pair for Figure 5.8 are no longer present, we may expect the cost-function to be suitable for minimisation. The advantage of the PPCR method is most obvious in the use of statistical cost-functions post-enhancement, all cost functions now appear reasonable and we would expect to be able to align each image pair. Figure 5.9 suggests that we should be able to use any cost-function in image registration using PPCR. A further advantage of the PPCR method is that it is iterative, we should be able to find a reasonable alignment in the first iteration as shown here, but further iterations should refine this alignment further.

5.4.2 Conclusion

The results of this section demonstrate the potential benefit on cost-function minimisation of PPCR when registering contrast enhanced images. In the cases shown, for very simple displacements, the PPCR algorithm may be used to allow registration of images using cost-functions that do not cope with contrast-enhancement. For the type of deformation applied here, it appears that we might achieve a good registration using information theoretic cost-functions, however for more complex displacements this cannot be guaranteed. The PPCR algorithm is also iterative, so subsequent deformations are refined and we do not stop by simply registering to the first set of target images generated by PCA, this is important from the perspective that we wish to register a large number of images into a common coordinate system.

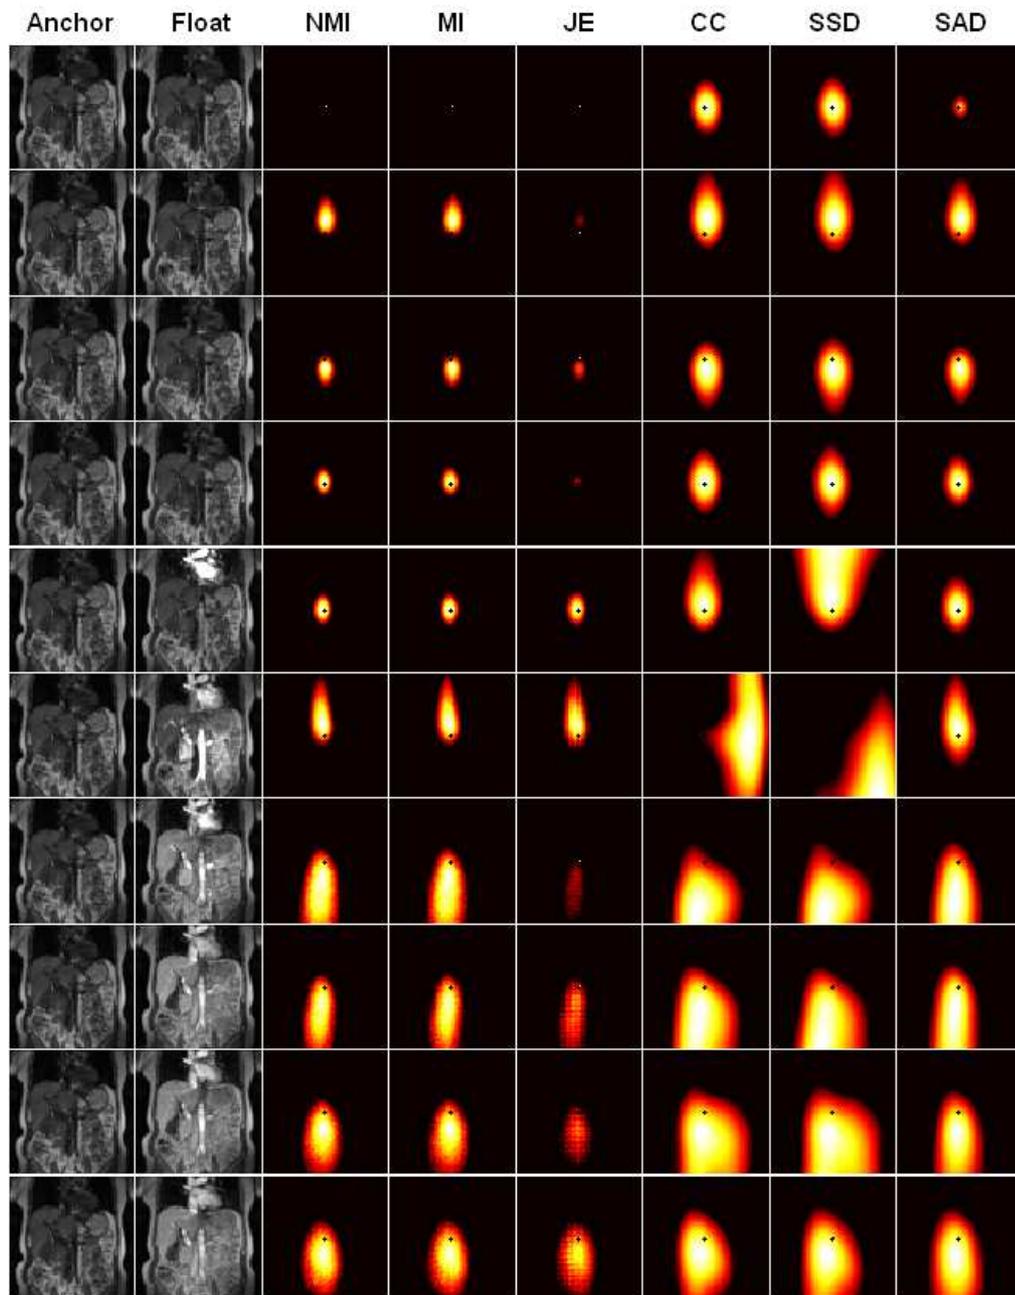


Figure 5.8: Direct registration cost-function minimisation spaces for the first ten images from a DCE-MRI dataset. For ease of presentation, the negative-log values for MI, NMI and CC are shown; here light corresponds to good image alignment and dark to poor. The x-axis corresponds to a large medial-lateral force applied negative through positive from left to right, the y-axis is the equivalent for the superior-inferior force. The cost function comparison is taken between the corresponding image on the far left, and the float in the neighbouring column subject to the given deformation (*see text for clarification*).

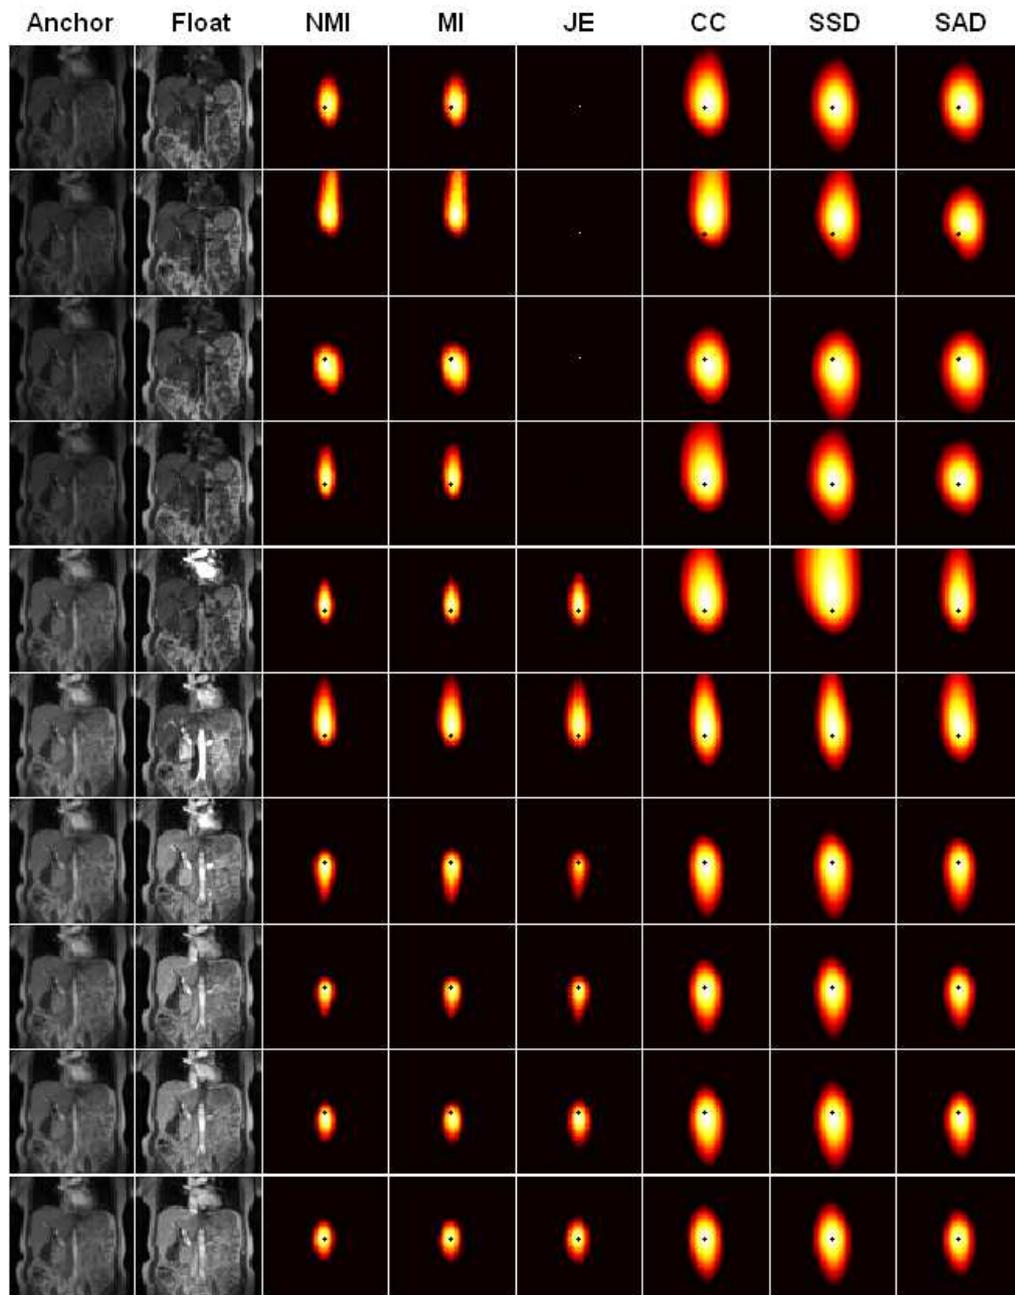


Figure 5.9: PPCR cost-function minimisation spaces for the first ten images from a DCE-MRI dataset. For ease of presentation, the negative-log values for MI, NMI and CC are shown; here light corresponds to good image alignment and dark to poor. The x-axis corresponds to a large medial-lateral force applied negative through positive from left to right, the y-axis is the equivalent for the superior-inferior force. The cost function comparison is taken between the corresponding image on the far left and its neighbour, subject to the given deformation (*see text for clarification*).

5.5 Choice of Anchor Image in Conventional Registration - Revisited

This section briefly revisits Section 4.5 to include the Cost-Function Matrix Mean (CFMM) PPCR result concerning the choice of anchor image in the conventional registration of DCE-MRI. In comparison, the choice of anchor image in direct image registration provides the co-ordinate system for all images. When using PPCR for image registration it may be necessary to convert between co-ordinate systems when analysing the result of PPCR with reference to either the original images or those produced by a different registration method. This is because registration by PPCR is to a coordinate frame formed in the process of the registration algorithm, which is likely to be different from the coordinate frame of any of the individual images.

The results of Figure 5.10 now include NMI-CFMM values after registration by PPCR. It is seen that PPCR allows a higher value of NMI-CFMM than any choice of anchor image in conventional registration. This provides some evidence that registration by PPCR to an iterated coordinate system allows improved registration of groups of DCE-MR images when compared to registration to the coordinate system of a single image.

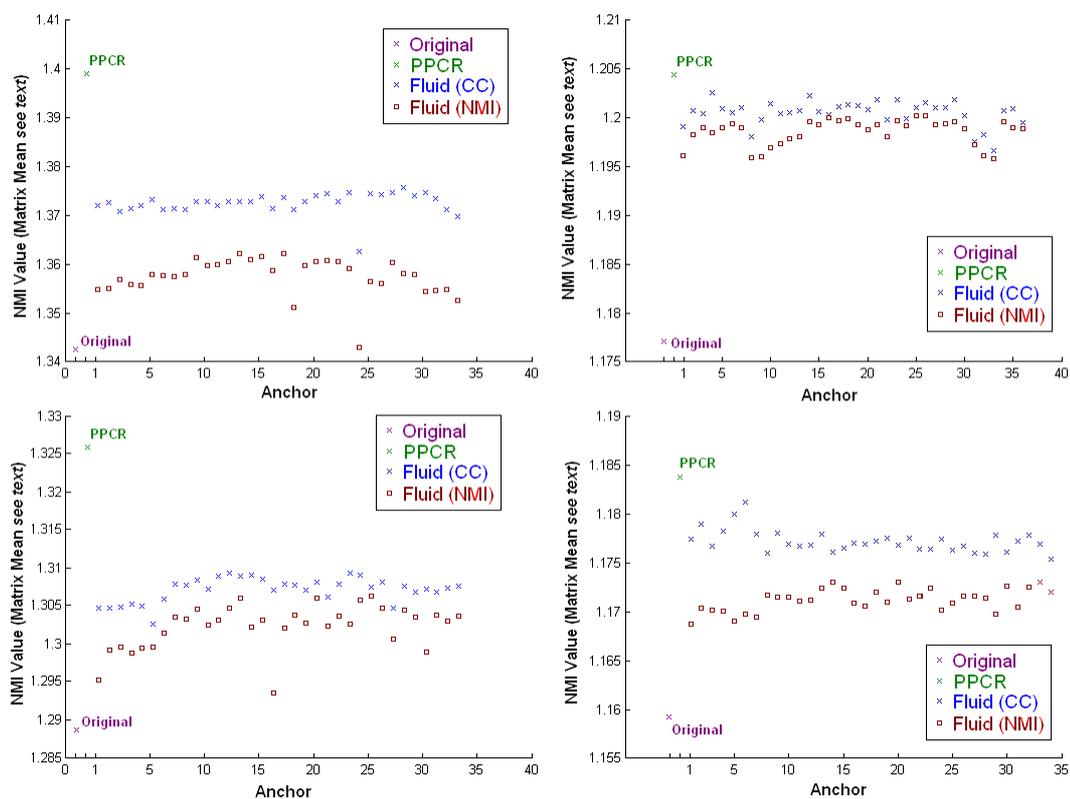


Figure 5.10: Assessment of registration result of variations in target image selection using Cost Function Matrix Mean (see Section 4.4). For four separate datasets from Table 1.1, fluid registration proceeds using the n th image as the anchor image. The NMI Matrix Mean is shown for the result of registration using Cross-Correlation (blue) and Normalised Mutual Information (red). Also shown are the original NMI Matrix Mean before registration (magenta) and after registration using PPCR (green).

5.6 Analysis of the PPCR Algorithm Progress

5.6.1 Changing the number of Principal Components

The iterative nature of PPCR makes the process slow. For a dataset with n images, instead of running $(n - 1)$ registrations to a single anchor image chosen from the dataset as with direct registration, the number of registrations becomes $n(n - 1)$. However, since the later principal components do not contain much variance and therefore may not contribute noticeably to the final anchor image intensities, it may be beneficial to terminate the PPCR process once the total variance of the original dataset contained in the new anchor images reaches a threshold value. This section investigates the use of such a stopping criteria and the potential CPU-time benefit.

In addition to the stopping criteria discussed above, there is a more general way of monitoring the progress of each PPCR iteration. Analogous to the continual updating of the fluid registration deformation, we can track the deformation field as it changes. This is not necessarily a good way of finding a stopping criteria since the size of the displacements in the current deformation field do not necessarily predict the size of displacements in the following deformation field. The reason is that the inclusion of a later principal component may produce larger changes in the target images than the previous component whilst still containing a smaller amount of dataset variance - this is likely to be the case where one image in the dataset has quite different intensity variations from the other images (as seen in bolus arrival images).

Results

Figure 5.11 shows the progress of the deformation towards the final deformation field with each PPCR iteration for four datasets from Table 1.1. The total absolute residual over all images is calculated and divided by the total number of pixels. The result for each of the four datasets demonstrates a steady decrease in total absolute residual towards the final position. The curves are slightly convex, suggesting that later iterations contribute slightly less towards the final deformation field. However, it is not obvious that the PPCR process should be terminated before the final iteration.

Conclusion

The results of Figure 5.11 do not suggest that there is a benefit to be gained from terminating the PPCR algorithm at an early iteration. Early termination may prevent the registration reaching the true final deformation field that is only achieved once all principal components are included. The final deformation field should be considered preferable because of the inclusion of later principal components which allow the continued refinement of the registration deformation field. The following section (Section 5.6.2) discusses an alternative stopping criteria more suited to the nature of the PPCR algorithm and deformation field progression.

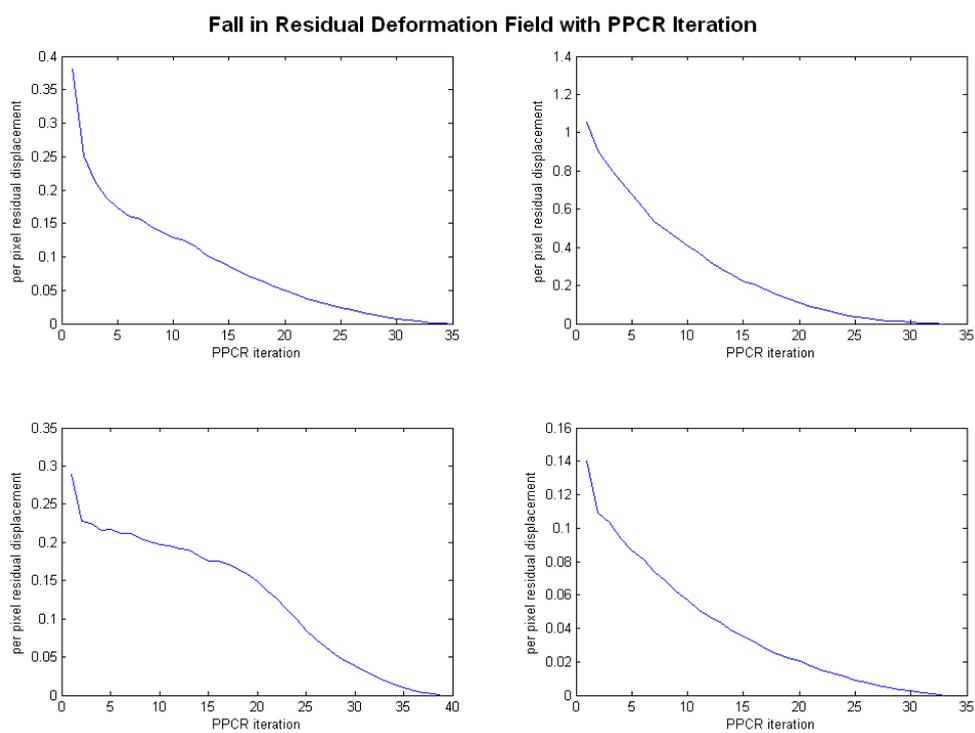


Figure 5.11: Example curves showing the approach of the PPCR deformation field towards the final deformation. Graphs show the per-pixel residual and the approach of the residual towards the final deformation position for four separate 2D DCE-MRI datasets from Table 1.1.

5.6.2 Changing the number of Registration Iterations

As discussed in the previous section, for a dataset of n images there are a maximum of $n(n-1)$ registrations when using the PPCR algorithm. However, the progressive nature of PPCR suggests that since the early registrations are crude, we need not run them for many internal iterations. For instance if a single registration takes 400 fluid-equation iterations, then we can limit linearly the number of internal iterations in the registration so that they only do, for instance, 20 internal iterations before the next PPCR iteration. This section investigates the effect on accuracy of altering the internal registration iteration number.

We vary the maximum number of intrinsic registrations between 0-400. The final deformations are analysed and the sum of squared differences found between the standard 400 iteration maximum displacement field and the reduced-iteration displacement field.

Results

Table 5.1 shows the total residual deformation when using a low number of internal registration iterations. The default maximum number of iterations used in the internal fluid registration algorithm is 400, hence the residual difference of the final deformation field from this result is shown. Equivalent graphs for the two DCE-MRI datasets (see Section 1.3) are shown for the approach of the deformation to the maximum iteration case with increasing iteration number.

Table 5.1: PPCR Registration for varying maximum number of internal registration iterations

Internal Registrations	Time	SSD between 400i result (Livdt-04b)	SSD between 400i result (Livdt-07b)
0	0	2.374	2.049
10	10n(n-1)	0.223	0.121
20	20n(n-1)	0.118	0.091
40	40n(n-1)	0.054	0.059
60	60n(n-1)	0.034	0.025
80	80n(n-1)	0.024	0.024
100	100n(n-1)	0.024	0.028
120	120n(n-1)	0.026	0.025
140	140n(n-1)	0.021	0
160	160n(n-1)	0.018	0
180	180n(n-1)	0.027	0
200	200n(n-1)	0.023	0
300	300n(n-1)	0	0
400	400n(n-1)	0	0

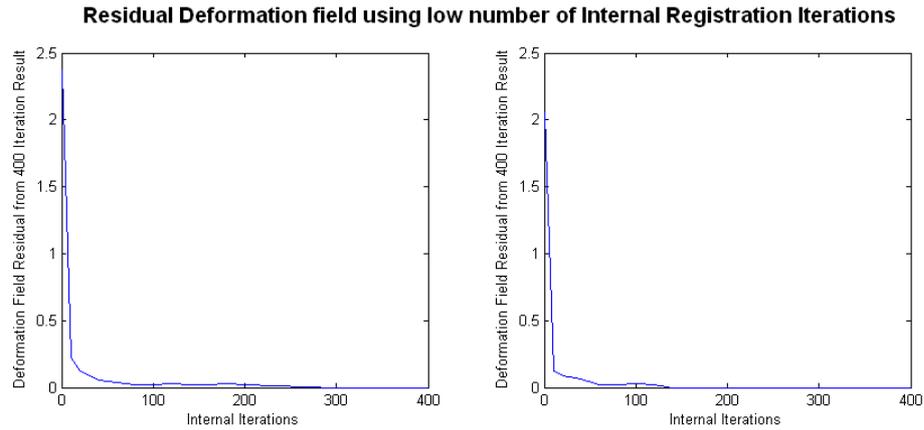


Figure 5.12: Example curves showing the approach of the PPCR deformation field towards the maximum deformation found when varying the number of internal registration iterations up to a maximum of 400 iterations in the component fluid registration algorithm. Graphs show the absolute residual (divided by total number of pixels) and the approach of the residual towards the maximum deformation position for two DCE-MRI datasets from Table 1.1.

Conclusion

The results shown in Figure 5.12 present a strong result (at least for the two cases tested) for shortening the total CPU runtime by reducing the maximum number of internal registration iterations. This should be contrasted with the result of the previous section (Section 5.6.1). The results suggest that it may be feasible to reduce the number of registration iterations by a substantial amount, whilst retaining a good approximation of the final result. Although the number of iterations is fixed throughout the algorithm, it may be preferable to adjust the number of iterations dynamically or to increase the number of iterations as image detail is added with increasing numbers of principal components.

5.7 Future Adjustment for large variations in Contrast Enhancement

As seen in Figure 5.2, extreme enhancement features are not well-represented in the first principal components. Therefore in this case the PPCR method is unable to fully match the contrast variation and the cost-function gradients may include forces that may result in mis-registration of enhancing features. This will only occur in the situation that a few images that contain very different features or contrast from the other images of the dataset. This section discusses one possible fix for this disturbance by adding the effect of extra components to particular images within the dataset.

The target images are checked for suitability using the following algorithm: 1. find the Sum of Squared Differences cost-function between each float and anchor pair; 2. Find the mean and standard deviation of this spread of cost-function values; 3. If the cost-function of individual float-anchor pairs is more than one standard deviation from the mean, then add an additional principal component in the formation of *only* this image; 4. Iterate this process until all float-anchor Sum of Squared Differences are within one original standard deviation. Although the effect may be to cause under-registration of the affected images, the modification should help prevent the mis-registration that may otherwise occur.

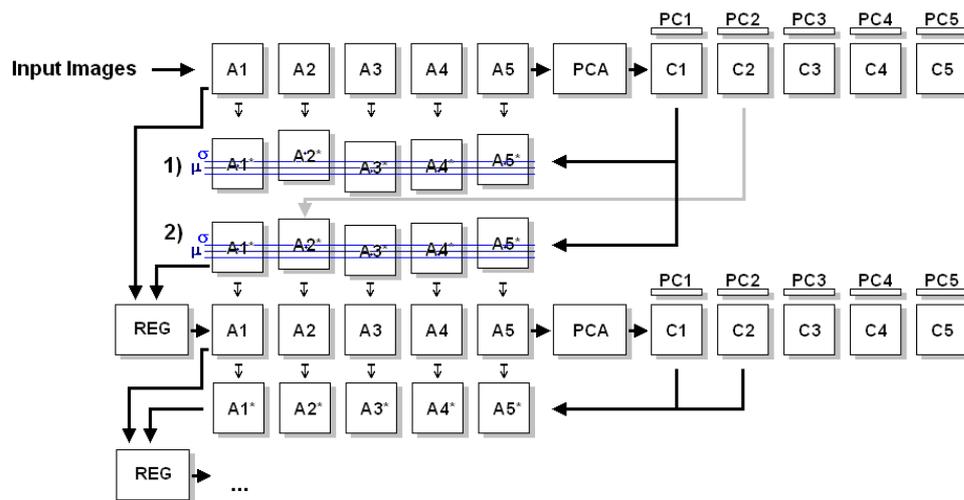


Figure 5.13: Illustration of the PPCR adjustment process. After the formulation of principal components and weighting maps in the first iteration, we form target images from the first principal component. **Step 1** we find the mean (and standard deviation) sum-of-squared differences (SSD) between all float-anchor pairs. **Step 2** if any of the float-anchor pairs have an SSD value outside of one standard deviation from the mean SSD, we add further principal components until the SSD of that float-anchor pair is below one standard deviation from the mean float-anchor SSD.

5.8 Conclusion

This chapter has discussed some of the expected benefits of the PPCR algorithm and modifications that might be made to improve performance. The use of the model in Chapter 3 has been used to show that certain types of motion are not well suited to the PPCR algorithm. One of the assumptions of the PPCR algorithm is that early principal components are dominated by enhancement characteristics and later ones by motion. In the case of smoothly periodic motion, early principal components contain representations of the pixel intensity fluctuations associated with objects moving into and out of that pixel. However, the requirements for this type of corruption to occur under PPCR require motion to be periodic. The results of the DCE-MRI simulations suggest that PPCR is effective for repeated breath-hold data. When comparing the PPCR algorithm to direct registration of images to a single anchor image, revisiting the choice of anchor image data from Chapter 4 suggests that even if it were possible to choose the best target image *a priori*, the resulting registration performance is exceeded by using the PPCR algorithm (Figure 5.10).

With regard to the long PPCR run-time, some performance benefits can be produced. Although terminating the number of principal component iterations early does not seem to confer a benefit, setting the number of iterations in each fluid registration can be used to achieve a time-saving with less of an effect on the final outcome. This may make the algorithm a desirable addition to a conventional registration algorithm when registering groups of images with little time penalty.

Section 5.4 shows the effect of PPCR on the cost-function minimisation space. For the simple deformations shown, PPCR produces a cost-function space that is better-defined for minimisation by all cost-functions (Figure 5.9). The use of a fluid or b-spline method should not influence the performance of the PPCR algorithm; as discussed in Chapter 4 it is the formation of the cost-function space that is important when registering contrast-enhanced images, the transformation model is used to regularise the cost-function minimisation to generate a desirable (e.g. diffeomorphic or smoothed) deformation. A low degree-of-freedom model such as an affine parametrisation may also be implemented but restrictions on the allowed deformations prevent the effect of the benefit of PPCR in local regions of contrast variation.

Chapter 6

Registration of Breath-hold Dynamic Contrast Enhanced MRI

The following chapter will analyse the results of applying the PPCR algorithm to real DCE-MRI data. The algorithm is first applied to a study of twenty-seven 2D datasets from seven patients each with at least one follow up scan (Table 1.1). Section 6.1 will compare the results of registration by direct registration to the first image in the dataset and registration by PPCR. Section 6.2 will analyse six full 3D DCE-MRI datasets using pharmacokinetic analysis to determine the success of different image registration methods.

6.1 Registration of 2D DCE-MRI Using PPCR

6.1.1 Introduction

The PPCR algorithm is now further discussed by comparing it to those cases where conventional registration causes artefactual misalignment of contrast-enhanced images. By inspection of images by blinded observers and through basic pharmacokinetic model fitting, we can evaluate some of the benefits of the PPCR method. The PPCR method outlined in Chapter 5 is compared to a simple, single registration to the first image in each dataset, also registered with the same Eulerian fluid registration process.

The liver is analysed by implementing non-rigid registration methods based on a fluid equation. The use of a fluid-equation based registration over other transformation models is of little importance at this stage. As discussed in previous chapters, failure of image registration in the case of DCE-MRI is due to the cost-function not the transformation model. The same fluid registration algorithm is used for direct fluid registration and within the PPCR algorithm. The fluid-equation is balanced using image derived forces calculated from an image similarity measure, in this case cross correlation, a measure normally considered suitable for same-modality images. The use of cross-correlation is due to empirical observation rather than theoretical considerations since it appears to perform a better registration in those cases in which the registration is correct. This was discussed briefly in Chapter 5. An implementation of this approach as developed by Crum (Crum et al., 2005) is used based on original work from Christensen (Christensen et al., 1997). The images are analysed using two registration schemes, the direct fluid registration and the PPCR scheme (Melbourne et al., 2007b). Registration in the case of 2D data might be affected by through plane motion: in the case of objects moving from the field of view, this information is likely to be encapsulated in the principal components leading to the algorithm generating anchor images maintaining an absence or presence of these features. Application of the subsequent 2D image registration will be less likely to result in mis-registration.

Other Registration Methods

Some other possible registration methods may confer an advantage in the case of DCE-MRI data and they are discussed here.

Registration to Mean Image If the mean image is used as the anchor, we may proceed by registering all images to this mean-image as used for comparison in Buonaccorsi *et al* (Buonaccorsi et al., 2006). It is possible to imagine an iterative scheme by which the mean-image is updated after a few registration steps. However, the reduction in image resolution by the summation of the original image intensities may result in either under-registration, due to features being aligned to their position in the mean-image, or mis-registration of features. Mis-registration could result from the formation of the joint-image histogram using the mean-image. Dispersion in the joint image histogram will be increased, the loss of definition in the pixel intensity mappings will smooth the bin contents so we might expect weaker force-gradients in the resulting registration. As briefly discussed in Section 5.2, registration to the mean image may be considered a nullth PPCR prior to the inclusion of any principal components and therefore generating anchor images without contrast matching.

Grouped Fluid Registration A further method can be devised that attempts to mediate spurious registrations by assuming that an average of the equivalent registration paths between *float to anchor* and *float to neighbouring time-point to anchor* will provide a more robust registration. For images in a dataset \mathbf{A} at timepoint t , we might expect the registration $A(t) \mapsto A(0) \equiv A(t) \mapsto A(t-1) \mapsto A(0) \equiv A(t) \mapsto A(t+1) \mapsto A(0)$. The reasoning is that mis-registration artefacts are unique to each float image, so taking an average, via the registration of its nearest temporal neighbours will reduce spurious misalignment since these will not be present in the remaining registration paths. Using this method, contrast-enhancement induced misalignment will be reduced by suppression of registration artefacts associated with the individual registration paths, this is in contrast to the PPCR algorithm, where ideally, misalignment due to contrast enhancement should not occur since we are contrast-enhancement matching.

Registration to Mean Position An alternative algorithm to registration to the mean image is registration to the mean position. As discussed, registration to the mean image may result in a redefining of image intensities and mis-registration to averaged boundaries and features. By careful formulation of the force gradients (see Chapter 4) it is possible to add information from multiple images. For information based cost-functions, histograms may be formed from all images in the dataset; histogram bin incrementation would preserve information from image intensity values since no intensity averaging is required. The result can be imagined as an averaging of the image force gradients as opposed to the averaging of the images followed by the calculation of force gradients as in registration to the mean. The process may be iterated, updating the force gradients from each image periodically. This concept is not explored further in this thesis.

6.1.2 Method

An analysis of the success of registration is developed by visual inspection of the similarity of image features using image intensity difference images and additional blind evaluation. Pixel time-intensity curves are inspected for residual motion artefacts and intensity consistency. The Cost Function Matrix Mean evaluation criteria developed in Section 4.4 is also used to analyse the results of image registration.

An indirect validation of the registration method can be determined using a model-fitting algorithm. Significant work has been done using the extended Kety Model used by Buonaccorsi (Buonaccorsi et al., 2006). The standardisation and interpretation of the parameters is covered by Tofts (Tofts et al., 1999) and problems with uncertainty in the model fitting are discussed by Buckley (Buckley, 2002). In order to assess the registration, without needing to determine an arterial input function, the slow variation in the wash-out phase can be fitted to a function such as (6.1), which is a de-parametrised interpretation of the post-enhancement Kety model. This model does not attempt to fit to the bolus arrival, the shape of which is useful in determining K^{trans} , and therefore does not assess the impact of reduction in model-fit error on determination of this parameter.

$$\mathbf{A}_{ij}(t) = B_{ij}e^{-b_{ij}t} \quad (6.1)$$

For each pixel \mathbf{A}_{ij} , the parameters B_{ij} and b_{ij} can be estimated using a non-linear, least-squares fitting routine. The result will only be used to gauge registration success. We expect pixels that are well-registered, post-enhancement, to exhibit a monotonic variation in intensity with no rapid intensity fluctuations. This is due to redistribution of contrast-agent around the body after a finite bolus injection. In the case of poor registration, artificial artefacts cause pixel-wise intensity fluctuations that do not fit this model and the sum-of-squared-differences (SSD) between original pixel and fitted curve provides an indication of registration error (6.2) (Hayton et al., 1997). If the registration is successful, the curve-fitting will be improved, the intensity profile will be smoother and the SSD will be reduced (Hayton et al., 1997), (Buonaccorsi et al., 2005).

$$SSD_{ij} = (\mathbf{A}_{ij}(\mathbf{t}) - B_{ij}e^{-b_{ij}\mathbf{t}}) \cdot (\mathbf{A}_{ij}(\mathbf{t}) - B_{ij}e^{-b_{ij}\mathbf{t}}) \quad (6.2)$$

6.1.3 Results

Figure 6.1 presents registration results from two situations. In the top row, registration by simple fluid and PPCR both correct for motion. In the bottom row, an example of the failure of simple fluid registration is presented. Evidence for mis-registration in the case of a simple fluid registration is found in Figure 6.1g, a subtraction of two images registered to the first image in this dataset. In the liver, Figure 6.1g shows evidence of artefactual fluctuations in tumour position in the upper lateral portion of the liver. The figure presents only the difference of two time points, but the effect repeats throughout the dataset and this would pose a serious problem for successful analysis.

Registration by Progressive Principal Component Registration (PPCR) in Figure 6.1h shows removal of the liver registration artefacts in Figure 6.1g (represented by less difference signal in the highlighted region). Comparison with the no-registration cases in Figure 6.1b shows that motion artefacts throughout the image are also reduced or removed. This is evidence that PPCR allows successful registration of DCE-MRI datasets. A comparison in the presence of motion is represented in Figures 6.1a to 6.1d. The registration to the first image in the sequence shows removal of much of the superior-inferior displacement artefact in the liver.

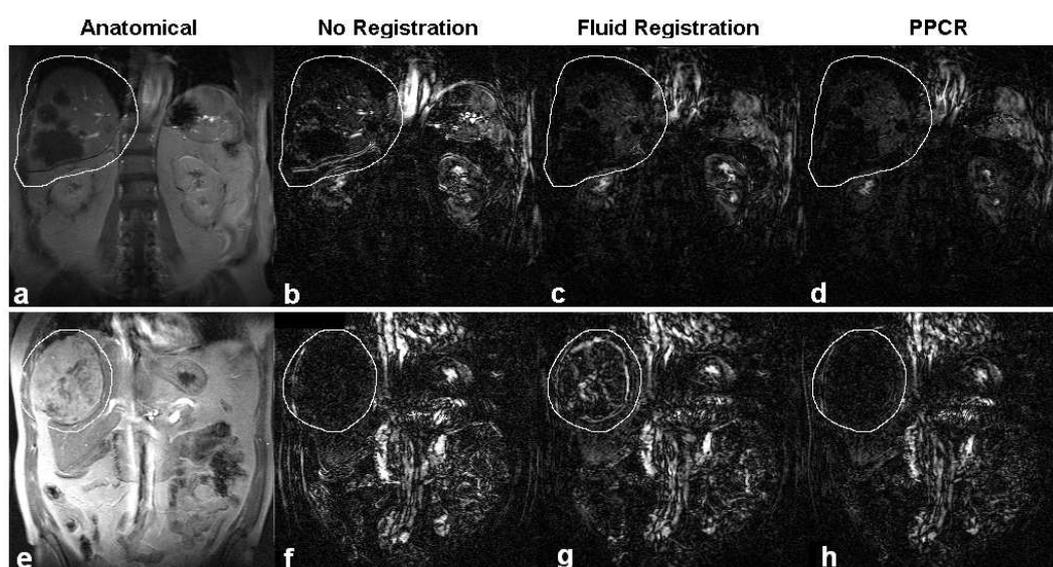


Figure 6.1: Absolute-difference images demonstrating registration failure for fluid registration to first image in dataset and correct PPCR, of two post-contrast images. Images **a** and **e**, anatomical images for reference. The drawn region is the same across a row and provides a visual guide. Images **b** and **f**, no registration for comparison. Images **c** and **g**, registration to first image in dataset, **c**, correct gross-registration of the liver and **g** artefactual tumour motion. Images **d** and **h**, registration by PPCR with improved motion correction and reduced artefacts.

The effect of registration on motion and artefact production is demonstrated in Figure 6.2 for four different patient datasets. In the first example, registration of the large tumour in the superior-lateral portion of the liver proved difficult for the basic fluid registration to first image in the dataset (Figures

6.2a and 6.2b), resulting in increased intensity fluctuations. PPCR registers the tumour correctly as shown. For Figures 6.2c and 6.2d the PPCR correctly registers the motion to the same extent as the simple registration, removing the real motion artefacts and therefore resembles the case of registration to the first image in the dataset. Figure 6.2c shows both registration methods working correctly for a moving bright region (associated with vasculature). Figure 6.2d demonstrates the correct registration of the diaphragm position by both methods at timepoint 26. Figures 6.2e and 6.2f show correct registration of motion by both methods with the exception of one intensity point that is mis-registered by the simple fluid registration scheme (timepoint 25). Figures 6.2g and 6.2h are sagittal images from a further patient. Figure 6.2g shows correct registration by both methods. Differences between methods in Figure 6.2h are due to mis-registration by the simple fluid registration scheme near an enhancing tumour boundary. In the majority of pixels, the smoothness of the intensity profile appears smoother, reflecting reduced motion-induced discontinuities.

Results from 22 datasets were compared using a blind-evaluation process (Table 6.1). Each of four operators were presented two movies side-by-side, from a selection of three movies of a particular dataset (the unregistered images, registration-to-first-image-in-dataset and registration by PPCR). The operator chose which movie they preferred or recorded no preference. A preference is characterised by reduced motion and reduced evidence of artefacts, particularly in the liver region. A group of four operators familiar with the process of image registration, although unfamiliar with DCE-MRI of the liver, evaluated 89 movie-pairs. The results of Table 6.1 show a preference for the PPCR method over both the unregistered datasets and the registration to the first image in the dataset. It should be noted that these are the preferences of image registration specialists and may differ from those of radiologists or clinicians. Three example datasets are included on the supplementary CD (see Appendix E). All movie files have the same format from left to right. The left-most movie is the original, unregistered DCE-MRI dataset. The second is registration using direct fluid registration using cross-correlation to the first (unenhanced) image in the dataset. The third is the result of registration using the PPCR algorithm. The file *movie-3-01.avi* demonstrates the correction of liver position between breath-holding using direct fluid registration and PPCR; *movie-3-02.avi* demonstrates the correction of the position and shape of a mass in the inferior liver by PPCR (third movie from left) and the second movie from the left in *movie-3-03.avi* shows failure of direct fluid registration in a contrast-enhancing mass in the superior liver.

The sum of squared differences (SSD) between time-dependent data pixels and model time-series of equation 6.1 are expressed as a percentage of the SSD value relative to the unregistered case (Figure 6.3). The PPCR almost always outperforms the simple registration to first image in the dataset. The apparent success of the registration methods appears to be patient dependent, patient 2 appears well-registered by both methods. For patient 6, PPCR provides a significant advantage. This is due to problems with the fluid registration, which can visibly distort tumour boundaries. The reason for the anomalous results of patient 1d and patient 4 is not immediately clear. Patient 1d is visually well-registered by both methods. Patient 4 appears registered to the same level of success by PPCR as with registration to the first image in the dataset. Visual inspection of movies of all the registration results reveals that PPCR has failed to

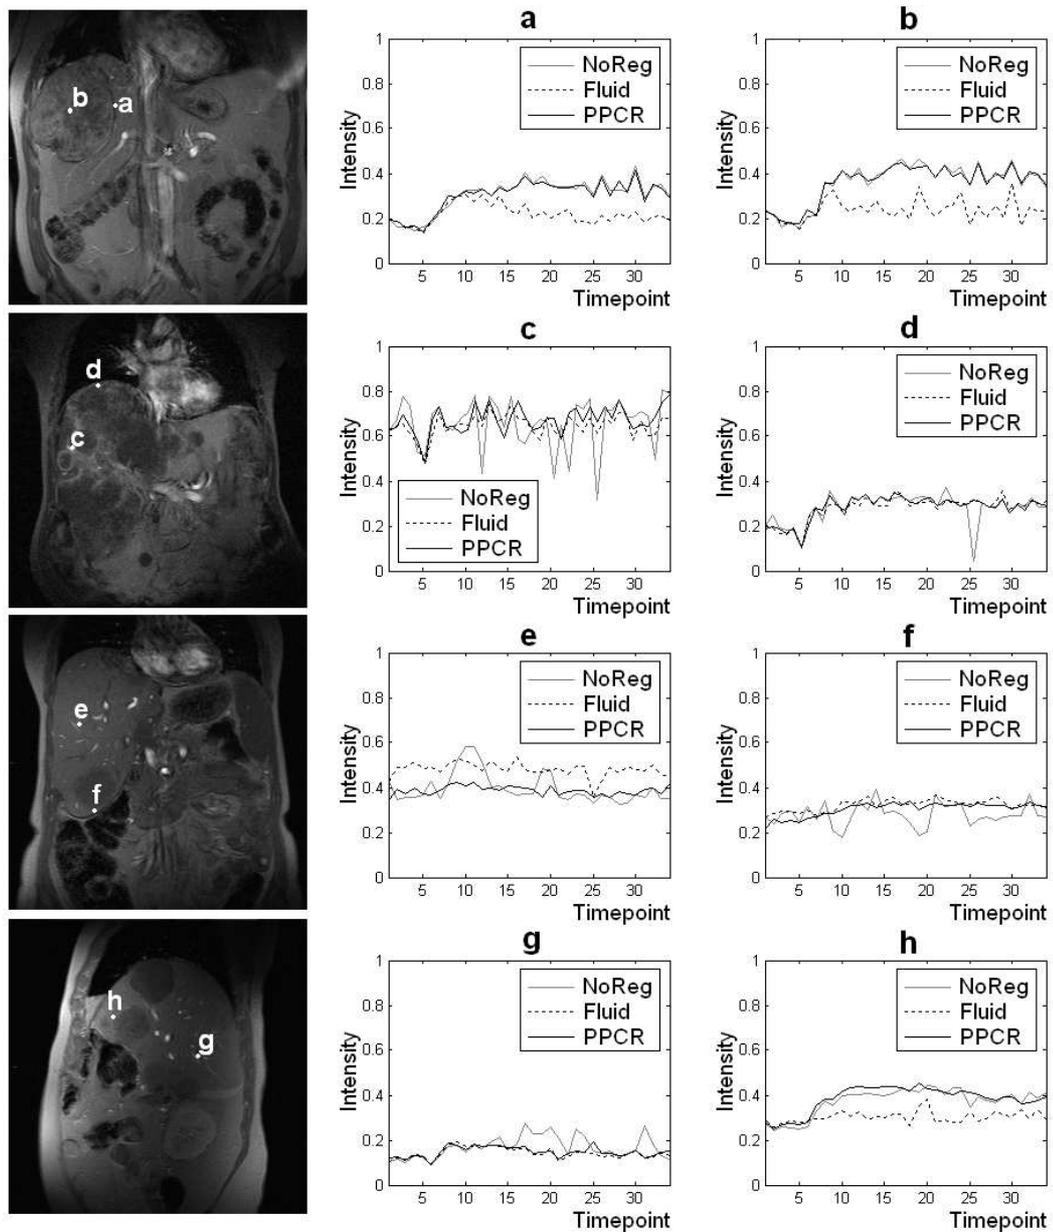


Figure 6.2: Plots of pixel intensity with time for pixels selected from anatomical images (left-hand column). The unregistered cases (NoReg) are shown for comparison. Artefactual oscillations in fluid timecourse for a) tumour artefacts due to nearby boundary motion & b) tumour artefacts due to tumour mis-registration (see Figure 6.1g). c) corrected motion of bright region within liver & d) corrected diaphragm position with large inferior displacement at timepoint 26. In this case both registration methods (Fluid & PPCR) identified the large displacement. Again for different patients, e) a correctly registered bright region and f) correct registration by both methods of inferior liver motion. Sagittal images, g) correct registration by both methods and h) correct registration by PPCR and mis-registration close to a tumour boundary by the simple fluid registration scheme.

Table 6.1: Blind-evaluation of different registration methods. Each row contains scores comparing two registration types, representing the number of cases each registration was preferred over the other, and the number of cases in which there was no preference. Permutations are between either the unregistered case (NoReg), registration to first image in the dataset (Fluid) and to PPCR (PPCR).

Registration	Instances Preferred	Registration	Instances Preferred	No Preference
NoReg	0	Fluid	28	5
Fluid	0	PPCR	25	5
PPCR	25	NoReg	0	1

correctly register a small number of images in this dataset. This was not detected by the blind-evaluation process because the randomly picked pairs did not include this example. This may be because the principal component analysis has not completely separated motion artefacts from contrast-enhancement intensity changes and so the generated target images are not entirely motion-free. Again for patient *3b* the SSD values are very similar. Visual inspection reveals the registrations are also very similar, although a comparison of PPCR with fluid registration did not occur in the blind-evaluation. In the majority of cases, PPCR effectively de-couples motion induced intensity changes from contrast-enhancement induced changes, allowing a registration unencumbered by contrast-enhancement intensity variations. Registration by PPCR allows improved curve-fitting, which in principle allows superior model-fitting and physiological parameter extraction.

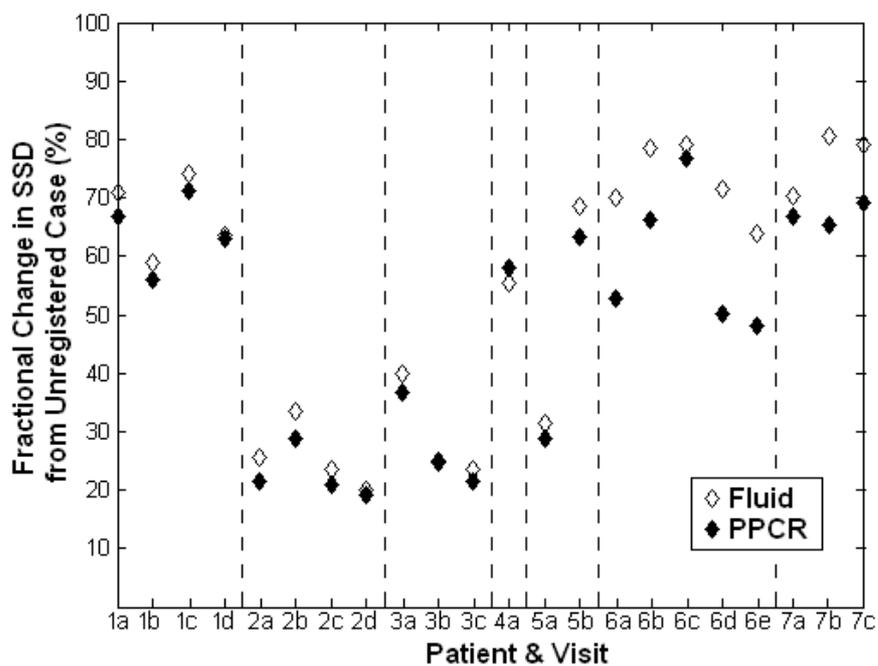


Figure 6.3: Results of fluid registration to the first image in the data set or PPCR. The comparison uses the sum of squared differences (SSD) between the image pixel data and the decaying exponential model of the post-enhancement phase. Results are shown as percentages of the respective SSD in the case of no registration.

Analysis of Registration Success using Cost-Function Matrix Mean

The cost-function matrix mean (CFMM) analysis measure discussed in Section 4.4 may be used to evaluate registration success. Here we will continue to use normalised mutual information (NMI) as the evaluation method by which we will compare the different registration methods. Although the formulation of NMI is extremely general, care must be taken when comparing values due to the non-linearity between NMI value and probabilistic image similarity. Since the normalised mutual information by itself may only be used to provide a relative measure of image similarity, the CFMM must also be used in this way. Figure 6.4 shows the percentage change in CFMM value for each patient dataset shown in Table 1.1 after registration to the first image in the dataset and after PPCR for all iterations (PPCR-end). Also shown is the maximum CFMM value if the PPCR algorithm was terminated at the iteration with the highest CFMM value (PPCR-best). In all cases the CFMM value is higher for PPCR than for a simple, single fluid registration. This is a result of the formulation of the measure, which rewards circumstances in which all images are similar to one another, rather than all being similar to a single image. The PPCR algorithm may be expected to give higher values since component registrations proceed to anchor images that contain information from the entire time-series.

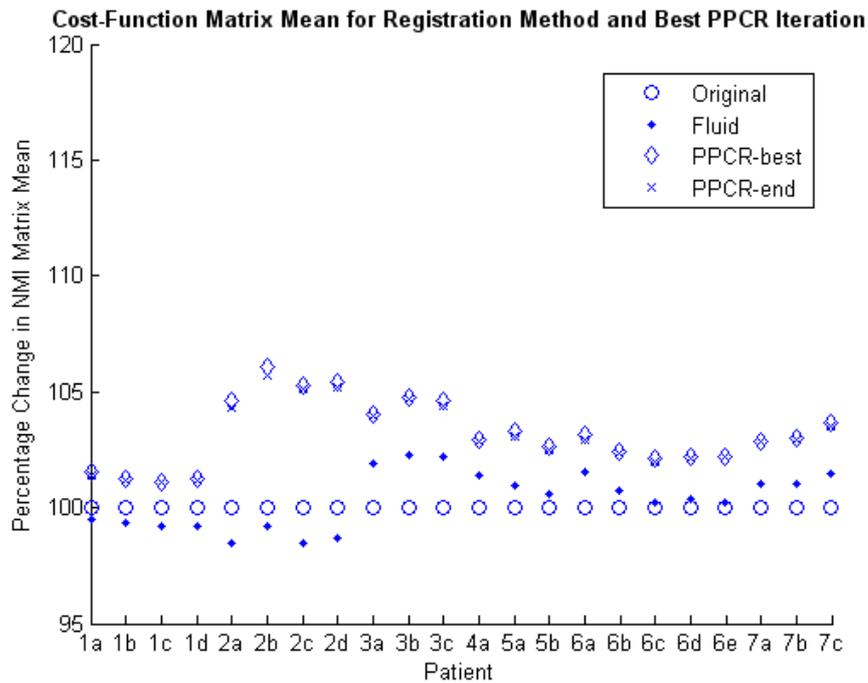


Figure 6.4: Corresponding percentage change to the Cost-Function Matrix Mean (CFMM) values for Unregistered Data, Registration by direct-fluid registration and by PPCR. Also shown for PPCR is the highest CFMM found during the iterations.

6.1.4 Conclusion

The advantage of the PPCR method over conventional registration is that it allows the use of information from the entire dataset to guide the image-wide deformations. In the case of DCE-MRI data, the

early registrations are guided by the simplest uptake profiles (the early eigenvectors explain the biggest changes in intensity) that explain the data. This allows a registration process to take place which is refined with each iteration. The early, simple uptake profiles are expected to be relatively noise free because early principal components describe the major, long-term trends in the data. These components provide little information about the true pharmacokinetic uptake profiles of individual pixels. Although it is difficult to extract useful pharmacokinetic information from the principal components in this form, they represent a suitable way to generate intermediate data during a refining, iterative registration. Principal component sets calculated from the partially registered data are successively less likely to have random motion noise in later components and this is the justification for the use of higher components in later registrations.

The method relies on being able to separate motion and contrast-enhancement artefacts in order to proceed successfully. If the registration cannot successfully do so, it may fail. Conventional registration of DCE-MRI is complicated by the changing intensity structure of the images and simple image-to-image registration methods may fail, producing artefacts, due to the difficulty of selecting appropriate target images. Progressive Principal Component Registration allows image-by-image registration to a partially compressed dataset in which motion artefacts are suppressed in a series of target images generated to resemble the original dynamic data. The use of principal components analysis circumvents the use of a pharmacokinetic model not only avoiding the problems of accurate model-fitting but permitting the use of much larger target areas of the dataset. Therefore, PPCR is a successful method for the model-free registration of large region-of-interest DCE-MRI datasets.

6.2 Registration of 3D DCE-MRI Using PPCR

6.2.1 Introduction

This section continues analysis of the PPCR algorithm using the six 3D DCE-MRI datasets from Table 1.2. These datasets are at a temporal resolution of 12s consisting of a 6s held-breath volume acquisition and a further 6s breathing interval. The slice profile in the outside slices is particularly poor, and although included in subsequent registrations we will not analyse these regions further. The passage of the contrast agent bolus is particularly well defined in these datasets (see Figure 6.5), demonstrating the bolus passage through the heart with two defined periods of hepatic enhancement. Since these images are acquired at breath-hold, motion between subsequent timepoints due to breath-hold depth inconsistency remains a problem, and will impact on pharmacokinetic analysis.

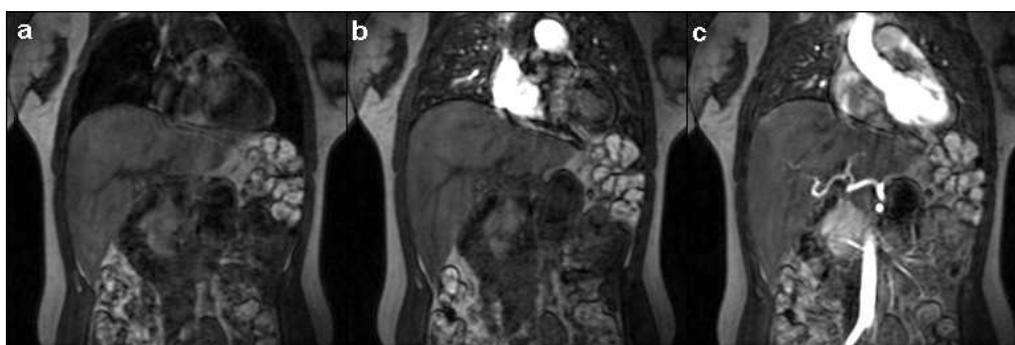


Figure 6.5: Example Contrast Enhanced Images of the Abdomen. Pre-enhancement image **a** for comparison with image showing passage of bolus through heart **b** and subsequent portal enhancement of the liver **c**.

6.2.2 Method

Because the PPCR algorithm uses no information about pixel spatial location, only the number of pixels, the extension of the PPCR method to 3D is trivial, no modification is needed other than those considering memory management which are not fundamental to the algorithm. As such, the intermediate fluid registrations are extended to full 3D and run-times are kept low by implementing the algorithm in parallel form for use on a computer cluster; the datasets may be registered within a few hours (i.e. overnight).

The datasets used here are of both higher spatial and temporal resolution than those used previously, and we can apply a full pharmacokinetic analysis to the liver in order to extract pharmacokinetic parameters. Analysis of these parameters, and the error on these parameters, before and after registration will indicate the performance of registration by PPCR. The DCE-MRI image analysis package is the MRIW software provided by the Institute of Cancer Research ((d'Arcy et al., 2006), (Parker et al., 1998))

The MRIW software will be applied over a manually segmented region of the liver from the central slice of both the pre and post-registration images, an additional registration of a low-flip angle image to the first pre-enhancement image is also included for T_1 estimation. A pharmacokinetic model is applied to each pixel to find parameter maps for values of K^{trans} , v_e and the hepatic perfusion index (HPI), for each registration method. The pharmacokinetic model will incorporate a hepatic dual cosine arterial

input function composed of an hepatic arterial term and a portal term (Equation 3.11 to 3.13 discussed in Chapter 3 (Woolrich et al., 2004), (Orton et al., 2008)). The dual blood supply to the liver is fitted by finding the HPI and the constants of the dual-cosine arterial input function are shown in Table 3.2.

6.2.3 Analysis by MRIW

For direct fluid registration and PPCR (and where available a proprietary registration algorithm), the MRIW software is used to calculate parameter maps for K^{trans} , v_e , and the hepatic perfusion index, HPI and the pixel residual. These are displayed in the following figures.

Figure 6.6 shows the results for Patient 2. Registration by PPCR demonstrates improved model-fitting by way of a reduction in the pixel-residual maps.

Figure 6.7 for Patient 5 contains parameter maps calculated from a dataset with a relatively large amount of motion. Motion artefacts present themselves in the images as discrepancies at the boundaries of features such as the large hepatic masses and the diaphragm. Some residual motion remains at the superior boundary of the liver. Due to the residual differences in liver position between some frames, the model-fitting routine is unable to successfully fit to these areas when they occur inside the blue boundary. Also included in this Figure is the result of a proprietary software non-rigid registration provided with the MRI scanner used to acquire the results. Unfortunately without details of the algorithm it is difficult to tell why the registration is performing poorly, as seen in the pixel-residual maps which introduce model-fitting errors for the inferior liver and the diaphragm. Direct image registration confers no obvious improvement to the parameter maps. There are also additional acquisition related artefacts due to the inclusion of an obsolete scanner software patch designed to adjust zero-filled data in k-space. These artefacts are also present in Figure 6.8.

Figure 6.8 is a further dataset in which the in-built scanner registration method is used. For the pixel-residual maps of this patient we see a minor improvement using the PPCR algorithm but for the in-built registration method we see a large increase in model-fitting errors. The reason for the poor pixel-fitting of the in-built algorithm is clear when the images are inspected: the centre of the liver is actively mis-registered between images and so in this case the results should not be used.

The total reduction in model fit residual for each dataset is summarised by Table 6.2. The results for each dataset summarise the visible changes to the model-fit residuals in the segmented hepatic regions of Figures 6.6 to 6.8. Direct fluid registration often achieves little improvement in residual model-fit. This is in contrast to the PPCR method which is often able to reduce the model-fit residuals by over 10% across the liver.

Three example datasets are included on the supplementary CD (see Appendix E). All movie files have the same format from left to right. The left-most movie is the original, unregistered DCE-MRI dataset. The second is registration using direct fluid registration using cross-correlation to the first (un-enhanced) image in the dataset. The third is the result of registration using the PPCR algorithm. If a fourth movie exists, this is the result of a scanner-based image registration algorithm. Since the data from Table 1.2 is 3D, only the central slice is shown. The file *movie-3-04.avi* demonstrates minor improvements to the correction of liver position between breath-holding (compare with Figure 6.6). The

Table 6.2: Change in total model fitting residual error for each dataset in Table 1.2 after each registration method for the segmented regions in Figures 6.6 to 6.8.

Dataset	Unregistered	Direct Fluid Registration	PPCR	Proprietary Registration
1	100%	-2%	-19%	
2	100%	-6%	-16%	
3	100%	-3%	-8%	
4	100%	0%	-13%	
5	100%	+1%	-19%	-4%
6	100%	+2%	-17%	+10%

files *movie-3-05.avi* and *movie-3-06.avi* include the results of erroneous registration by the in-built algorithm (fourth movie from the left) and minor improvements in registration are seen for PPCR (third movie from the left). Mis-registration of the central region of the liver in both these examples leads to an increase in model-fit residual seen in Figure 6.7 and 6.8 respectively.

It is also possible to inspect for some indication of what difference registration makes to the estimates of DCE parameters such as K^{trans} . By selecting small regions of interest it is possible to analyse changes to parameter statistics before and after registration. The heterogeneity of regions of interest, particularly pathology, makes the interpretation of changes to cursory statistics difficult. Because of this, we inspect histograms of the K^{trans} parameter to reveal changes over the region of interest. Figure 6.9 shows the selection of three hepatic regions of interest within dataset 6 to be analysed further (see Figure 6.8 for comparison). Each region displays a weakly enhancing core and a rapidly enhancing boundary, characteristic of tumour tissue. Figure 6.10 plots histogram statistics for K^{trans} for each of the three regions. For each region prior to registration there is a binomial distribution of the K^{trans} parameter reflecting low values in the centre of the region of interest surrounded by higher values. Registration by direct fluid registration and PPCR maintain this distribution. Median K^{trans} values are observed to decrease slightly by 3-4% after registration by PPCR, this may be due to a reduction of intermediate K^{trans} values between tumour core and tumour rim brought about by improved alignment. Model fit residuals in these regions are reduced by approximately $\frac{1}{3}$ in each case using PPCR. The scanner based algorithm performs poorly and resolution of the binomial K^{trans} distribution is lost in the three cases; the median K^{trans} value is increased by 20% for region 3 suggesting mis-registration has lead to over-estimation of this parameter. This is visible as a loss of contrast between region centre and region periphery in the associated K^{trans} image in Figure 6.8.

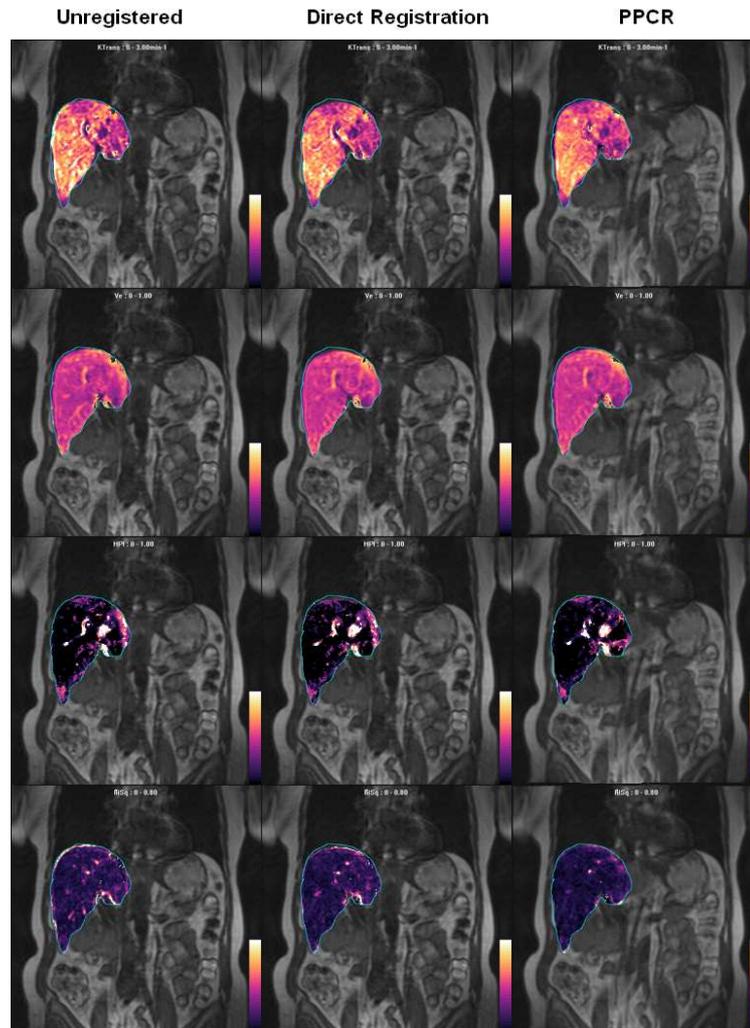


Figure 6.6: Patient 2: Model fitting results using dual-cosine arterial input function showing parameter maps for (**rows**): $K^{trans}(min^{-1})$ (range 0-3), v_e (range 0-1), HPI (range 0-1) and Pixel Residual (range 0-8). Each column represents registration by **Column 1**: Unregistered Data, **2**: Direct Fluid Registration and **3**: PPCR.

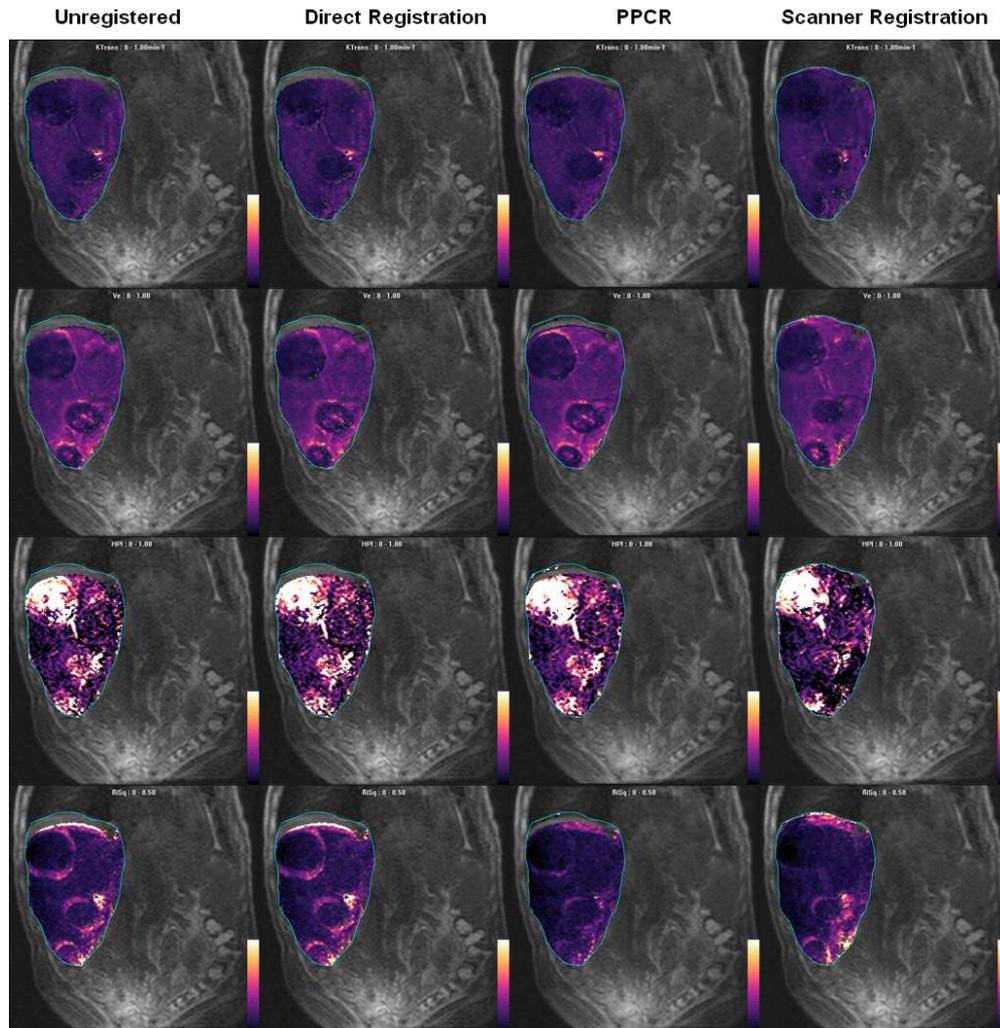


Figure 6.7: Patient 5: Model fitting results using dual-cosine arterial input function showing parameter maps for (rows): K^{trans} (min^{-1}) (range 0-1), v_e (range 0-1), HPI (range 0-1) and Pixel Residual (range 0-0.5). Each column represents registration by **Column 1**: Unregistered Data, **2**: Direct Fluid Registration, **3**: PPCR and **4**: Automatic registration on scanner using unknown algorithm.

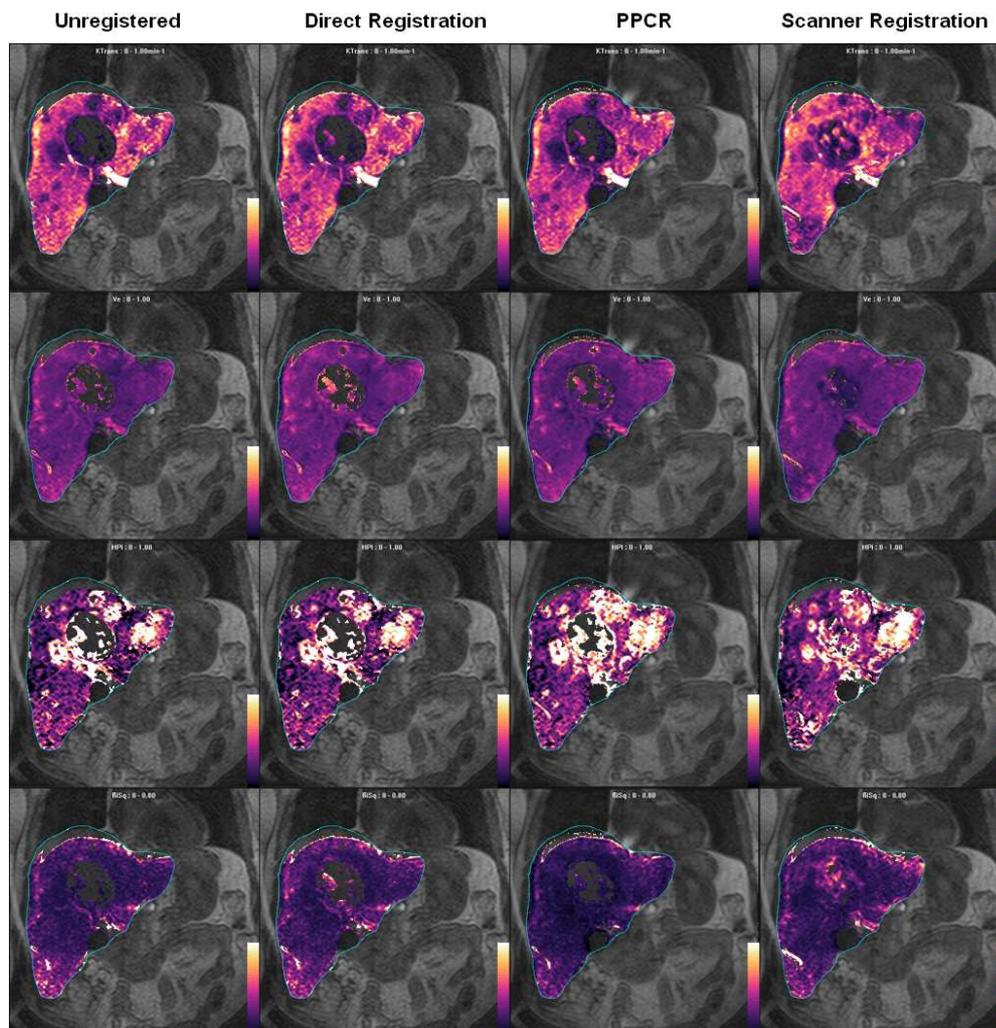


Figure 6.8: Patient 6: Model fitting results using dual-cosine arterial input function showing parameter maps for (rows): K^{trans} (min^{-1}) (range 0-1), v_e (range 0-1), HPI (range 0-1) and Pixel Residual (range 0-0.8). Each column represents registration by **Column 1**: Unregistered Data, **2**: Direct Fluid Registration, **3**: PPCR and **4**: Automatic registration on scanner using unknown algorithm.

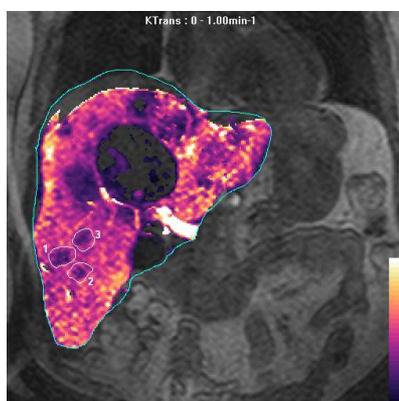


Figure 6.9: Selection of local regions of interest for analysis of changes to pharmacokinetic parameters before and after registration. See Figure 6.8 for statistics over the whole liver.

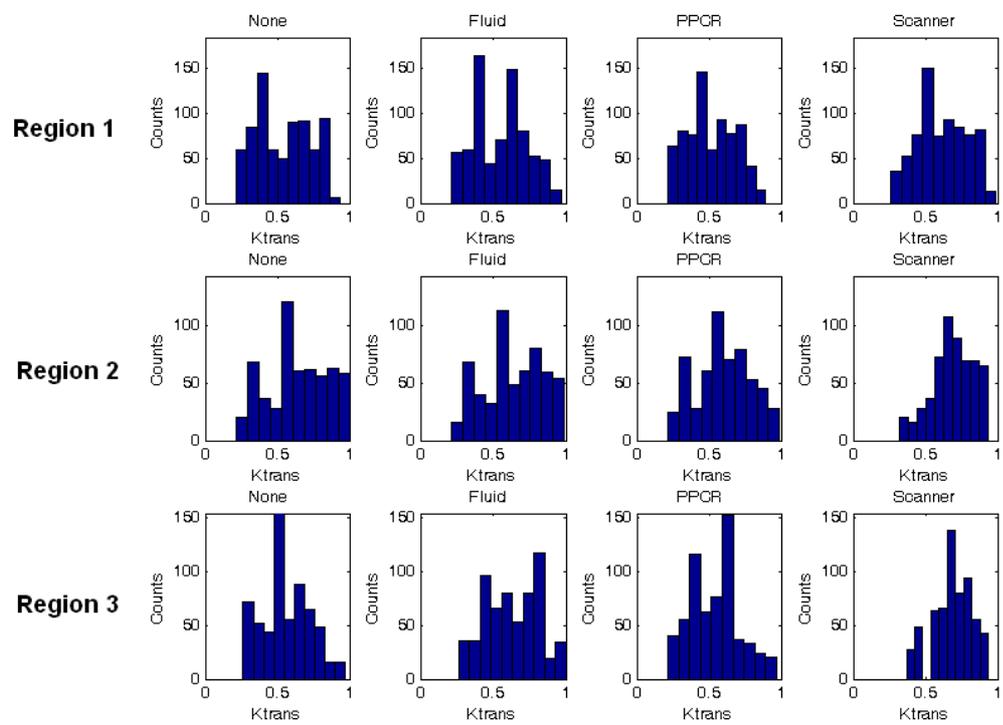


Figure 6.10: K^{trans} (min⁻¹) parameter distribution histograms for each of the three regions (rows) in Figure 6.9. Columns: results after each registration method as labelled.

6.2.4 Analysis by Cost-Function Matrix Mean

The normalised mutual information cost-function matrix mean (NMI-CFMM) can also be used to assess the performance of 3D registration by PPCR. Each of the registration methods analysed above has a CFMM value calculated. The differences between registration methods are shown in Table 6.3 as absolute NMI-CFMM values and with the percentage change compared to the unregistered case shown in brackets. These results are broadly in line with the conclusions from the parameter maps of the previous section. Direct registration often produces a small improvement to the CFMM value. Registration by PPCR will allow double the improvement. The registration of the two datasets using the in-built scanner registration algorithm does not provide an improvement. Care must be taken with these results, since as discussed before, maximisation of NMI does not necessarily correspond to correct image registration.

Table 6.3: DCE-MRI Patient 3D Scan Data. NMI Cost function matrix mean values after registration by each method (the percentage improvement relative to the unregistered case is shown in brackets).

Patient	Un-registered	Fluid	PPCR	Siemens
1	1.211	1.229(+1.5%)	1.255(+3.6%)	
2	1.194	1.218(+2%)	1.242(+4%)	
3	1.187	1.205(+1.5%)	1.228(+3.5%)	
4	1.183	1.206(+2%)	1.227(+3.7%)	
5	1.20	1.20(+0%)	1.22(+1.7%)	1.21(+0.1%)
6	1.18	1.18(+0%)	1.20(+1.7%)	1.18(+0%)

6.2.5 Conclusion

Results of this section suggest that PPCR may be applied to 3D data. The datasets used in this section are at higher temporal resolution and images are acquired at breath-hold. The time penalty of the PPCR algorithm has been offset by using a computing cluster to submit registrations in parallel at each principal component iteration. Full PPCR registrations take approximately 4hours (for datasets of 20 images) and as a result the time penalty is not prohibitive. Use of the model-fitting algorithm in the MRIW software has allowed an estimation of the improvement made by registration (and its effect on the model-fitting). The PPCR algorithm has been shown to allow improved model-fitting by reduction in the model-fit residuals. The PPCR algorithm gives a quantifiable benefit over other registration methods.

Chapter 7

Registration of Diffusion Weighted MRI

The acquisition of Diffusion Weighted MR images may be confounded by both patient motion and machine eddy currents. In the brain, the resulting images are often corrected using an affine registration, which is often thought appropriate in the brain due to the nature of the artefacts. Here, two non-rigid registration schemes are compared to the result of affine registrations: a single fluid registration of the individual diffusion directions; and a Progressive Principal Component Registration. All registrations are full 3D. Twelve DW-MRI datasets consisting of 128x128x64 volumes from 15 diffusion directions are registered by each method (see Section 1.3) and the different results combined to produce fractional anisotropy maps. These maps are then inspected for improved feature appearance and artefact reduction. The affine registration demonstrates a modest improvement in the twelve cases. Image alignment by single fluid registration causes lateral brain features to appear sharper at the expense of poor deformations of the medial brain. Registration by PPCR demonstrates both improved demarcation of lateral brain features and preservation of medial features such as the corpus callosum. Figure 7.1 shows the b0 image and 15 diffusion directions (labelled) for a slice from one of the pre-registration datasets.

7.1 Introduction

Diffusion Weighted MRI is an important tool for brain connectivity imaging and is increasingly being applied to other organs of interest. However, due to the length of a scan and eddy currents, diffusion weighted images often become spatially misaligned. Subsequent analysis of the images, such as their combination into a fractional anisotropy map or principal diffusion direction image is then compromised. Within the brain, eddy current distortion is regarded as being geometrical and can be countered by an affine registration of each diffusion direction image into a common coordinate system. Since scanner eddy current induced distortion artefacts should only cause relative scaling, shifts and shearing between images (and only in two dimensions) and the patient can only move their head with approximately rigid translations and rotations, we would expect affine registration to be appropriate. If we were to use non-rigid registration we are likely to find that since features such as fibre tracts appear in multiple diffusion images with different contrast (dependent on the gradient of the diffusion direction), that local registration of these regions would be difficult. This is analogous to the contrast changes due to contrast-agent concentration in DCE-MRI disturbing the assumptions of dispersion in the joint image

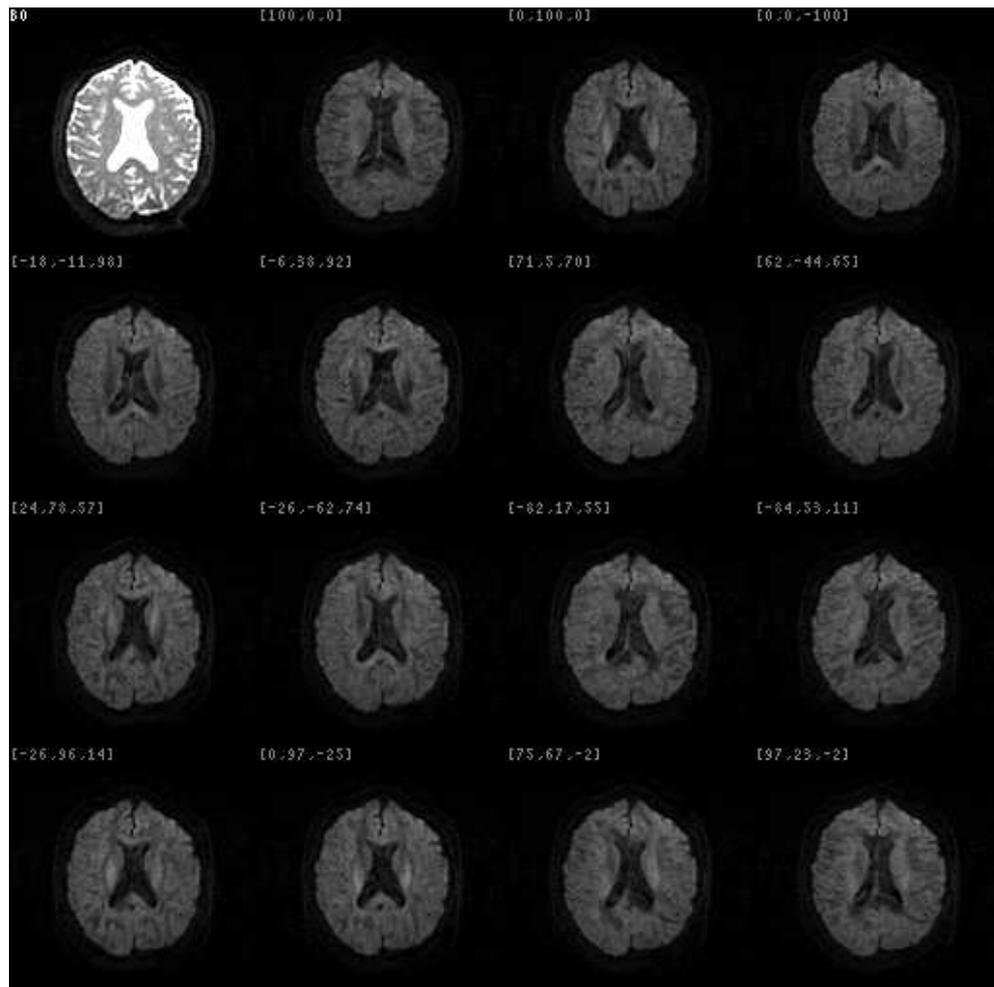


Figure 7.1: Example DWI Slice, comprised of a B0 Image and fifteen diffusion direction images (with gradient direction labelled)

histogram. Whether the registration scheme is affine or non-rigid, registration will often be applied to images containing the same features with different contrast which may be a confounding factor. Resistance of the registration to artefacts will be improved using an affine registration, this may be one of the reasons it is used in preference to non-rigid registration, since an affine transformation permits no local changes. However, minimisation of the cost-function in an affine registration still leaves a susceptibility to contrast-change induced mis-registration. Since it is problematic to register different diffusion direction images to each other (and therefore assign prominence to a particular direction), we must register them to a common target image, often the b_0 image which often has a better signal-to-noise ratio but itself contains different contrast and information to the diffusion direction images. Attempts have been made to register individual gradient direction images ((Leemans et al., 2005), (Tao & Miller, 2006)) and there is also a large amount of work on the reorientation of tensors after their calculation ((Alexander et al., 2001), (Guimond et al., 2002), (Hecke et al., 2007)).

The fundamental diffusion imaging equation is given by Equation 7.1 where the signal for a given b -value (in units of $s.mm^{-2}$) is given by this relationship given the non-weighted signal S_0 and gradient direction \mathbf{g} . The diffusion tensor for a given region, \mathbf{D} , can be found by varying the gradient orientations. In three dimensions we have a 3x3 Tensor but since reflected gradient directions are redundant (i.e. $D_{xy} \equiv D_{yx}$), we need a minimum of 6 gradient values to determine the tensor, plus a b_0 map. In practice, many more gradient directions are often included. Performing an eigenvector analysis on the diffusion tensor produces a principal vector representing the dominant diffusion direction, which can be used to infer nerve-fibre orientation.

$$S(\mathbf{g}) = S_0 e^{-b\mathbf{g}^T \mathbf{D} \mathbf{g}} \quad (7.1)$$

The work in this Chapter applies the Progressive Principal Component Registration (PPCR) scheme (Melbourne et al., 2007b), a method that uses a principal components analysis to generate target images from a set of images containing the same information with changing contrast. In the case of DW-MRI, the method makes use of the overlapping image information from different gradient orientations to perform registration. The method applies a non-rigid fluid registration at each step as used in previous chapters (see (Crum et al., 2005)), which allows the potential for the registration of inter-subject cranial Diffusion Weighted MRI, registration to an atlas or extra-cranial Diffusion Weighted MRI. It is possible that the registration of the images may require the re-orientation of the gradient direction associated with each image; this can be done in a straightforward manner when using affine registration. For the non-rigid case, a change in direction can be found from the local deformation field. For a given pixel, we find the rotational component of the Jacobian of the deformation. For the results in this chapter we find the rotational components to be extremely small (sub one-degree) and so do not reorientate the gradient direction.

7.2 Method

We apply the PPCR method (Melbourne et al., 2007b) to 12 datasets. Each dataset consists of a 128x128x64 volume acquired from 15 diffusion directions with a b -value of $1000s.mm^{-2}$ and a corre-

sponding b_0 volume. The PPCR method is implemented in a full-3D implementation and is compared to both a 3D affine registration of each diffusion gradient direction volume to the corresponding b_0 volume and a 3D single fluid registration of each diffusion direction to the first diffusion direction - the registration cost function used in the case of affine registration and single fluid registration is Normalised Mutual Information (Studholme et al., 1999), the similarity measure that should be most appropriate here since the pixel intensity relationship is unknown. The affine registration provides a comparison of the PPCR method to an existing registration method and the single fluid registration allows a comparison of the PPCR and affine methods to a direct non-rigid registration scheme.

Application of PPCR to DW-MRI is analogous to its application to DCE-MRI. As with DCE-MRI, conventional cost-functions will seek to minimise artefacts due to both motion and contrast change with equal weight leading to poor registration. PPCR allows the generation of target images that are contrast matched to the original images. This is because there is shared information between images with different diffusion gradient directions - the gradients directions overlap with at least one of three orthogonal coordinate directions. Motion artefacts in particular do not overlap with gradient direction (although eddy currents will), they are unique to a particular image and are therefore unlikely to appear in early principal components. Figure 7.2 shows the results of rebuilding the dataset for three images (top row), using (on each successive row) 1-4 principal components. In the case of diffusion images, principal components represent functions of signal intensity with direction. For a region of restricted diffusion a smoothly varying function of diffusion strength with angle might be disrupted by motion in individual images. In this case early principal components represent pixels with well-defined directional information and later components may be used to remove the image noise that is largely unique to a particular diffusion direction. Information from pixels with multiple dominant diffusion directions is likely to be encoded in a handful of early principal components; provided that noise does not dominate the numbers of pixels with well-defined signal-angle profiles, registration by PPCR may proceed.

A reminder of the fundamental PPCR equation is shown in Equation 7.2, in which the result (the new best registered data) at a given iteration, n , is given by the registration of the best registered data from the previous step, registered to artificial images generated from a temporal principal components analysis of the best registered data from the previous step, rebuilt using n principal components. Registration of source images to artificial target images uses a fluid registration algorithm (Crum et al., 2005).

$$\mathbf{A}_{ij}^{n+1} = \mathbf{A}_{ij}^n \mapsto \sum_{m=1}^{n < T} (\mathbf{U}_m^n \cdot \mathbf{A}_{ij}^n) \mathbf{U}_m^n \quad (7.2)$$

The results of the three registration methods are combined into both fractional anisotropy maps and principal diffusion direction RGB images so that differences in the resulting features can be observed. Due to the combination of the 15 separate, registered diffusion direction images into a single fractional anisotropy map, it is difficult to analyse the direct contribution of the individual deformation fields to the final result.

A comparison of the consistency of the fractional anisotropy before and after registration can be used to provide evidence of registration success. Work by Bai (Bai & Alexander, 2008) compared fractional anisotropy values calculated from two sets of 30 images from a 60 diffusion direction dataset

before and after registration. We choose to use a leave-one-out analysis since we are also constrained by having only 15 direction images; this also removes the influence of bias when choosing which combination of directions to leave out. Calculation of the fractional anisotropy uses combined information from the remaining 14 different diffusion directions. For each diffusion direction in turn, we calculate the fractional anisotropy without that direction. This produces 15 values of the fractional anisotropy for each pixel. If the registration is successful we would expect the fractional anisotropy to be consistent despite the removal of a particular direction; hence we inspect the standard deviation of the fractional anisotropy values for each pixel over a central region of the brain (as defined by an ellipsoidal binary mask of radius $\frac{1}{4} \times \frac{1}{4} \times \frac{1}{2}$ of the image dimensions in $[x, y, z]$ on the centre of the image) on the assumption that smaller variations in FA value correspond to good alignment of those diffusion directions used in its calculation. The central region of the brain refers the result of a mask formed by a sphere located in the centre of the image. Larger image mis-alignment errors occur in the lateral brain, some of which may not be included as a result of masking so the following results might not show the increased benefit of image registration in these errors. The sphere size is defined so that spurious tensor calculations in the skull and air are ignored.

Also provided are the results of the tensor fitting calculation using the least-squares fitting method (see (Kingsley, 2006)) in which the system matrix is comprised from the gradient vectors. The total tensor-fit residuals over the central region of the brain are shown for each dataset. The central region is calculated as described above.

7.3 Results

Qualitative registration results are presented in Figures 7.3, 7.4 and 7.5. These figures are best viewed electronically. The images presented here demonstrate image registration success, or otherwise, by allowing inspection of the brain features calculated after conversion of the diffusion direction images into fractional anisotropy and principal diffusion direction images (Melbourne et al., 2008b). Quantitative registration results are provided in the following sections.

7.3.1 Visual Inspection of Fractional Anisotropy Maps

Figure 7.3 demonstrates the benefit of image registration in a subject in which there is substantial motion between diffusion directions. The resulting fractional anisotropy calculations in this region, particularly in the anterior brain, provide little fibre tract information. The implementation of affine or fluid registration improve the demarcation of fibre tracts slightly, but the application of the PPCR method appears to have well-recovered fibre-tracts in this area.

Figure 7.4 is a slice from the same subject but superior to that in Figure 7.3. Improved feature resolution and apparent visual noise reduction are seen throughout the slice using the affine and fluid-based registrations, but PPCR appears to show further improvements, particularly in the lateral brain.

Figure 7.5 from a different subject to those presented above, but with less severe artefacts as seen in the unregistered fractional anisotropy images. Registration of the diffusion direction images allows

improved artefact reduction, for instance as seen here in the ventricles. Analysis of the fluid registration result seems to show mis-registration of the corpus callosum. This feature appears to be the result of mis-registration as a result of registering component direction images that contain strong signal in one image and weak in the other relative to the global intensity level. As a result the strong image features are contracted in the resulting registration. The absence of this feature in the other registration methods suggests that problems may be encountered in naive direct registration. Despite this error, fluid registration shows improved feature resolution in the medial brain, but this is matched by the PPCR result, and exceeded, particularly with reference to the direct-registration failure in the corpus callosum.

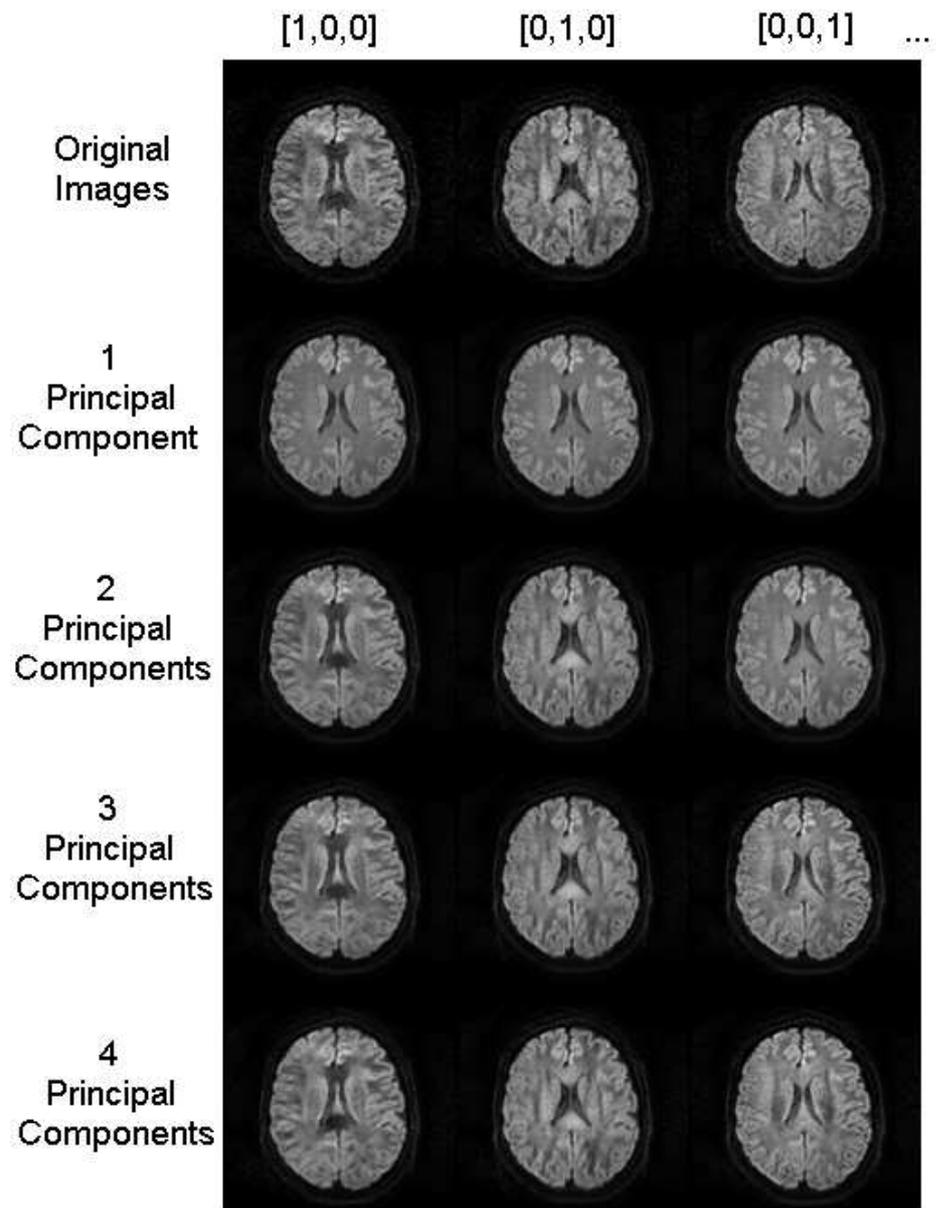


Figure 7.2: **Top row:** Images from a DW-MRI dataset for gradient directions 1 to 3. **Subsequent Rows:** Images rebuilt using, on each row, 1,2,3,4 components respectively.

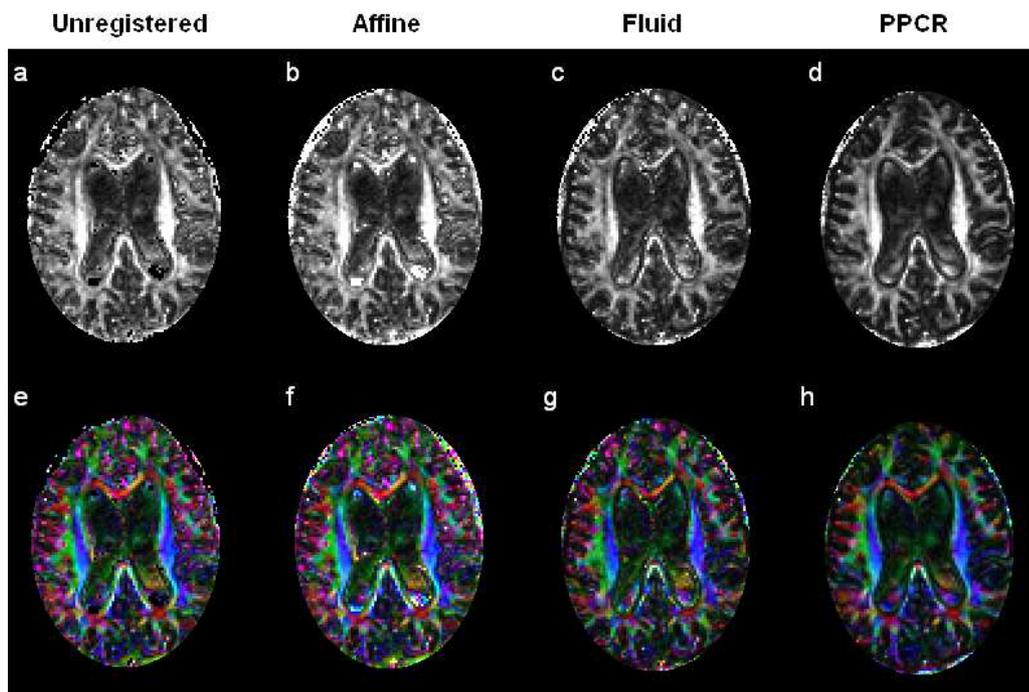


Figure 7.3: (Subject 175) Demonstration of apparent improved pixel resolution and feature demarcation with registration method in Fractional Anisotropy (FA) and Principal Diffusion Direction Images. Note increased demarcation of features in the anterior brain, with increasing improvement using affine, fluid and PPCR registration protocols respectively. Also note a decrease in spurious features in the ventricles with registration method.

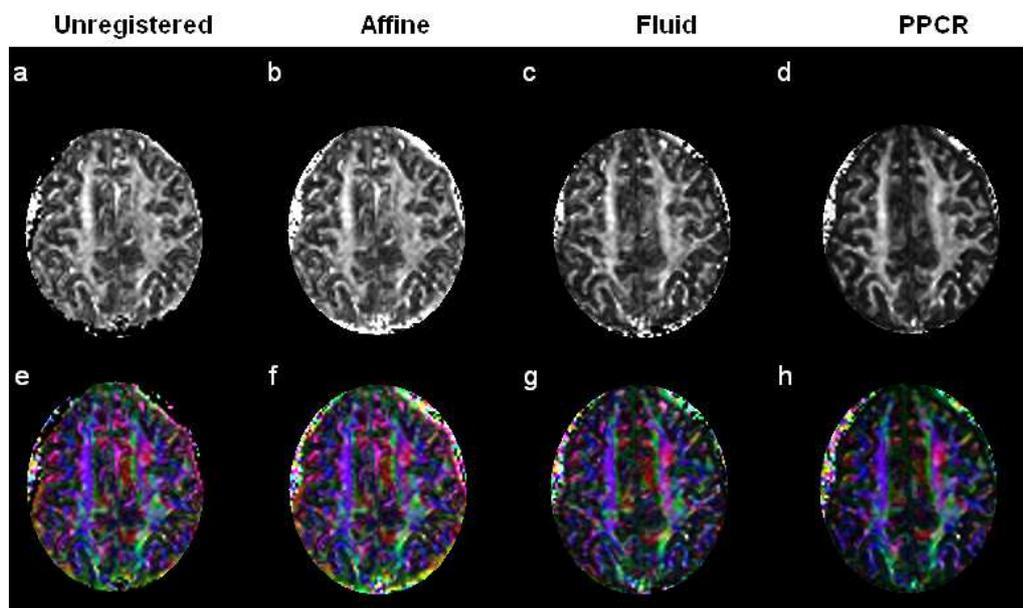


Figure 7.4: (Subject 175) Demonstration of apparent improved pixel resolution and feature demarcation with registration method in Fractional Anisotropy (FA) and Principal Diffusion Direction Images. Note improved resolution of features throughout this brain slice with increasing registration algorithm complexity. Serious motion artefacts associated with this particular patient have been removed most successfully using PPCR, making the dataset suitable for further analysis.

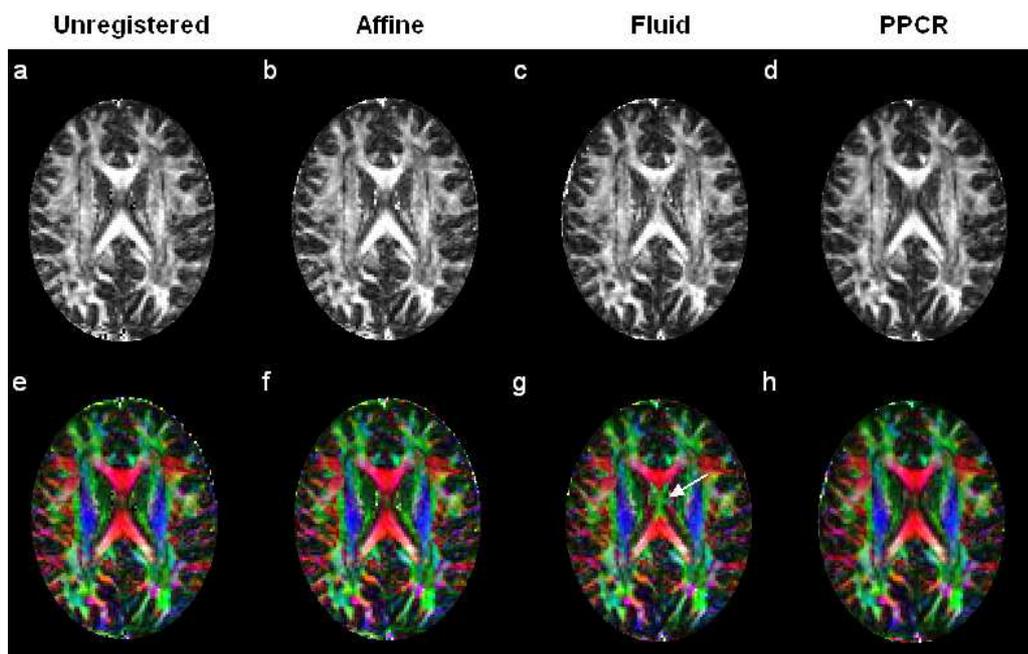


Figure 7.5: (Subject 52) Demonstration of apparent improved pixel resolution and feature demarcation with registration method in Fractional Anisotropy (FA) and Principal Diffusion Direction Images. Fractional Anisotropy Images **a** to **d**: note disruption of corpus callosum using simple fluid registration, but preservation under PPCR. This disruption is likely to be due to erroneous through plane registration of the corpus callosum edges in the diffusion direction images, **e** to **h**. Note also reduction in noise in ventricles under PPCR in comparison to other methods. Definition of lateral brain features is enhanced in both non-rigid registration algorithms.

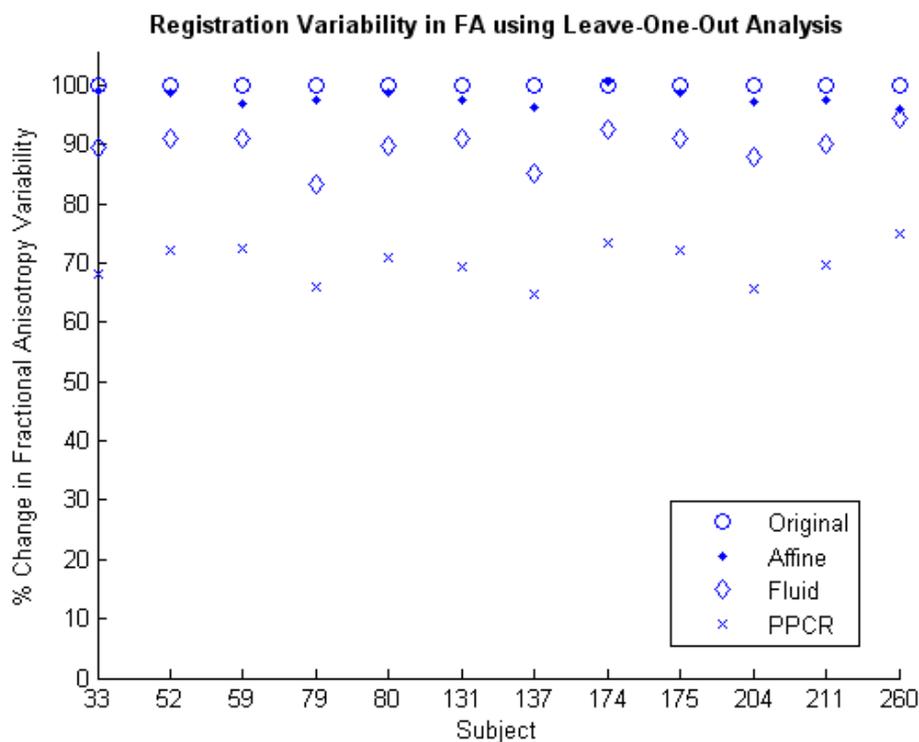


Figure 7.6: Analysis of variation in Fractional Anisotropy value for a leave-one-out analysis of the fifteen diffusion directions. Values are expressed as a percentage of the original FA standard deviation for each registration method *see text*.

7.3.2 Analysis of Fractional Anisotropy Variation

Figure 7.6 contains the results for the leave-one-out fractional anisotropy consistency analysis described in the methods section. In all cases, the PPCR algorithm is seen to improve the fractional anisotropy calculation by reducing the variability in its result by up to 35%. Registration by the affine method produces only a modest improvement suggesting that the registration algorithm is unable to cope with the type of deformations required. Fluid registration produces an improvement over the affine case. The calculation of the fractional anisotropy variability in this way suggests that the non-rigid registration methods are producing both visually and quantitatively superior results.

The results in Figure 7.6 may contain the effects of an overall reduction in fractional anisotropy across the region of interest. Motion artefacts in regions of low signal are likely to cause the resulting FA value to appear high. Reduction of these motion artefacts reduces the observed anisotropy, reducing the observed FA in these regions. Conversely, a reduction of motion artefacts due to the image registration procedure in regions of high anisotropy might have less impact on the calculated FA which will remain high. The combined impact should result in improved contrast, partly explaining the improved feature demarcation in Figures 7.3 to 7.5.

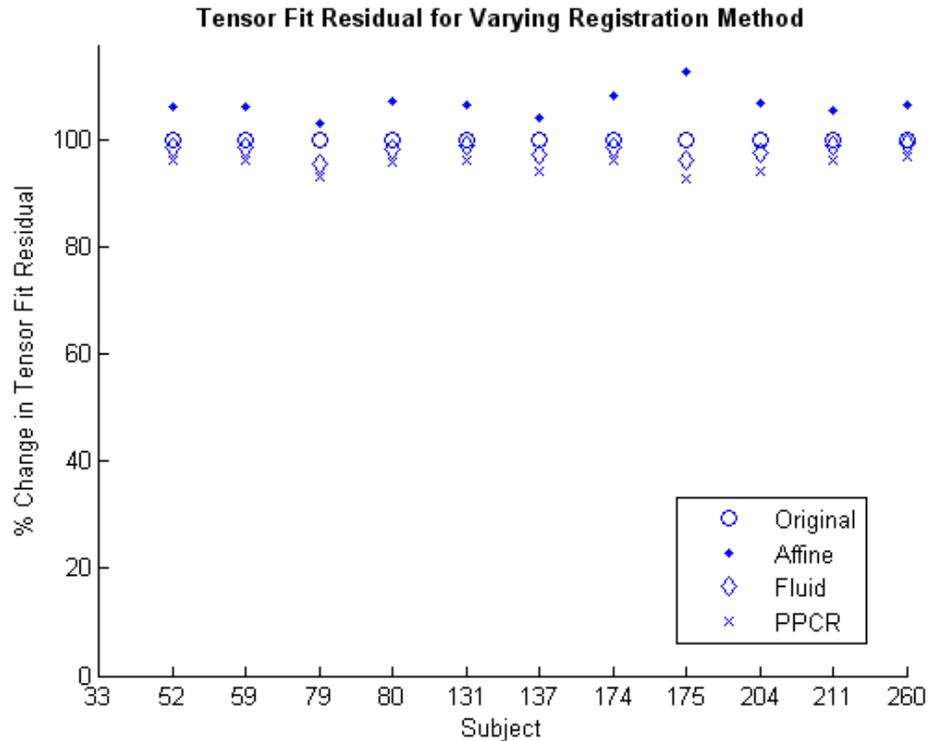


Figure 7.7: Analysis of tensor fitting residual after fitting to 15 diffusion directions. Values are expressed as a percentage of the original residual of the unregistered images for each registration method.

7.3.3 Analysis of Tensor Fitting Residuals

Figure 7.7 contains the results of the residual in the tensor fitting procedure in signal space. Values represent the average residual over a central sphere of the dataset as described above. Affine registration in this case produces results that make tensor fitting more difficult. This may be the case at the extreme edges of the sphere due to the constrained nature of the deformation and the resulting increase in displacements away from the image centre. Non-rigid registration can be seen to allow improved tensor fitting for both direct fluid registration and PPCR.

Affine or Non-Rigid Image Registration?

An important question is the use of non-rigid image registration in an application in which affine registration is considered adequate. The PPCR algorithm does not explicitly require a non-rigid image registration method and therefore we may in future substitute an affine registration method. The operation of the PPCR algorithm allows conventional image similarity measures to be used on images with varying contrast in non-rigid applications where local volume change would otherwise be a problem. The imposition of a global affine deformation model is likely to be reasonably robust to the local contrast changes seen when inspecting DCE-MRI or DW-MRI images although the cost-function may still be inappropriate. In an affine registration there can be no local contraction of enhancing features that might result from non-rigid registration and so there should be little justification for a PPCR-affine algorithm. The results of the previous sections suggest that PPCR is able to successfully non-rigidly register

images in which there is reason to believe the number of degrees of freedom is restricted. In the case of atlas-based image registration and inter-subject registration, non-rigid registration will be necessary.

7.4 Conclusion

The PPCR method demonstrates improved registration of Diffusion Weighted MR images when compared to simple affine registration or when using a naive fluid registration scheme. The use of affine registration in the alignment of cranial DW images is regarded as adequate, but this is particularly unlikely to be the case for extra-cranial organs such as the liver. The computational time penalty of the PPCR method, when compared to affine registration, is large and a handicap when considering that improvements with affine registration are marked and detectable. However, the PPCR method has a higher success rate when compared to the single fluid registration method and is preferable in this instance. The PPCR method works as a result of combining overlapping information from the non-orthogonal gradient directions. Without this overlap, the images would not be suitably similar in order to generate principal components indistinguishable from noise and the PPCR registration would likely fail. The method allows enhanced feature detection and reduced noise by reducing motion and eddy current artefacts, which is a benefit to subsequent analysis such as calculating the fractional anisotropy (as shown) and also when considering tractography. The use of PPCR on much higher angular resolution data might be limited by the increasing violation of the linearity assumption of the principal components, an isotropic distribution of orientated pixels will no longer be present. However, determination of the exact threshold is left as future work.

7.5 Diffusion MRI of the Liver

7.5.1 Introduction

Diffusion Imaging is increasingly finding application in oncology to measure diffusion coefficients in regions of tumour. Here the emphasis is not on taking many diffusion gradient directions to form a directional diffusion tensor, but on quantifying the isotropic diffusion coefficient (in this case the Apparent Diffusion Coefficient, ADC). Equation 7.3 shows the diffusion signal with a b -value, b and ADC value, D given the un-weighted signal S_0 . In order to find the ADC, data from different b -values are taken. Hence an ADC map can be found for example as in Equation 7.4 for two different b values. For many b values, the ADC can be found by least-squares fitting, which will be presented further in the methods.

$$S(b) = S_0 e^{-bD} \quad (7.3)$$

$$ADC = \frac{1}{b_2 - b_1} \log\left(\frac{S(b_1)}{S(b_2)}\right) \quad (7.4)$$

ADC measurements have been investigated for both tumour determination and for the assessment of changes to ADC brought about by anti-angiogenic therapies. Disruption to high density neo-vasculature cannot easily be said to either increase or decrease the ADC although necrosis is considered to increase the local ADC (Provenzale et al., 2006). The disorganised growth of a tumour is believed to result in isotropic diffusion without a preferred orientation over the size of an imaging pixel. The relatively highly cellular environment of tumour tissue, when compared to healthy tissue, is expected to be restrictive to diffusion. However, the expected changes to the ADC values are somewhat complicated and have not been fully characterised as reported separately by Koh and Vandecaveye ((Koh & Collins, 2007), (Vandecaveye et al., 2007)). Regions of tumour growth are likely to be disorganised and increased cellular density might be expected to result in a decrease in the observed isotropic diffusion coefficient (Vandecaveye et al., 2007). Since we are assessing areas of isotropy, values between the three orthogonal directions should be the same up to the noise value. Between different b -values, the isotropy measure should simply be related by scaling. Therefore, conventional cost-functions (at least those that accommodate overall intensity change (e.g. cross-correlation but not sum-of-squared differences) should remain appropriate. Assessing diffusion coefficients in the liver is made difficult by patient breathing motion, requiring either gated acquisition or post-processing techniques such as image registration (Kwee et al., 2008). This section applies image registration methods on order to observe changes to the calculation of pixel-wise ADC values. Figure 7.8 shows example liver images when varying the b -value: liver tumours appear as focal increases in the signal throughout the liver.

7.5.2 Method

Data assessed here consists of information from seven different b -values between $0-750s.mm^{-2}$, each non-zero value is assessed in three orthogonal directions generating a total of 16 images (volumes) (see Table 1.4). As reported by Koh (Koh & Collins, 2007), the spread of b -values is susceptible to flow artefacts at lower b -values ($b < 150s.mm^{-2}$) (particularly relevant for the highly vascular hepatic environment) so vascular regions (an extreme example is the aorta) might be assigned inappropriate ADC values when using a mono-exponential function such as in Equation 7.3.

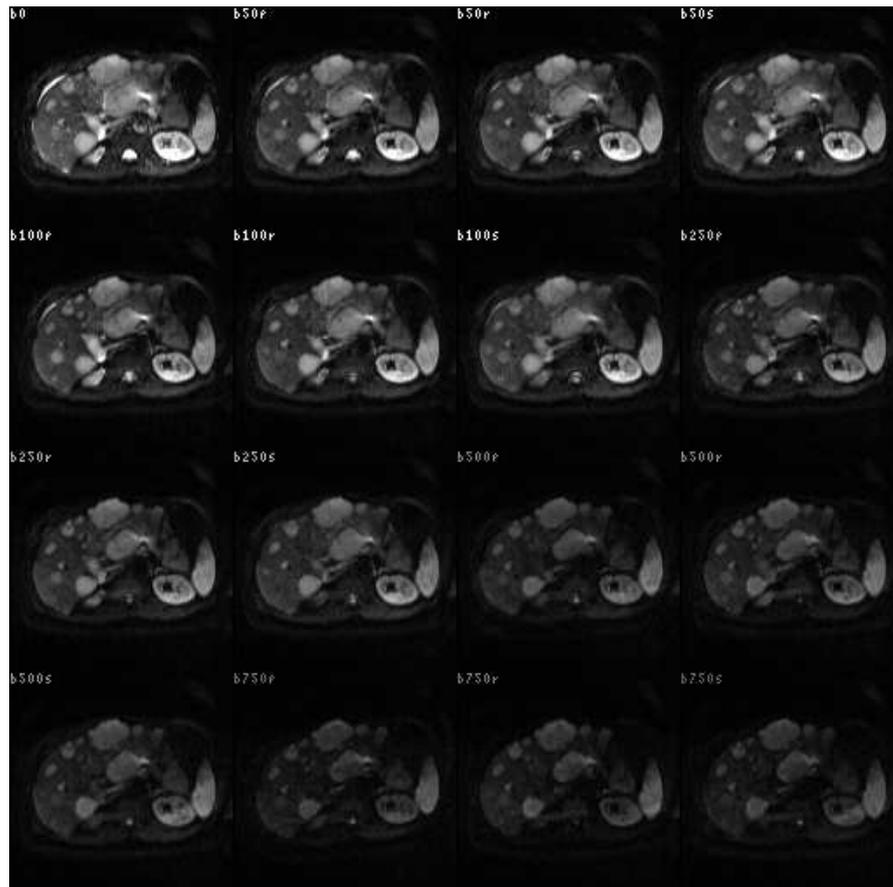


Figure 7.8: Example Abdominal DWI Slice, comprised of a B0 Image and fifteen diffusion images corresponding to different b values and orthogonal directions prs (labelled). The imaging signal decreases with increasing b -value (Equation 7.3), liver tumours appear as focal increases in the signal.

In contrast to the previously assessed DCE-MRI data and multiple diffusion direction cranial imaging, we are now assessing data where there should be no *local* contrast change. As discussed by Koh (Koh & Collins, 2007) we do not expect directional anisotropy in the ADC within areas of interest, although this might not always be the case, for instance near to blood vessels.

Calculation of the ADC, is found by non-linear least-squares fitting; finding the optimal value of D in the solution of 7.5. The estimation of D is dependent on the spread of b -values, noise and the stability of the signal to flow artefacts at low b -values. In this section we do a linear fit to the log of the signal intensities. Although this might introduce b -value dependent errors in the fitting, the bias towards higher values might be thought to counteract the lower confidence in low b -value data as a result of possible perfusion artefacts.

$$\min \left[\sum_n^{N_b} (S(b_n) - S(b_0)e^{-b_n D})^2 \right] \quad (7.5)$$

Table 7.1: Total ADC Residual Before and After Image Registration

Dataset	Number of images	Unregistered	Fluid Registration	PPCR
1	7	100%	94%	63%
2	16	100%	80%	71%

7.5.3 Results

Figure 7.9 shows calculated ADC maps before and after registration with corresponding fitting residual images in Figure 7.10 for the dataset with 7 b-values. Visible improvements are not immediately clear when assessing the the accuracy of the ADC values throughout the slices shown, however slices are shown for completeness. Inspection of the ADC fit residual reveals some improvement after registration. Figures in Table 7.1 are calculated over the central region of the dataset only (as in the previous section, defined by an ellipsoidal binary mask of radius $\frac{1}{4} \times \frac{1}{4} \times \frac{1}{2}$ of the image dimensions in $[x, y, z]$ on the centre of the image). A reduction in residual is seen using the PPCR method in the residual maps in the right hand column of Figure 7.10. Corresponding figures for the total residual are presented in Table 7.1 showing a fall in ADC residual over the entire volume of 29% for this case. Direct registration sees a reduction in residual of 20% in this case, as may be expected following the discussion above in which we would expect image registration using unmodified cost-functions to be successful. A reduction in ADC residual is also seen for dataset 1 in Figure 7.1 although we are fitting to fewer b -values.

7.5.4 Conclusion

This section presents preliminary work on the registration of diffusion MRI outside of the brain. The increased use of diffusion weighted imaging in oncology will require sophisticated image processing techniques to reduce the influence of patient motion. In particular for the abdomen, the use of non-rigid registration algorithms is likely to be necessary. This section provides evidence that standard non-rigid registration techniques should be suitable, however one explanation for improved results using PPCR is that the iterative nature of PPCR and registration to a common coordinate frame rather than that of a single particular image can be used to provide a better result (Bai & Alexander, 2008). Perhaps the implementation of PPCR in the registration of any large group of images (given the conditions imposed in Chapter 5) should be considered when using images in other applications.

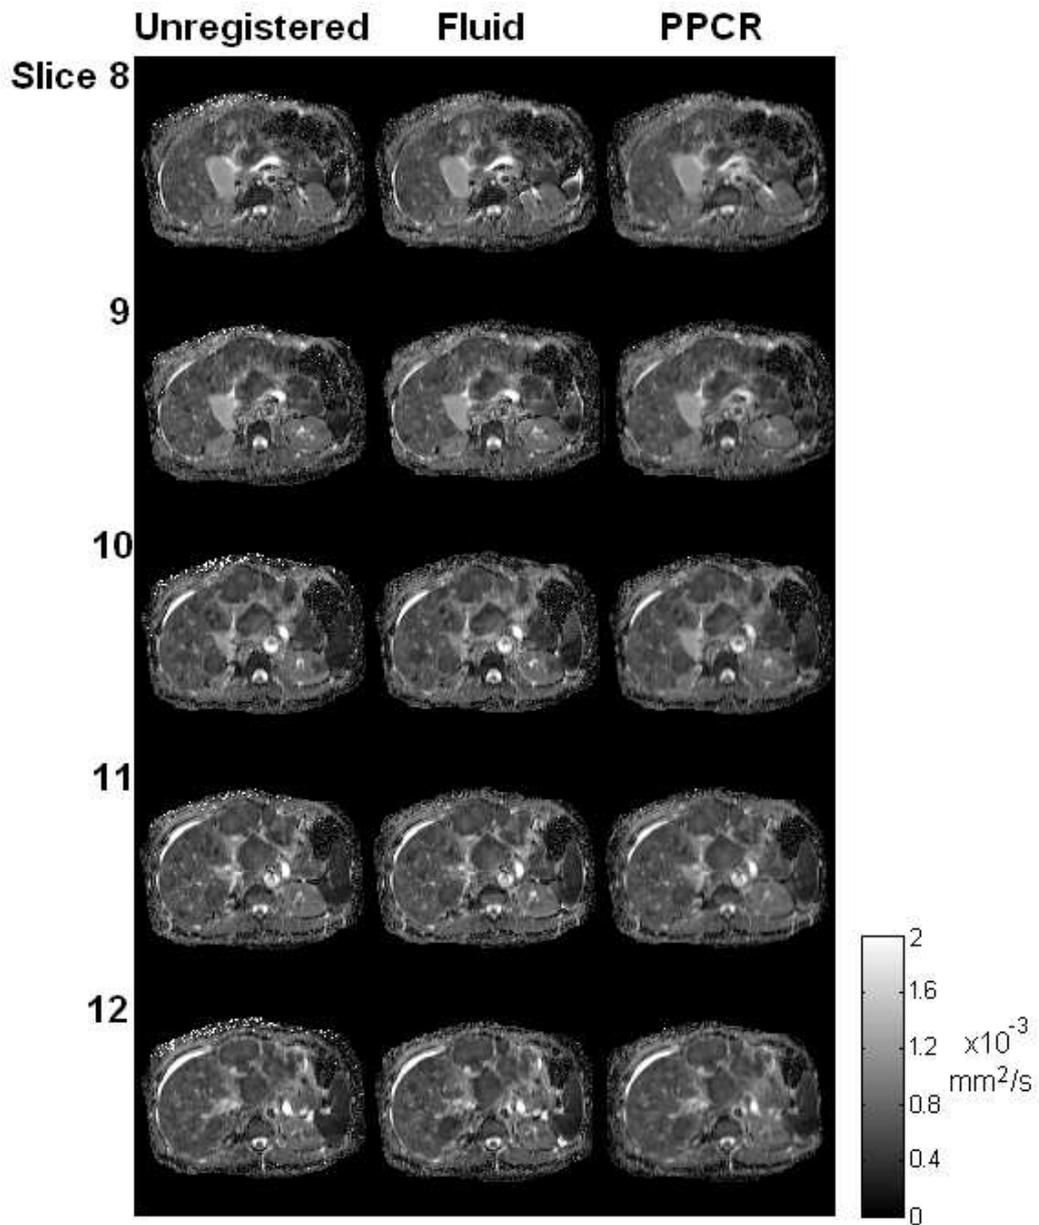


Figure 7.9: ADC maps for slices from Patient 1 from Table 1.4: **Column 1** before registration, **Column 2** after fluid registration using NMI cost function, **Column 3** after registration by PPCR.

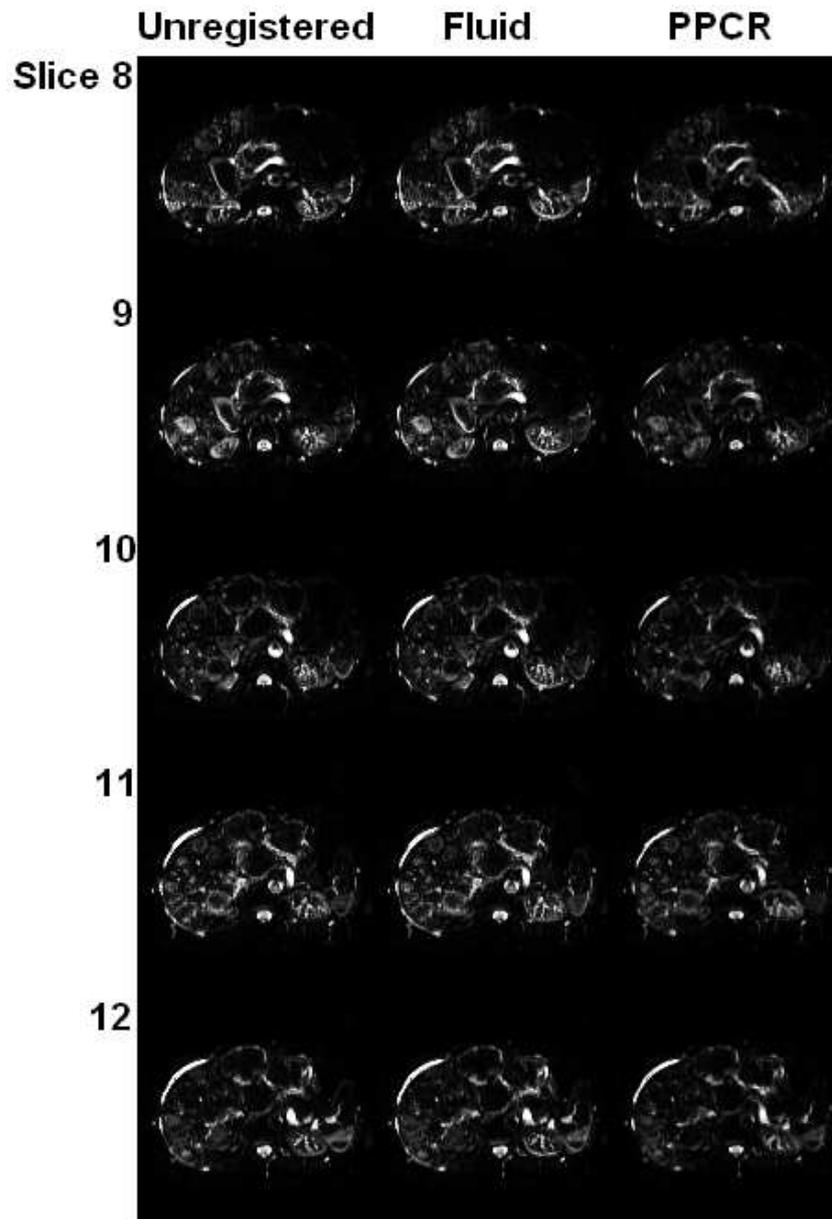


Figure 7.10: ADC residual maps for slices from Patient 1 from Table 1.4: **Column 1** before registration, **Column 2** after fluid registration using NMI cost function, **Column 3** after registration by PPCR. Showing some evidence of reduction to model-fit residuals after registration.

Chapter 8

Kullbach Leibler Assisted Image Matching and Patching (KLAMP)

8.1 Introduction

A Dynamic Contrast Enhanced MRI dataset consists of many imaging frames, often both before and after contrast injection. Registration may be driven by minimising joint image histogram dispersion. Force gradients used to drive registration algorithms are derived from changes that reduce the dispersion in the joint image histogram. Differences between joint image histograms between images before and after contrast arrival can be compared. We investigate if pixels contributing to joint image histogram dispersion by contrast change alone might be separated from those due to motion changes. As discussed in Chapter 5 the joint histogram between pre and post enhancement images contains dispersion due to both motion and enhancement, both of which are minimised during registration. If we assume that there are gross changes to the joint image histogram between pre and post enhancement due to enhancement processes, then we can seek to minimise those changes by comparison of the histogram distributions. Work related to intensity alteration has been produced by Weisenfeld (Weisenfeld & Warfield, 2004): the authors sought a functional multiplicative relationship between intensities in two images in order to correct for intensity biases in the MRI acquisition and therefore improve segmentation techniques. The concept is similar to that proposed in this chapter, but here we seek to reduce the influence of large local contrast variations on image registration.

Figure 8.1 shows the generation of joint image histograms: one between two pre-enhancement images and the second between a pre-enhancement and post-enhancement image. Image registration force gradients are calculated to reduce dispersion in the joint image histograms, however in the case of the pre and post-enhancement images, dispersion in the joint image histogram is generated from both motion and contrast-enhancement and it is possible that image registration will attempt to shrink or distort enhancing regions as discussed in Chapter 4.

8.2 Method

We describe the joint image histogram (normalised so that it may be considered a probability distribution of intensity values) of two pre-enhancement images, \mathbf{A}_{pre1} and \mathbf{A}_{pre2} as $\mathbf{JIH}(\mathbf{A}_{pre1}, \mathbf{A}_{pre2})$. Similarly,

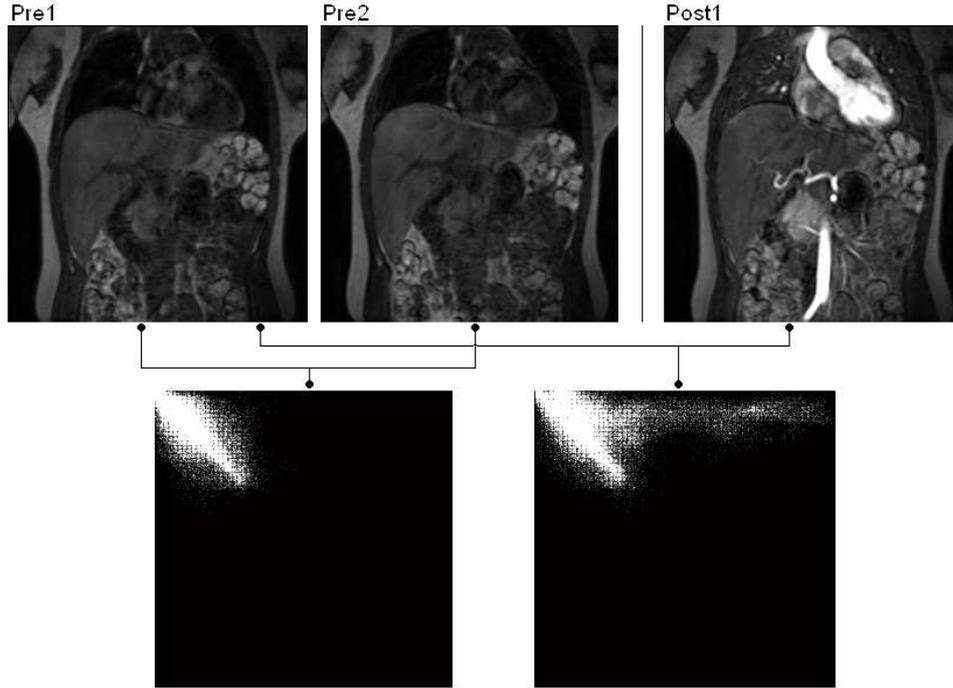


Figure 8.1: Formation of joint image histograms between a pre enhancement images and either a pre or post enhancement image. Force-gradients that reduce the dispersion in the histogram formed using the post enhancement image may result in shrinkage of enhancing features.

we consider the normalised joint image histogram between a pre-enhancement image, \mathbf{A}_{pre1} , and a post enhancement image, \mathbf{A}_{post1} , multiplied by a binary pixel mask, \mathbf{M} giving $\mathbf{JIH}(\mathbf{A}_{pre1}, \mathbf{M} \cdot \mathbf{A}_{post1})$. We optimise this mask so that the divergence between the pre and post joint image histograms, as measured by the Kullbach Leibler divergence (KLD - Equation 8.3), is minimised (Equation 8.2). For an enhancing image, masking pixels that are enhancing reduces their impact on the formation of image force gradients as discussed in Chapter Four. In this sense our pre-enhancement histogram is a *training histogram* used to modify the *enhancement histogram* so that it can be used for motion correction. Our pre-enhancement histogram should contain dispersion typical of the motion between two images and should therefore be calculated between pre-enhancement images only. Any dispersion in this joint image histogram can be considered to be motion related. By comparing the pre-pre and pre-post joint image histograms we can estimate which contributions to the pre-post joint image histogram are due to enhancement and suppress their influence in subsequent image registration procedures.

$$KLD(\mathbf{JIH}(\mathbf{A}_{pre1}, \mathbf{A}_{pre2}) \parallel \mathbf{JIH}(\mathbf{A}_{pre1}, \mathbf{M} \cdot \mathbf{A}_{post1})) \quad (8.1)$$

$$\min[KLD(\mathbf{JIH}(\mathbf{A}_{pre1}, \mathbf{A}_{pre2}) \parallel \mathbf{JIH}(\mathbf{A}_{pre1}, \mathbf{M} \cdot \mathbf{A}_{post1}))] \quad (8.2)$$

The mask can be found using an automatic method inspecting the small change in Kullbach Leibler divergence (KLD) brought about by removing individual pixels; the approach used here is similar to that used by Crum *et al* (Crum et al., 2005). If the removal of a particular pixel reduces the KLD, then the

corresponding mask position is set to zero. The small change in KLD can be found by considering the removal of a particular pixel from the histogram, hence the associated bin contents are reduced by 1. This procedure is non-iterative and fast and the derivation is given in Equations 8.3 to 8.8 for corresponding intensity bin contents in the training histogram n_i and enhancement histogram m_i .

$$KLD = \frac{1}{N} \sum_{i=1}^{bins} n_i \log\left(\frac{n_i}{m_i}\right) \quad (8.3)$$

$$KLD = KLD_{i \neq j} + \frac{1}{N} [n_j \log(n_j) - n_j \log(m_j)] \quad (8.4)$$

$$KLD = KLD_{i \neq j} + \frac{1}{N} [(n_j - 1) \log(n_j - 1) - (n_j - 1) \log(m_j - 1)] \quad (8.5)$$

$$KLD = KLD_{i \neq j} + KLD_j + dKLD \quad (8.6)$$

$$dKLD = \frac{1}{N} [n_j \log\left(\frac{m_j(n_j - 1)}{n_j(m_j - 1)}\right) - \log\left(\frac{n_j}{m_j}\right)] \quad (8.7)$$

$$dKLD \approx \frac{1}{N} \log\left(\frac{m_j}{n_j}\right) \quad (8.8)$$

From the steps outlined above, we expand Equation 8.5 to Equation 8.7 and find the change in KLD ($dKLD$). For a large number of pixels N and large bin contents n_i and m_i (therefore $n_i \approx n_i - 1$), the expression for $dKLD$ can be represented as Equation 8.8.

It may not be appropriate to simply remove those pixels that result in a reduction in the KLD. By inspection of test data, this process also removes pixels that are not enhancing. Theoretically this is due to discrepancies between the training histogram and enhancement histogram that encompass motion not captured by the training histogram. It is therefore necessary to define a threshold $dKLD$ so that if removing a particular pixel reduces the total KLD by more than the threshold amount, we mask that pixel. The assumption is that contrast enhancement intensity changes contribute more to the KLD value than motion artefact intensity changes; it follows that contrast enhancement intensity changes that contribute to a large KLD are generating dispersion in the enhancing joint image histogram that is greater than that due to motion artefacts. The reduction of contrast-enhancement induced joint image histogram dispersion is likely to result in mis-registration, so its removal may be advantageous. Using this method, the pixel mask should remove the influence of contrast-enhancement on the force gradient by removing from the analysis those pixels that are changing intensity value between pre and post enhancement images in a fashion unlike those between the pre-enhancement images.

If registration was carried out between the pre enhancement image and the masked post enhancement image, we would still have spurious force gradients at the mask boundaries. To reduce or remove this effect it is necessary to calculate a patch image to fill the gap produced by the mask. One method of doing so is as follows: for each masked pixel, we look at the image intensity it should have, given the joint image histogram formed by the masked post-enhancement image and the pre-enhancement image (see Figure 8.2). The most likely value (that with the highest probability) is then given to each masked pixel to create a patch. Selecting the most likely value is preferable to drawing from the associated intensity probability distribution since this would introduce histogram dispersion that may produce spurious force-gradients; taking the most likely value prevents this. The effect of the patch is to reduce or remove the effect of erroneous force gradients at the mask edges. This arises because the masked pixel has been

given the most likely intensity value and so generating a force across it is unlikely to further minimise (or maximise) the cost-function.

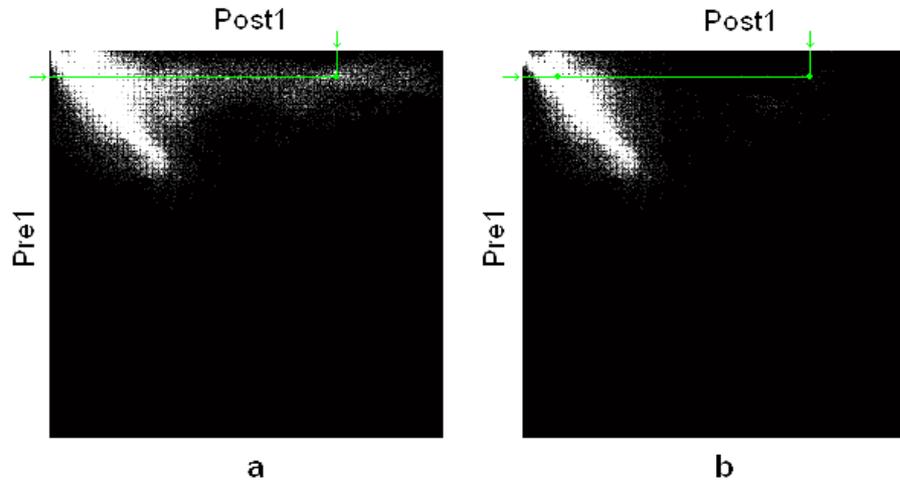


Figure 8.2: Selection of most-likely pixel intensity values for the masked image patch. Masked pixels have a new value calculated by considering their most likely value from the new joint image histogram between pre and masked post-enhancement images (as demonstrated by the green line markings, arrows show the original intensity values, the green dot the replaced intensity value). **a** Original joint image histogram and **b** Joint image histogram after masking.

The process described above is termed **K**ullback **L**eibler **A**ssisted image **M**atching and **P**atching (KLAMP).

Results are shown for two different variations of the KLAMP method. First we inspect the effect on final image registration result using the non-rigid fluid registration method described in Appendix C where we submit an anchor image and a masked and patched float image. Second we modify the fluid registration algorithm so that the KLAMP method is directly incorporated (see Appendix D). To analyse registration performance, we use a real pre-enhancement anchor image and a real post-enhancement image subject to an additional deforming force (as per Chapter 3). The additional perturbing force generates motion artefacts larger than the existing motion between the real pre and post enhancement images. Therefore we have a reasonable approximation to the gold standard deformation that the registration algorithms should recover. Analysing manually segmented images of the liver, heart and aorta, we can investigate how well we recover the gold-standard images and quantify mis-registration.

8.3 Results

Results are presented for direct analysis of the force gradients, selection of the pixel rejection threshold and analysis of the deformation fields using a pre-enhancement anchor image and masked and patched post-enhancement float image. We also investigate the effect of the registration after incorporation of the KLAMP algorithm into a fluid registration implementation as in Appendix D.

8.3.1 Inspection of Driving Force Gradients using KLAMP

Figure 8.3 demonstrates the formation of an image mask and its effect on reducing contrast-enhancement induced dispersion on the joint image histogram between pre and post-enhancement images.

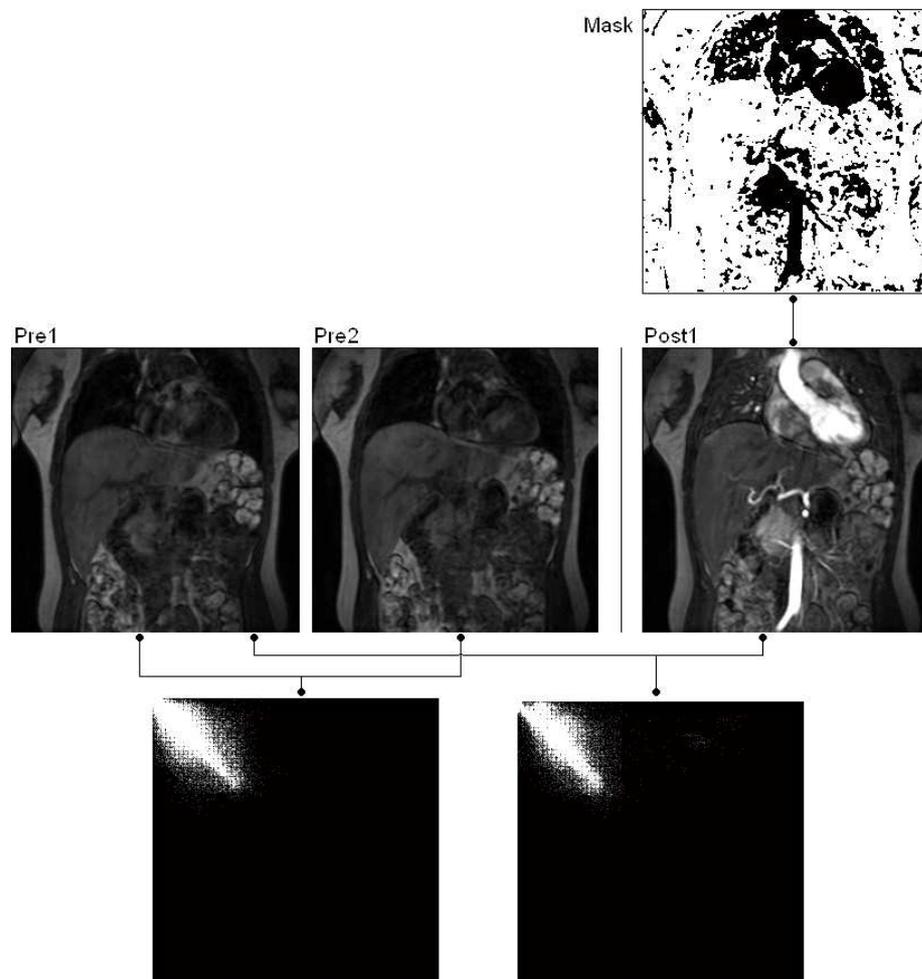


Figure 8.3: Formation of joint image histograms between two pre enhancement images and a masked post enhancement image. Formation of the image mask as discussed in the text reduces force-gradients that are likely to contribute to shrinkage of enhancing features.

Figure 8.4 demonstrates the described masking and patching method and the effect on the image registration force gradients. Force gradients are calculated by consideration of maximising the Normalised Mutual Information as shown in Appendix B. Force gradients for the registration of the unmodified post-enhancement image to the pre-enhancement image are wrong (Figure 8.4): in addition to medial-lateral

correction of abdominal wall movement and superior-inferior liver/diaphragm displacement there are pinching force gradients on both the aorta and the left-ventricle/ascending aorta. The formation of NMI force gradients when using the masked post-enhancement image and the pre-enhancement image are different (Figure 8.5), but as discussed in the methods section, there are difficulties in calculating force gradients over the boundaries between masked and unmasked pixels. These effects are reduced or removed using masking and patching (Figure 8.6). In the masked and patched method we maintain the force gradients correcting medial-lateral abdominal wall movement and superior-inferior liver/diaphragm displacement but have removed force gradients associated with the pinching of enhancing features. There remains some residual difficulty in the left-ventricle which may need to be considered on the eventual image registration. The application of regularisation (e.g. finding the displacement from these force gradients using the fluid-equation) will smooth 'noisy' force gradient regions but is likely to preserve the consistent forces that correct the major abdominal wall movement and superior liver displacement.

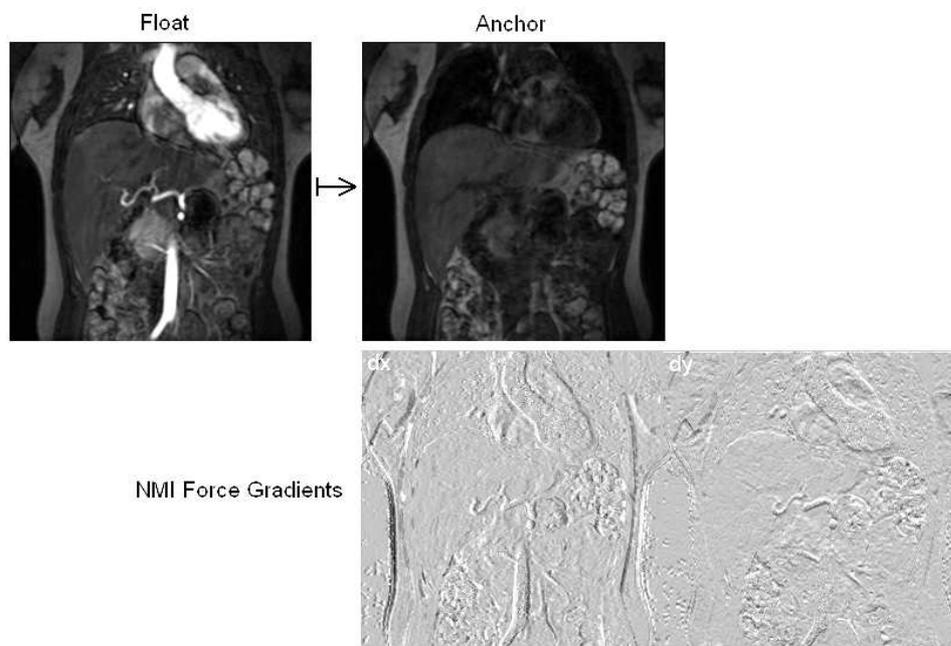


Figure 8.4: Formation of Normalised Mutual Information force gradients between unaltered post-enhancement and pre-enhancement images. Note the pinching gradients around the heart and aorta in addition to abdominal wall and diaphragm position corrections.

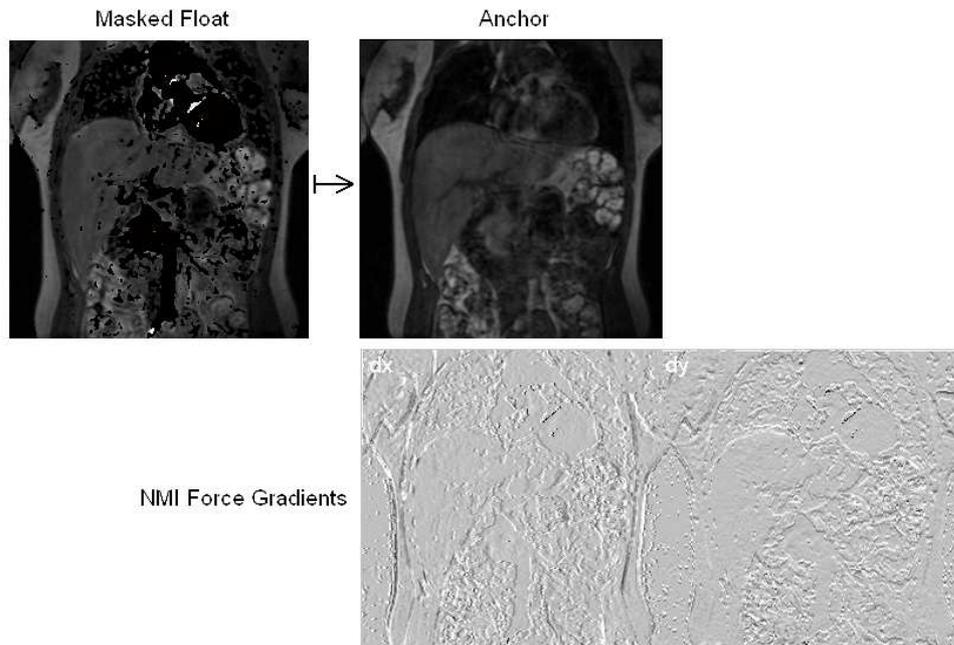


Figure 8.5: Formation of Normalised Mutual Information force gradients between masked post enhancement image and pre-enhancement image. Note the absence of pinching force gradients in the heart and aorta but appearance of flat regions with visible gradients at the boundaries. The mask used is identical to that shown in Figure 8.3.

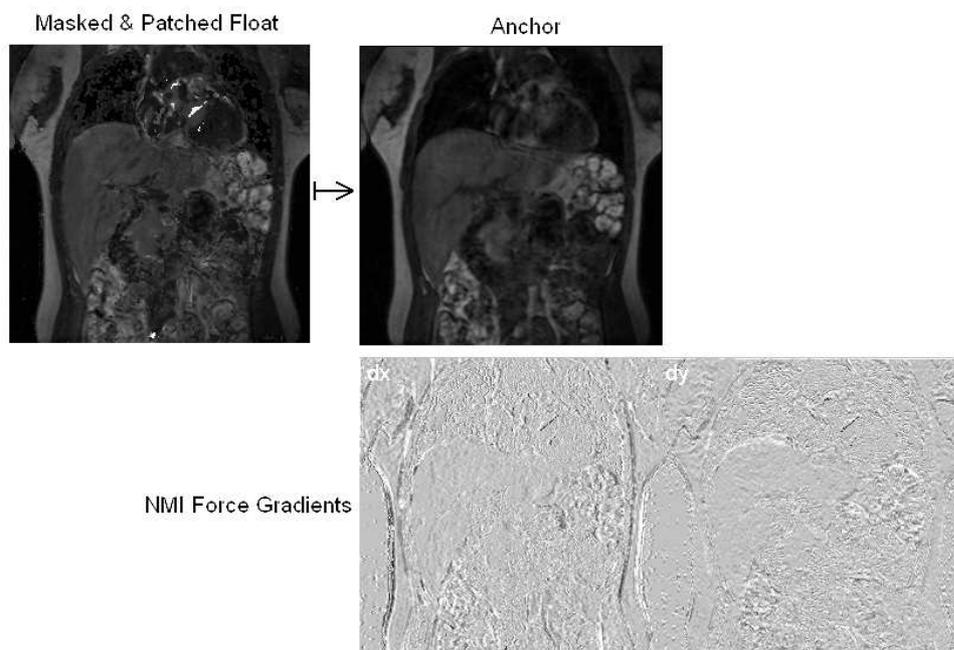


Figure 8.6: Formation of Normalised Mutual Information force gradients between masked and patched post-enhancement image and pre-enhancement image. Note the absence of pinching force gradients and its replacement with local noise in masked regions. Abdominal wall and liver position corrections are maintained. The mask used is identical to that shown in Figure 8.3.

8.3.2 Setting the KLD Threshold in KLAMP

The selection of the threshold in the previous section was an empirical choice having a value of $\log \frac{n_i}{m_i} = -1$ for the bin contents n and m at intensity i in the training and enhancement histograms respectively. This threshold value is now tested, and varied between 0 and -2 in steps of 0.25. Figure 8.7 shows the resulting image masks (morphologically dilated by one pixel). Selection of a threshold value of -1 appears to provide a trade off between masking enhancing features such as the aorta and masking features due to unmodelled histogram differences. This should be tested further before stronger conclusions can be formed.

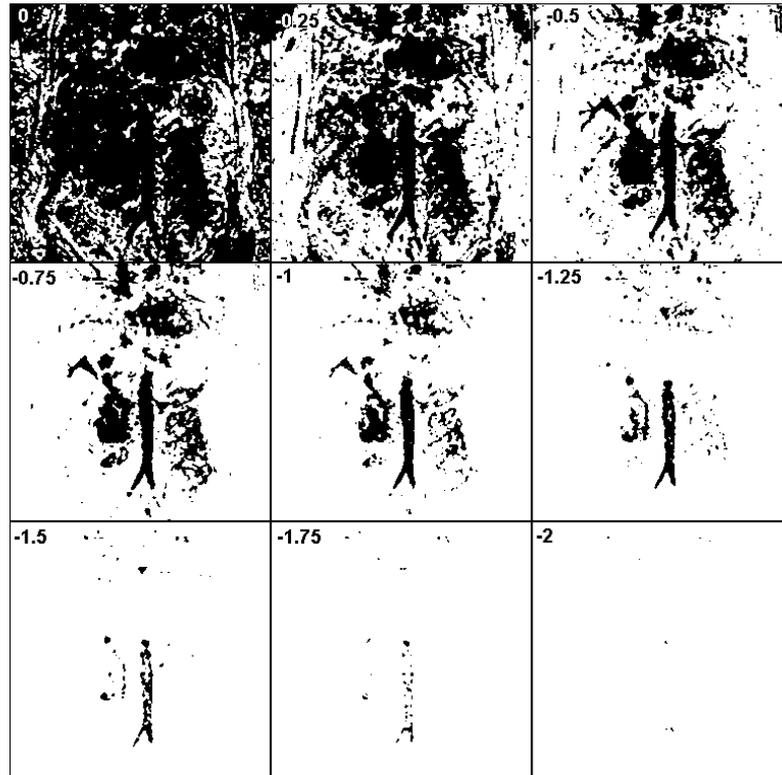
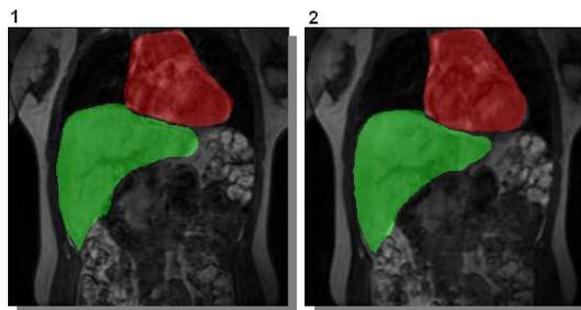


Figure 8.7: KLAMP mask formation for threshold values between 0 and -2 using step size 0.25 (Pixels are masked if $\log \frac{m_i}{n_i} < threshold$).

8.3.3 Float Image Pre-Processing using KLAMP

Hand-segmented contrast enhanced images are selected from the data in Table 1.2 as demonstrated in Figure 8.8. Three segmentation types are selected representing the heart, the enhancing aorta and the liver. Since the images are in reasonable alignment prior to registration, we impart an additional deformation using the motion model in Chapter 3. The deformation is applied to both the enhancing images and the training image (image 2) so that a good training histogram is found. We now have a gold-standard against which to compare the results of image registration. Mis-registration of enhancing features may result in a reduction in the area of segmented enhancing features (volume in 3D) between gold-standard and registered images. Correct registration should result in an increased intersection of the segmented liver between gold-standard and post-registration contrast-enhanced images. Image registration is carried out between post-enhancement and pre-enhancement images with and without KLAMP to produce deformation fields which may then be applied to the segmentations in order for a comparison with the 'undeformed' gold-standard.

Pre-Enhancement Images



Enhancement Images

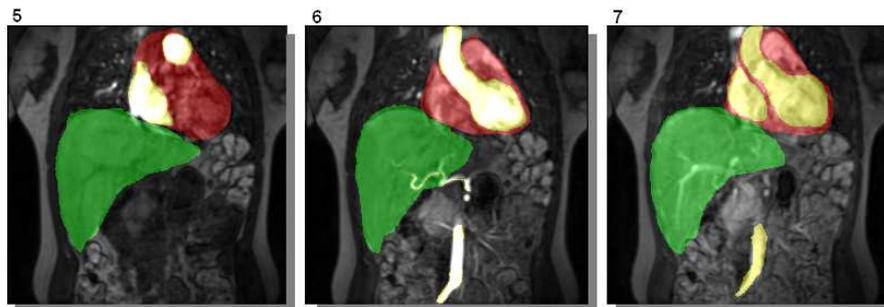


Figure 8.8: Segmentation of enhancing features and the liver for analysis of registration performance. Image 1 represents the anchor image to which we register enhancing float images 5 to 7. Image 2 is the additional training image used in the KLAMP method. Motion artefacts between images 1 and 2 are used to generate the training joint image histogram.

The following section will provide an analysis of the KLAMP method as applied within the fluid registration framework of Appendix C. Table 8.1 shows results for the total area of heart and aorta and the intersection of liver pixels before deformation and after deformation and subsequent registration

with and without the KLAMP method. There is evidence that the unmodified registration method causes shrinking of the enhancing heart and aorta by up to 18%. This incorrect shrinking of the heart and aorta is visible in both the registered and segmentation images and may be counteracted using the KLAMP method. The final row in the Table 8.1 suggesting that the shrinking is the same both with and without KLAMP is the result of poor threshold choice in the KLAMP image mask (the threshold is too high and no pixels are masked in this case). Registration success, as measured by an increased overlap of segmented liver pixels is represented in Table 8.1 columns 5 and 6, as a percentage of the maximum possible number of overlapping green pixels. The number of overlapping pixels between the original undeformed image and the registered image is improved in all cases both with and most cases without KLAMP. However, the KLAMP method appears to cause under-registration of the images, as marked by a much smaller increase in overlapping pixels. The reasons for this are investigated further by inspection of the modification to the potential cost-function minimisation space caused by the KLAMP algorithm.

Further results are shown in Table 8.1 Columns 7 and 8 for the total absolute difference between the applied deformation and correcting registration deformation field. Results are expressed as a percentage of the applied deformation; perfect deformation recovery would result in a zero absolute difference. The under registration using the KLAMP method is summarised by residual deformations that represent a large fraction of the original applied deformation. Recovery of the deformation without the KLAMP method is variable.

Some example images are shown in Figure 8.9 for the marked result in Table 8.1. Given the pre-enhancement anchor image, we register a post-enhancement float image to the anchor using normalised mutual information. The results for each registration method are shown in images **b,c,d** in Figure 8.9. There is visible pinching of the enhancing left ventricle in the registration without KLAMP. Corresponding segmentation images are shown in the second row and the recovery of the gold-standard deformation is shown in the third row in the colour difference image. Clearly visible in image **i** is the contraction of the ventricle (and also the descending aorta). Results for the corresponding segmentation image (image **j**) after registration using KLAMP show the under-registration suggested by the results of Table 8.2. The KLAMP method implemented as a float image pre-processing step appears to make registration difficult.

Table 8.1: Segmentation statistics for the change in area of heart and aorta and the intersection of the liver after image registration with and without KLAMP. Columns 7 and 8: Absolute Residual Deformation (as percentage of original deformation) after image registration with and without KLAMP. **see Figure 8.9.*

Dataset :	Area of Segmented Heart and Aorta (Without KLAMP)	Area of Segmented Heart and Aorta (With KLAMP)	Intersection of Liver (Deformed)	Intersection of Liver (Without KLAMP)	Intersection of Liver (With KLAMP)	Residual Deformation (Without KLAMP)	Residual Deformation (With KLAMP)
1 : 5 \mapsto 1	-13%	-1%	90%	99%	93%	33%	56%
1 : 6 \mapsto 1	-16%	-1%	91%	99%	92%	32%	82%
1 : 7 \mapsto 1	-7%	-4%	86%	97%	89%	33%	83%
2 : 5 \mapsto 1	-17%	-2%	87%	99%	88%	27%	67%
2 : 6 \mapsto 1	-10%	-2%	83%	98%	92%	31%	51%
2 : 7 \mapsto 1	-7%	-8%	88%	96%	90%	78%	98%
3 : 5 \mapsto 1	-11%	-0%	87%	95%	85%	50%	100%
3 : 6 \mapsto 1*	-14%	-1%	88%	98%	87%	48%	90%
3 : 7 \mapsto 1	-10%	-13%	88%	95%	88%	70%	98%

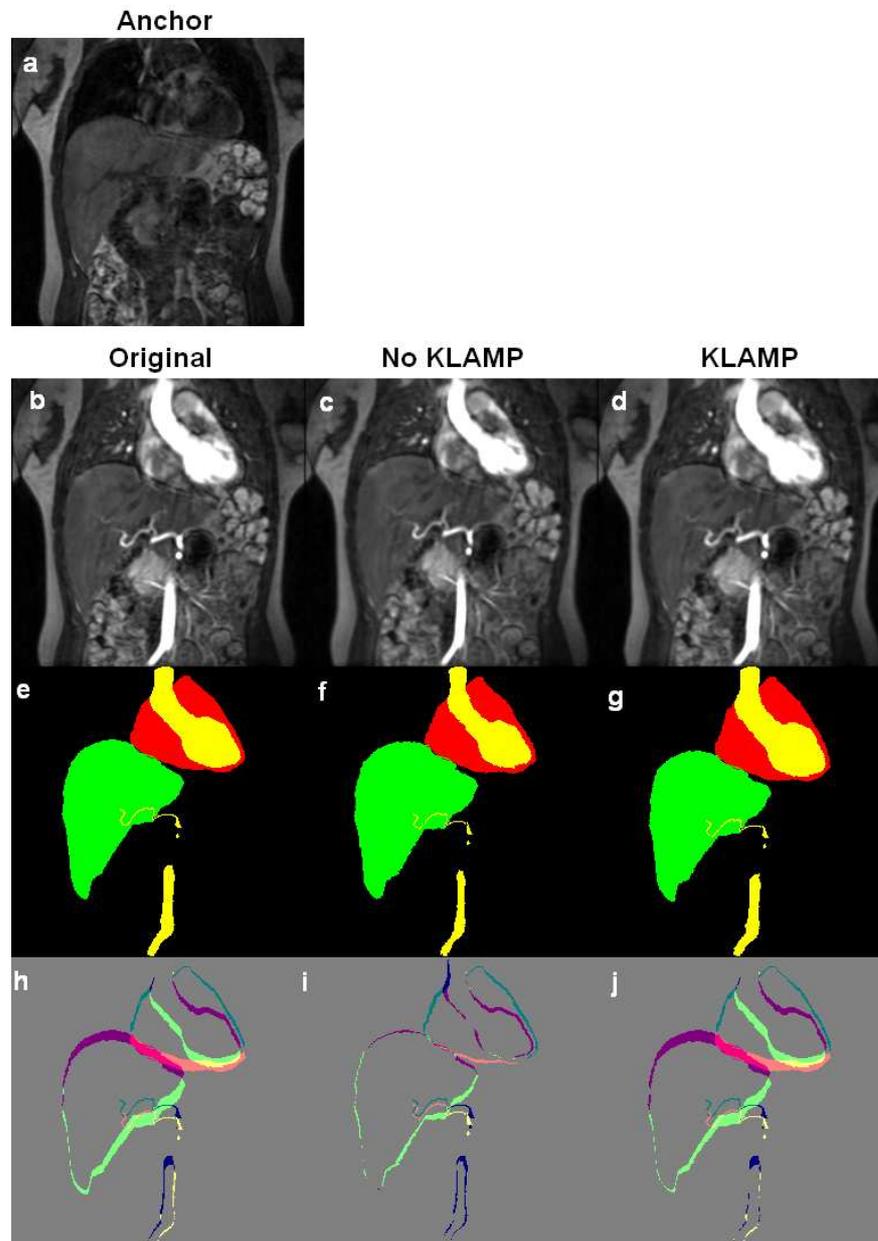


Figure 8.9: Example images from Table 8.2 Row 8. **a)** anchor image. **b)** original un-deformed image to compare with **c)** registration without KLAMP and **d)** registration with KLAMP. **Row 2:** Corresponding segmentation images after registration. **Row 3:** Segmentation difference images with gold-standard segmentation for each registration method.

Inspection of the effect of the KLAMP method on cost-function minimisation may reveal the reason for the apparent under-registration. Using the cost-function minimisation space analysis of Section 4.3 it is possible to inspect how easily we might expect to minimise our cost-function after KLAMP. Figure 8.10 demonstrates the changes to the minimisation space and the effect of the suppression of enhancing features. It appears that the inclusion of the KLAMP algorithm might make the minimisation of information theoretic cost-functions difficult. The reason for this is the patching of the image with the most likely intensity values. This is designed to reduce spurious force gradients in the masked image, but has the effect of increasing the value of the normalised mutual information, reducing the scope for cost-function improvement by registration and flattening the minimisation space.

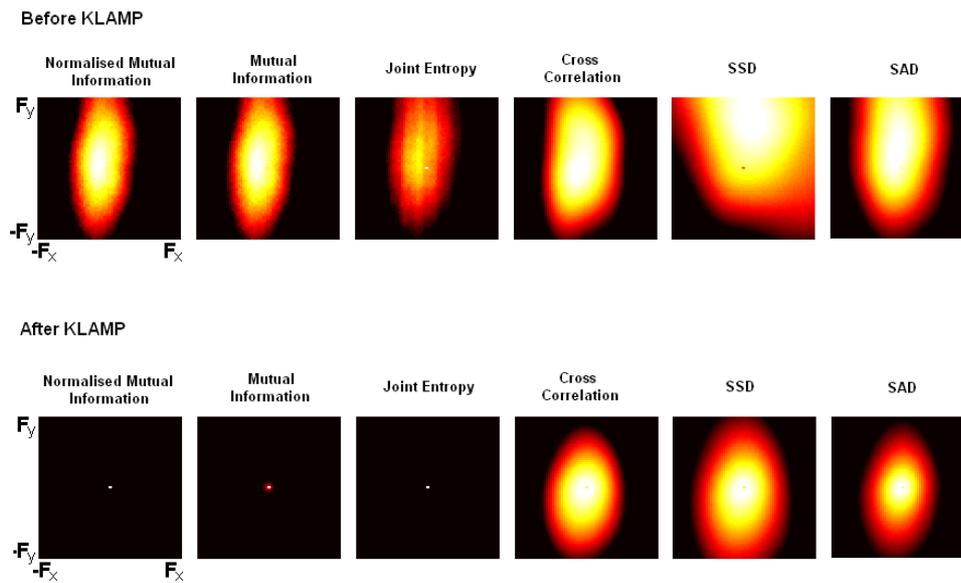


Figure 8.10: Cost-function minimisation spaces of a pre and post enhancement image. Top row: Unmodified images deformed by a force varying in x and y . Bottom row: Corresponding minimisation spaces for pre-enhancement image and KLAMPed post-enhancement image demonstrating a flattening of the cost-function space.

8.3.4 Joint Entropy Recalculation Using KLAMP

We now analyse the effect of incorporating the KLAMP algorithm into the image registration procedure. The previous sections necessitate a modification to the original KLAMP algorithm in order to maintain a cost-function that may be minimised. The use of the KLAMP method to produce a float image suitable for registration appears to make the cost-function difficult to minimise successfully. In order to maintain the minimisation of the cost-function we make the KLAMP method internal to the registration algorithm and update the mask and patch at each iteration. The resulting fluid-KLAMP algorithm and the modification to the cost-function formation is described in Appendix D.

From Chapter 2, when inspecting the entropies calculated when using mutual information (or normalised mutual information), we seek to maintain the marginal entropies H_A and H_F whilst reducing the joint entropy H_{AF} . Of these three terms, it is the joint entropy that is difficult to minimise successfully due to the formation of the joint image histogram that includes dispersion due to both motion artefacts and contrast enhancement. The marginal entropy H_F in the presence of contrast-enhancement is likely to be increased when using contrast-enhancement. Therefore it may be considered important not to influence the marginal probability distribution using the KLAMP algorithm. It is possible to use the KLAMP algorithm to modify the joint image histogram formation only, leaving the marginal float image histogram unchanged. This will alter the interpretation of the mutual information slightly since we are now considering a measure of mutual information given three images. In an information theoretic sense we are assessing the interaction information of the anchor, float and masked-float images, however a formal definition is not presented in this thesis and is left as future work. This re-definition of the mutual information should allow improved image similarity maximisation; we maintain an unmodified marginal float entropy and also minimise the effects of the contrast-enhancement dispersion formed in the joint-image histogram.

Table 8.2 columns 2 and 3 show the reduction in volume of the heart and aorta with and without the internal KLAMP algorithm. The volume reduction associated with enhancing features is minimised using the internal KLAMP algorithm. However there remains some residual volume reduction. The segmented liver intersection results in Table 8.2 columns 4 to 6 show improved results when using the internal-KLAMP method as compared to Table 8.1. Registration performance in many cases is comparable to results not including the internal-KLAMP algorithm. Results for the absolute deformation recovery residual are shown in Table 8.2 Columns 7 and 8 confirming an improvement to the re-alignment of the deformed images both with and without internal-KLAMP although the registration fails to register in two cases. This under-registration is the reason for the lowest results for the change in segmented area of heart and aorta also seen in Table 8.1.

Some example images are shown in Figure 8.11 for the marked result in Table 8.2. Given the pre-enhancement anchor image, we register a post-enhancement float image to the anchor using normalised mutual information. The results for each registration method are shown in images **b,c,d** in Figure 8.11. There is visible pinching of the enhancing left ventricle in the registration without KLAMP. Corresponding segmentation images are shown in the second row and the recovery of the gold-standard deformation

is shown in the third row in the colour difference image. Clearly visible in image **i** is the contraction of the ventricle (and also the descending aorta). These features of mis-registration are not present in the corresponding segmentation after registration using KLAMP (image **j**), although there appears to be slight under-registration of the liver position and small over correction of the position of the aorta. These results correspond well to Table 8.2.

Table 8.2: Segmentation statistics for the change in area of heart and aorta and the intersection of the liver after image registration with and without internal-KLAMP. Columns 7 and 8: Absolute Residual Deformation (as percentage of original deformation) after image registration with and without internal-KLAMP. *see Figure 8.11

Dataset :	Area of Segmented Heart and Aorta (Without KLAMP)	Area of Segmented Heart and Aorta (With KLAMP)	Intersection of Liver (Deformed)	Intersection of Liver (Without KLAMP)	Intersection of Liver (With KLAMP)	Residual Deformation (Without KLAMP)	Residual Deformation (With KLAMP)
1 : 5 \mapsto 1	-13%	-10%	90%	99%	99%	33%	33%
1 : 6 \mapsto 1	-16%	-12%	91%	99%	98%	32%	32%
1 : 7 \mapsto 1	-7%	-5%	86%	97%	96%	33%	36%
2 : 5 \mapsto 1	-17%	-0%	87%	99%	87%	27%	97%
2 : 6 \mapsto 1	-10%	-6%	83%	98%	96%	31%	35%
2 : 7 \mapsto 1	-7%	-6%	88%	96%	96%	78%	71%
3 : 5 \mapsto 1	-11%	-7%	87%	95%	94%	50%	53%
3 : 6 \mapsto 1*	-14%	-5%	88%	98%	95%	48%	53%
3 : 7 \mapsto 1	-10%	-1%	88%	95%	88%	70%	93%

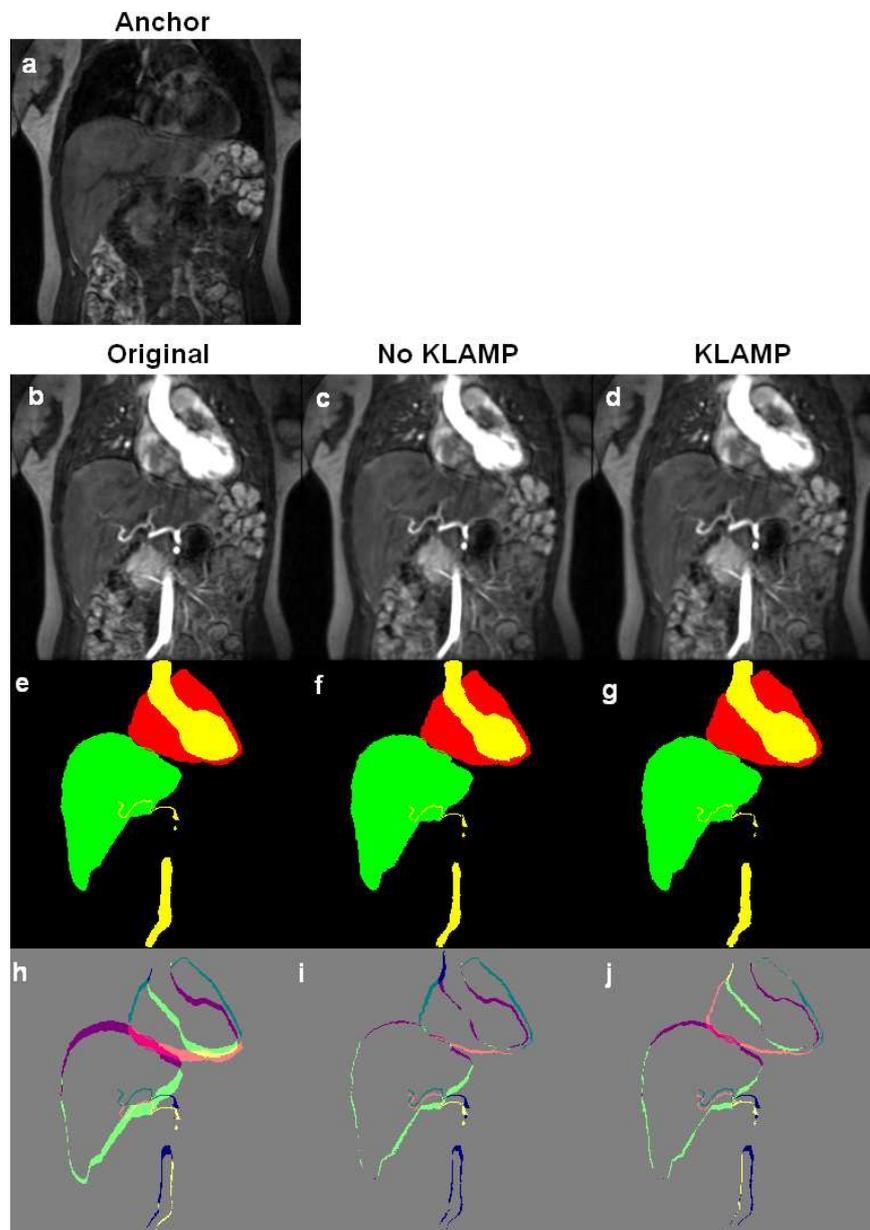


Figure 8.11: Example images from Table 8.2 Row 8. **a)** anchor image. **b)** original un-deformed image to compare with **c)** registration without KLAMP and **d)** registration with KLAMP. **Row 2:** Corresponding segmentation images after registration. **Row 3:** Segmentation difference images with gold-standard segmentation for each registration method.

8.4 Conclusion

The development of the Kullback Leibler Assisted image Matching and Patching (KLAMP) method has suggested that correct registration of images containing enhancing features may be achieved by a modification to the underlying cost-function intensity statistics. The method is suitable for low numbers of images, perhaps those containing first pass bolus information as shown in the examples in this chapter. The method requires at least two pre-enhancement images so that a training joint image histogram free of contrast-enhancement related dispersion can be constructed. The formation of the training joint image histogram is not necessarily limited to only two pre-enhancement images. It is conceivable that a wider range of motion artefacts could be allowed when using larger numbers of training images. Setting the threshold mask value is currently a largely empirical exercise chosen to eliminate contrast-enhancing features whilst maintaining artefacts due to motion discrepancies between the training and enhancement joint image histograms. Further investigation might be needed in order to define a theoretical cut off, given a set of anchor, float and training images.

The KLAMP method was originally devised as a method to allow the registration of a pre-processed post-enhancement masked and patched float image to a pre-enhancement anchor image without modification to the following registration algorithm. However, the formulation of the patched image appears to make the cost-function difficult to minimise. This is the result of masked pixels being given their most likely intensity value. Unfortunately this is necessary since we do not want spurious force gradients to remain in the image registration. Modifying the image registration algorithm so that the KLAMP method is incorporated within it allows the direct modification of the fluid equation driving force gradients. The following solution to the fluid equation (the regularisation) produces an acceptably smooth deformation across the force gradient mask boundaries.

The results in this chapter, although fairly preliminary, suggest that the method might be able to allow improvements to registration accuracy in the presence of contrast-enhancement. Unfortunately the results are not conclusive enough to suggest that the KLAMP algorithm in its current form is either robust or suitable for a wide range of enhancement characteristics. Further work is required to analyse the best implementation of the algorithm but results have shown that it is possible to achieve better results if the KLAMP algorithm is incorporated as a modification within an existing registration algorithm. Implementation of the KLAMP algorithm into the registration algorithm allows adaptive modifications to be made to the method. For instance, it would be possible to allow adaptive setting of the KLAMP threshold, adaptive setting of the bin number or an adaptive choice of cost-function depending on the properties of the component entropies H_A , H_F and H_{AF} .

Chapter 9

Conclusion

9.1 Summary

This thesis has presented work towards the successful registration of images in which conventional registration algorithms are ill-equipped to cope with local contrast changes. Chapter 2 presented an introduction to commonly used registration algorithms and described the cost-functions that are minimised (or maximised) during registration and the associated popular transformation models. Chapter 2 also detailed the motivation and physiological basis of dynamic contrast enhancement with exogenous Gd-DTPA. Chapter 4 presented a discussion of the limitations of conventional cost-functions ranging from the restrictive image intensity assumptions of the sum of squared difference cost-function to the more general description of image similarity given by normalised mutual information. Chapter 4 was used to provide evidence that registration failure is likely to be caused by inappropriate cost-function choice. Analysis was carried out using the novel inspection of cost-function minimisation spaces and the direct visual examination of cost-function gradient formation. In particular, Chapter 4 showed that the cost-functions discussed are likely to fail when local contrast changes occur since they are unable to distinguish between poor image similarity due to motion artefacts and changes to image similarity due to contrast enhancement. This discussion motivated the need for the Progressive Principal Component Registration (PPCR) algorithm in Chapter 5, which developed a model-free, full field-of-view registration technique without the need for model-fitting or segmentation. The PPCR algorithm was published (Melbourne et al., 2007b). The method allows improved registration performance by producing registration anchor images that are contrast-matched to their respective float images. As a result, the minimisation of a given cost-function between float image and PPCR anchor image is likely to be robust to contrast-enhancement mis-registration.

Chapter 4 also introduced a generalisation of image similarity for groups of images in the form of the Cost-Function Matrix Mean. This allowed a measure of the registration of all images within a dataset to one another. This was in contrast to the alternative formulation, comparing the registration of images to only a single image. The CFMM was used throughout the thesis as a measure of registration performance. Within the thesis, the formulation of the CFMM included values from the diagonal elements of the cost-function matrix, a true generalisation should only consider off-diagonal elements so that for the case of only two images, the CFMM is simply the cost-function value between the two images.

Chapter 3 developed a general simulator for producing dynamic contrast enhanced MRI datasets incorporating both image deformations and image contrast-change. The 2D model was built up of three main components: a motion model; a contrast-enhancement model with specific properties for simulating pathology; and a signal model for converting contrast-agent concentrations into MR image intensities. The model was designed with the evaluation of registration algorithms as its principal function in this thesis: results are used in Chapters 4, 5 and 8. The motion model is suitable for this purpose but the inclusion of a biomechanical model would be a desirable future step. The contrast enhancement model makes use of recent work on the representation of contrast-enhancement properties in the liver using suitable arterial input function modelling. Improvements to the enhancement curve modelling could be made using a more sophisticated image segmentation and a fuzzy tissue classification to vary the pharmacokinetic parameters. Additions to the signal modelling could be made that would allow for intra-image intensity variations and the analysis of the resulting k-space artefacts. A basic form of this 2D DCE-MRI simulator was presented in (Melbourne et al., 2008a), a full 3D extension of the work might prove to be useful as future work.

Chapter 5 presented results toward the explanation for improved registration performance using PPCR by discussing the generation of contrast-matched images to which registration can proceed using conventional cost-functions. The reduced influence of contrast-enhancement on image similarity allows registration using an unmodified registration algorithm internal to the PPCR algorithm. This was shown by inspection of the joint-entropy values calculated from datasets generated with varying motion corruption and contrast enhancement. The iterative nature of the PPCR algorithm allows refinements to the registration to be made with the inclusion of higher order principal components. The PPCR algorithm was initially tested on a simulated dynamic contrast enhanced MRI dataset incorporating an elastic deformation model and a model of contrast enhancement as discussed in Chapter 3. Variation of the strength and periodic nature of the motion model and the dose of contrast-agent revealed circumstances in which the PPCR algorithm was expected to allow improved registration to those methods that do not make allowance for local contrast changes. Part of this Chapter was published in (Melbourne et al., 2008a).

Chapter 5 also investigated the inclusion of some computational performance benefits into the PPCR algorithm. The time-limiting step is the individual registrations, run after the inclusion of each principal component. In particular, it was found that the number of iterations in each registration can be reduced whilst the set of registrations at each PPCR iteration can be run in parallel. The further parallelisation of the algorithm by running each registration on a graphical processing unit (GPU) will also make the algorithm fast. GPU-parallelisation of the fluid registration algorithm is a particularly desirable step; if the Cahill (Cahill et al., 2007b) method is used, the algorithm can make use of existing GPU algorithms for the Fast Fourier Transform and for convolution making registration much more rapid than on a conventional single CPU. This use of the fluid algorithm has yet to be implemented and would be beneficial to wider registration applications.

Chapter 6 applies the PPCR algorithm to real DCE-MRI data making use of the datasets described in Chapter 1. The benefits of the PPCR algorithm on visual appearance are discussed and the use of the

algorithm in statistical testing of motion within a dataset was presented. Visual inspection of difference images and visual comparison of registration methods by trained observers were used as method evaluation. Further analysis using a generalised cost-function for multiple image datasets (Cost Function Matrix Mean) and crude pharmacokinetic curve-fitting suggested that the PPCR algorithm gives more acceptable results to the direct fluid registration algorithm. The application of PPCR to 3D DCE-MRI datasets reveals an improvement in their registration over the use of a direct registration algorithm. The use of complex model-fitting using the independent MRIW software (d'Arcy et al., 2006) to the contrast-agent concentration time-curves showed improvement after registration by PPCR when inspecting the reduction in residual model-fit. Both the PPCR algorithm and the direct registration method showed superior registration (or resistance to mis-registration) than that seen using a proprietary scanner-based registration package.

Application of the PPCR algorithm to a superficially quite different application was presented in Chapter 7. On closer inspection, the reasons for failure of registration algorithms on both DCE-MRI and DW-MRI are analogous - local contrast changes disrupt the assumptions of conventional cost-functions which may lead to either under-registration or mis-registration. The improvement on image registration of DW-MRI datasets produced by the PPCR algorithm were outlined in this chapter. Analysis is carried out using visual inspection of registered slices and analysis of down-stream fractional anisotropy statistics. Analysis of the fractional anisotropy used a leave-one-out calculation of each diffusion direction and reveals reduced variation in fractional anisotropy variability. Part of this work was published in (Melbourne et al., 2008b). The PPCR algorithm is also applied to the registration of diffusion weighted images of the liver. Here the analysis concerned differences in diffusion b -value rather than gradient direction. As a result PPCR is not theoretically necessarily needed, as discussed in Chapter 7. However, inspection of the residual of model-fitting to the signal as a function of b -value, PPCR appeared at least as good as direct registration to a the zero-weighted b_0 image. Improved registration by PPCR may be the result of the algorithm acting to register all images towards a common coordinate frame rather than towards the coordinate frame of a single image, hence the CFMM will give a higher result when compared to direct registration.

Chapter 8 presented a novel algorithm for direct influence of cost-function gradients during image registration. The formation of this algorithm is a parallel development to the PPCR algorithm. It is developed in answer to problems found when there are large changes in image contrast profile (such as bolus arrival in the heart) relative to the remainder of images in the dataset. The PPCR method should be applied to groups of images under contrast enhancement but extreme changes in contrast in one or two images lead to principal components that do not necessarily enable contrast-matching of these images; hence this might result in mis-registration. Therefore the algorithm may be used complementary to the PPCR method. Analysis of the Kullback-Leibler divergence between joint image histograms in which both contain similar levels of motion, but only one contains contrast-enhancement allows the removal of contrast-enhancement by image matching and patching using the Kullback Leibler Assisted image Matching and Patching technique (KLAMP). Registration with reduced contrast-enhancement

influenced artefacts may proceed using either a standard registration algorithm of the pre-enhancement image and the KLAMPed post-enhancement image or by incorporation of the KLAMP algorithm into the registration algorithm to influence the driving force gradients.

Overall this thesis has demonstrated the limitation of image registration in difficult circumstances. The thesis has shown the influence of contrast-enhancement on cost-function gradient formulation by using simulated motion corrupted, and contrast-enhanced, MRI developed specifically for registration testing. Two new algorithms have been proposed, designed specifically for the registration of contrast enhanced data; performance benefits have been demonstrated for both. The PPCR algorithm has been demonstrated on subject data with both contrast variation due to extrinsic contrast-agent injection and diffusion gradient direction selection. Further benefits of the PPCR algorithm may be envisaged when considering the registration of any group of images into a common coordinate frame; some evidence was provided that this may be the case for diffusion weighted MRI using multiple b-values. The contribution of these algorithms has stretched the realm of application of image registration and the process of algorithm development during this thesis has revealed the requirements for novel, flexible image registration algorithms. The development of the PPCR and KLAMP algorithms allows improved registration in circumstances previously considered liable to mis-registration.

9.2 Future Work

A natural extension of the work outlined in this thesis is the combination of the PPCR and KLAMP algorithms. In cases in which contrast change is well represented by early principal components, the PPCR algorithm should operate without KLAMP. However, in those circumstances where one or two images contain contrast changes unlike those in the remainder of the dataset, the KLAMP algorithm may need to be incorporated to mediate against large local contrast variation. Some of the work in Chapter 4 and Chapter 8 may allow the adaptive selection of a cost-function and the parameters of the cost-function. For instance, a suitable cost-function from those listed in Chapter 2 and Appendix B could be determined by inspecting the component entropies H_A , H_F and H_{AF} or adapting the number and breadth of the intensity bins.

An alternative to Principal Components Analysis for data analysis is Independent Components Analysis. Independent Components are found by considering their contribution to a signal by source separation. Due to the fact the components are independent, they have no preferred ordering in terms of importance nor relative scaling (see (Milles et al., 2008)). These factors make PCA a preferable option for use in the PPCR algorithm. A generalisation of PCA is the Principal Geodesic Analysis concerning shapes (Fletcher et al., 2004). Further investigation of this generalisation in the context of image registration, and the corresponding generalisation of the PPCR algorithm could provide interesting work but this is left to the interested reader.

The work in this thesis addresses only motion artefacts between images (PPCR is an inter-image registration method). During an MRI acquisition, there may be intra-image motion corruption. Recent work by White *et al* (White et al., 2008) has addressed this, but only for images without contrast enhancement. If contrast-agent intensity changes within the image acquisition are small, this method will

be applicable in its current form. It may also be straightforward to approximate linear-contrast variations. In the spirit of the paper by White *et al* it may be beneficial to use a set of contrast-change training images, formed using a pre-bolus contrast agent injection. On injection of the full-bolus, the training images may be used to correct intra-image motion corruption in the presence of local contrast changes. In addition to the PPCR algorithm, this would allow full intra- and inter- image registration of a DCE-MRI dataset. There may be some scope for a direct intra-image deformation recovery scheme using PPCR by operating a PCA scheme (or more suitable data-reduction scheme) on the separate k-space fragments. However, some basic preliminary work on this idea was inconclusive.

Image registration in general does not incorporate biological information into the modelling of deformation. B-spline or fluid registration algorithms treat all pixels as equal when imposing a transformation. This is primarily due to the difficulties and computational complexity of building individual large region of interest physiological motion models. Some types of motion are particularly difficult to implement; a major example is sliding motions which are present in the abdomen under the influence of the breathing cycle. A registration algorithm that incorporated these features could benefit from techniques used in the computer game industry to track objects and provide realistic physics under which the objects are influenced. This may include environmental forces and prevention of undesirable mesh intersections.

This thesis has purely addressed MRI imaging modalities. However, MRI techniques have a wide diversity in the properties that they measure. The PPCR algorithm does not intrinsically require MR images on which to operate; the method should be widely applicable to other modalities. One interesting application may actually be the combination of images from different modalities of the same features. This would be a good extension of the work presented here and there are no direct theoretical problems other than those relating to good principal component formulation. Recent work on perfusion tensor imaging (Frank *et al.*, 2008) and also functional diffusion tensor imaging (analysing changes to diffusion caused by repeated nerve fibre activation) may also provide a further future application which would combine the results of Chapter 6 and 7.

The model of DCE-MRI presented in Chapter 3 is purely macroscopic. Future work could investigate the microscopic concentration changes in contrast agent. By modelling diffusion of contrast agent along a concentration gradient between vascular and extra-cellular, extra-vascular compartments, it may be possible to directly predict MR signal intensities given the vascular properties of the local region, taking into account parameters such as vessel size, permeability and tortuosity.

Appendix A

Formulation of Fluid Equation

For the interested reader, the following is a derivation of the fluid equation. The viscous fluid equation used in image registration is found from this argument.

Starting by considering the force on a given unit piece of fluid, the total force \mathbf{F}_T is given by the pressure gradient across the piece of fluid (Equation A.1). We also add a general term for a potential term, for instance if our fluid existed in a gravitational field we would add a term dependent on its relative height in relation to neighbouring pieces of fluid and as a function of local mass density ρ .

$$\mathbf{F}_T = -\nabla P \quad (\text{A.1})$$

$$\mathbf{F}_T = -\nabla P - \rho \nabla \psi \quad (\text{A.2})$$

We also include a continuity condition that says the total amount of fluid is conserved, if fluid moves away from a particular point, the amount of fluid remaining will go down!

$$\nabla \cdot (\rho \mathbf{v}) = -\frac{\delta \rho}{\delta t} \quad (\text{A.3})$$

We now consider the total force \mathbf{F}_T on the left-hand side. This is the full derivative rather than the partial $\rho \frac{\delta \mathbf{v}}{\delta t}$ since we are considering a particular piece of fluid *as it moves* not a particular spatial point through which the fluid is passing. Hence we must use the full derivative for $\rho \frac{d\mathbf{v}}{dt}$, Equation A.5.

$$\mathbf{F}_T = \rho \frac{d\mathbf{v}}{dt} \quad (\text{A.4})$$

$$\mathbf{F}_T = \rho \left[(\mathbf{v} \cdot \nabla) \mathbf{v} + \frac{\delta \mathbf{v}}{\delta t} \right] \quad (\text{A.5})$$

Next we add in a viscous component that describes the fluids resistance to being deformed. The fluid has no resistance to shear but will slip past itself (Equation A.6). We form the viscous term from the stress tensor \mathbf{S}_{ij} which can be derived from considering two plates encompassing a piece of fluid. If we move one plate by applying a force relative to the other, we consider a linear change in velocity between the two plates (the relative velocity of the fluid at each plate is zero. The linear relationship is governed by the viscosity parameter μ . The second viscosity parameter λ is required to make the stress tensor complete and allows for a response to internal forces.

$$\mathbf{F}_T = -\nabla P - \rho \nabla \psi + \mathbf{F}_{visc} \quad (\text{A.6})$$

$$\mathbf{S}_{ij} = \mu \left[\frac{\delta \mathbf{v}_i}{\delta x_j} + \frac{\delta \mathbf{v}_j}{\delta x_i} \right] + \lambda \delta_{ij} (\nabla \cdot \mathbf{v}) \quad (\text{A.7})$$

$$\mathbf{F}_{visc,i} = \sum_{j=1}^3 \frac{\delta \mathbf{S}_{ij}}{\delta x_j} \quad (\text{A.8})$$

For an isotropic medium (Equations A.9 and A.10), the derivative of the stress tensor has only second order components and so may be expressed as Equation A.11 which is the most general form of a function consisting only of second-order derivatives - if the medium was anisotropic we would require modifications as shown previously in Equation 3.16.

$$\frac{\delta \mu}{\delta x} = 0 \quad (\text{A.9})$$

$$\frac{\delta \lambda}{\delta x} = 0 \quad (\text{A.10})$$

$$\mathbf{F}_{visc} = \mu \nabla^2 \mathbf{v} + (\mu + \lambda) \nabla (\nabla \cdot \mathbf{v}) \quad (\text{A.11})$$

Putting together the terms, the entire fluid equation from Equation A.6 appears in Equation A.12 and particular assumptions can be made depending on the desired properties of the fluid. Equation A.12 includes an additional spatially dependent force term $\mathbf{F}_{sim}(\mathbf{u})$ which for our purposes is generated by an image similarity measure.

$$\rho \left[(\mathbf{v} \cdot \nabla) \mathbf{v} + \frac{\delta \mathbf{v}}{\delta t} \right] = -\nabla P - \rho \nabla \psi + \mu \nabla^2 \mathbf{v} + (\mu + \lambda) \nabla (\nabla \cdot \mathbf{v}) + \mathbf{F}_{sim}(\mathbf{u}) \quad (\text{A.12})$$

In medical image fluid registration, our fluid is considered viscous so we can ignore inertial and pressure terms as small or slowly varying and Equation A.12 reduces to Equation A.13. However it is important to recognise the above assumptions that went into the derivation in order to arrive at the result in Equation A.13.

$$\mu \nabla^2 \mathbf{v} + (\mu + \lambda) \nabla (\nabla \cdot \mathbf{v}) + \mathbf{F}_{sim}(\mathbf{u}) = 0 \quad (\text{A.13})$$

Appendix B

Formulation of Cost-Function Gradients

B.1 Change in Sum of Squared Difference with Pixel Displacement

The gradient of the local sum of squared difference can be found by the derivation shown in Equations B.1 to B.3. The final result is equivalent to that used by Christensen (Christensen et al., 1996). The parameters are defined as follows for each pixel n within all pixels N : \mathbf{A} is the anchor image, $\mathbf{F}(\mathbf{v})$ is the float image deformed under the transformation \mathbf{v} . We seek the derivative of the function over the transformation.

$$SSD = \sum_{n=1}^N (\mathbf{A}_n - \mathbf{F}_n(\mathbf{v}))^2 \quad (\text{B.1})$$

$$\frac{dSSD_n}{dv} = 2(\mathbf{A}_n - \mathbf{F}_n(\mathbf{v})) \frac{d}{dv} \mathbf{F}_n(\mathbf{v}) \quad (\text{B.2})$$

$$\nabla SSD = (\mathbf{A} - \mathbf{F}) \nabla(\mathbf{F}) \quad (\text{B.3})$$

B.2 Change in Cross Correlation with Pixel Displacement

The un-normalised local gradient of the cross-correlation cost-function can be derived as in Equations B.4 to B.5. Again, the parameters are defined as follows for each pixel n within all pixels N : \mathbf{A} is the anchor image, $\mathbf{F}(\mathbf{v})$ is the float image deformed under the transformation \mathbf{v} . We seek the derivative of the function over this transformation.

$$CC = \sum_{n=1}^N (\mathbf{A}_n \cdot \mathbf{F}_n(\mathbf{v})) \quad (\text{B.4})$$

$$\frac{dCC_n}{dv} = \mathbf{A}_n \cdot \frac{d}{dv} \mathbf{F}_n(\mathbf{v}) \quad (\text{B.5})$$

$$\nabla CC = \mathbf{A} \nabla \mathbf{F} \quad (\text{B.6})$$

The result of Equation B.6 can be seen to be one of the terms from the SSD Equation B.3 and therefore we may expect registration performance to be similar. Discrepancies occur for large biases in the $\mathbf{F} \nabla \mathbf{F}$ (self-similarity) term.

If normalisation is included, the gradient must be modified for the normalisation as in Equations B.7 to B.9.

$$CC = \frac{1}{\|\mathbf{A}\| \|\mathbf{F}(\mathbf{v})\|} \sum_{n=1}^N (\mathbf{A}_n \cdot \mathbf{F}_n(\mathbf{v})) \quad (\text{B.7})$$

$$\frac{dCC_n}{dv} = \frac{1}{\|\mathbf{A}\|\|\mathbf{F}(\mathbf{v})\|} \mathbf{A}_n \cdot \frac{d}{dv} \mathbf{F}_n(\mathbf{v}) - \quad (\text{B.8})$$

$$\frac{1}{2} \frac{1}{\|\mathbf{A}\|} \left[\frac{1}{\|\mathbf{F}_n(\mathbf{v})\|} \right]^3 (\mathbf{A}_n \cdot \mathbf{F}_n(\mathbf{v})) \sum_{i=1}^N \frac{d}{dv} \mathbf{F}_i(\mathbf{v}) \quad (\text{B.9})$$

$$\text{since, } \frac{1}{\|\mathbf{F}(\mathbf{v})\|} = \frac{1}{\sqrt{\sum_{n=1}^N \mathbf{F}_n(\mathbf{v})}} \quad (\text{B.10})$$

$$\frac{d}{dv} \frac{1}{\|\mathbf{F}(\mathbf{v})\|} = -\frac{1}{2} \left(\frac{1}{\sqrt{\sum_{n=1}^N \mathbf{F}_n(\mathbf{v})}} \right)^3 \sum_{n=1}^N \frac{d}{dv} \mathbf{F}_n(\mathbf{v}) \quad (\text{B.11})$$

B.3 Change in Joint Entropy with Pixel Displacement

The following is a derivation of the procedure used in finding the gradient of an information theoretic cost function. The result of this derivation is presented by Crum in (Crum et al., 2003) without discussion of the assumptions made in its derivation. The final result allows cost-function gradients to be calculated extremely rapidly when compared to methods that are more mathematically robust and continuous, for instance using a Parzen window method, or generalising the derivation used here to a partial volume derivation.

We start with the derivation for Joint Entropy. Pixels located in an image contribute to a particular intensity bin in the image histogram and joint image histogram. For two intensity bins with histogram counts n_1 and n_2 in a histogram with N total counts, their contribution to the entropy \mathbf{S} is:

$$\text{JointEntropy} = \sum_{i=1}^{\text{bins}} \frac{n_i}{N} \log \frac{n_i}{N} \quad (\text{B.12})$$

$$JE = \frac{n_1}{N} \log \frac{n_1}{N} + \frac{n_2}{N} \log \frac{n_2}{N} + \dots \quad (\text{B.13})$$

$$JE = \frac{n_1}{N} \log \frac{n_1}{N} + \frac{n_2}{N} \log \frac{n_2}{N} + S_0 \quad (\text{B.14})$$

We now imagine spatial shifting a particular pixel by one pixel width (perhaps to the right). This moves a unit value in the joint image histogram between two bins (with arbitrary locations in the joint image histogram). We decrease the value of one intensity bin and increase the other. Equation B.15 demonstrates this by moving a pixel count from bin n_2 to bin n_1 , hence the total joint entropy is now given by this equation. The remaining contribution of all other bins to the joint entropy is summarised by S_0 .

$$JE = \frac{n_1 + 1}{N} \log \frac{n_1 + 1}{N} + \frac{n_2 - 1}{N} \log \frac{n_2 - 1}{N} + S_0 \quad (\text{B.15})$$

Expanding...

$$\frac{n_1}{N} \log \frac{n_1 + 1}{N} + \frac{1}{N} \log \frac{n_1 + 1}{N} + \frac{n_2}{N} \log \frac{n_2 - 1}{N} - \frac{1}{N} \log \frac{n_2 - 1}{N} \quad (\text{B.16})$$

And again...

$$\frac{n_1}{N} \log \frac{n_1}{N} + \frac{n_1}{N} \log \left(1 + \frac{1}{n_1}\right) \quad (\text{B.17})$$

$$+ \frac{1}{N} \log \frac{n_1}{N} + \frac{1}{N} \log \left(1 + \frac{1}{n_1}\right) \quad (\text{B.18})$$

$$+ \frac{n_2}{N} \log \frac{n_2}{N} + \frac{n_2}{N} \log \left(1 - \frac{1}{n_2}\right) \quad (\text{B.19})$$

$$- \frac{1}{N} \log \frac{n_2}{N} - \frac{1}{N} \log \left(1 - \frac{1}{n_2}\right) \quad (\text{B.20})$$

We see that two terms in B.20 are the original entropy value, hence,

$$\Delta JE = \frac{n_1}{N} \log\left(1 + \frac{1}{n_1}\right) + \frac{1}{N} \log \frac{n_1}{N} \quad (\text{B.21})$$

$$+ \frac{1}{N} \log\left(1 + \frac{1}{n_1}\right) + \frac{n_2}{N} \log\left(1 - \frac{1}{n_2}\right) \quad (\text{B.22})$$

$$- \frac{1}{N} \log \frac{n_2}{N} - \frac{1}{N} \log\left(1 - \frac{1}{n_2}\right) \quad (\text{B.23})$$

Re-writing (to expand the n_1 terms)

$$\Delta JE = \frac{1}{N} \left[n_1 \log\left(1 + \frac{1}{n_1}\right) + \log n_1 - \log N + n_1 \log\left(1 + \frac{1}{n_1}\right) \right] \quad (\text{B.24})$$

$$+ n_2 \log\left(1 - \frac{1}{n_2}\right) + \log N - \log n_2 - \log\left(1 - \frac{1}{n_2}\right)] \quad (\text{B.25})$$

Doing some cancellation & rearrangement gives:

$$\Delta S = \frac{1}{N} \left[n_1 \log\left(1 + \frac{1}{n_1}\right) + \log\left(n_1\left(1 + \frac{1}{n_1}\right)\right) \right] \quad (\text{B.26})$$

$$+ n_2 \log\left(1 - \frac{1}{n_2}\right) - \log n_2 - \log\left(1 - \frac{1}{n_2}\right)] \quad (\text{B.27})$$

A final round of rearrangement gives:

$$\Delta JE = \frac{1}{N} \left[n_1 \log\left(1 + \frac{1}{n_1}\right) + \log(n_1 + 1) \right] \quad (\text{B.28})$$

$$+ n_2 \log\left(1 - \frac{1}{n_2}\right) - \log(n_2 - 1)] \quad (\text{B.29})$$

We now see that for a stable value of the change in joint entropy we require that both n_1 and n_2 are big. This also means that we can approximate the *logs* using $\log(1 + x) = x$ for small x . this can be shown to reduce to Equation B.30.

$$\Delta S = \frac{1}{N} \left[\frac{1}{n_1} + \frac{1}{n_2} + \log \frac{n_1}{n_2} \right] \quad (\text{B.30})$$

The log ratio term dominates for large n_1 and n_2 leaving us only to consider the fractional change for moving the pixel a little way dx Equation B.34. We now implement the same process for $-dx$ which would involve a third intensity bin n_3 . This is the result shown in (Crum et al., 2003).

$$\Delta JE^+ = \frac{1}{N} \log \frac{n_1}{n_2} \quad (\text{B.31})$$

$$\Delta JE^- = \frac{1}{N} \log \frac{n_1}{n_3} \quad (\text{B.32})$$

$$\Delta JE^+ - \Delta JE^- = \frac{1}{N} \log \frac{n_1}{n_2} - \frac{1}{N} \log \frac{n_1}{n_3} \quad (\text{B.33})$$

$$\frac{dJE}{dx} = \frac{1}{N} \log \frac{n_3}{n_2} \quad (\text{B.34})$$

B.4 Change in Mutual Information with Pixel Displacement

For Mutual Information we proceed with an identical analysis for the effect of a pixel shift on the marginal entropy of the float image. We include the result from above for the joint image histogram bin counts n_2 and n_3 and also the corresponding change to the float image marginal entropy by moving

a pixel and its effect on the associated bin counts m_2 and m_3 . The normalisation is by the total number of pixels, N , which is the same for both the marginal and joint image histograms.

$$MI = H_A + H_F - H_{AF} \quad (\text{B.35})$$

$$MI = H_A + \sum_{i=1}^{bins} \frac{m_i}{N} \log \frac{m_i}{N} - \sum_{i=1}^{bins} \frac{n_i}{N} \log \frac{n_i}{N} \quad (\text{B.36})$$

$$\frac{dMI}{dx} = \frac{1}{N} \log \frac{m_3}{m_2} - \frac{1}{N} \log \frac{n_3}{n_2} \quad (\text{B.37})$$

$$\frac{dMI}{dx} = \frac{1}{N} \log \frac{m_3 n_2}{m_2 n_3} \quad (\text{B.38})$$

B.5 Change in Normalised Mutual Information with Pixel Displacement

The expression for Normalised Mutual information is given by a combination of Equation B.34 and B.38 to give Equation B.42

$$NMI = \frac{H_A + H_F}{H_{AF}} \quad (\text{B.39})$$

$$NMI = \frac{H_A + H_F - H_{AF} + H_{AF}}{H_{AF}} \quad (\text{B.40})$$

$$\frac{dNMI}{dx} = \frac{d}{dx} \left[\frac{MI}{JE} + 1 \right] \quad (\text{B.41})$$

$$\frac{dNMI}{dx} = \frac{1}{H_{AF}^2} \left[JE \frac{dMI}{dx} - MI \frac{dJE}{dx} \right] \quad (\text{B.42})$$

$$(\text{B.43})$$

B.6 Alternative Information Based Cost-Functions

Registration using Normalised Mutual Information is equivalent to maximising the symmetric uncertainty; how well **A** predicts **F** and vice versa (as discussed in Chapter 2). Historically, registration proceeds by deforming the float image **F** to match the anchor image **A** so that a symmetric cost-function is not necessary since we only deform in one direction. If we were to maximise the uncertainty coefficients individually we are able to choose the direction which is deemed more meaningful. Here we consider the uncertainty coefficients of the pixel intensity probability distributions. The results are shown for each uncertainty coefficient in Equations B.44 to B.47. If we consider the uncertainty of the anchor image **A** given the float image **F** we get the force gradient Equation B.45 which is equivalent to maximising the mutual information as in Equation B.38 up to a scale factor. The alternative is to consider the uncertainty of the float image **F** given the anchor image **A** (Equation B.46). Finding the force gradients results in Equation B.47; this is equivalent to the Normalised Mutual Information gradient, apart from the substitution of the marginal entropy H_F for the joint entropy H_{AF} . To summarise, if registration proceeds such that the float image is deformed towards the anchor, we may use an asymmetric cost function, in the case described here we are maximising how *certain* we are of the float image given the anchor image. If we were to deform the float and anchor images towards each other, perhaps a symmetric coefficient is more meaningful, particularly if we desire an invertible or diffeomorphic deformation; in this case NMI may be more suitable for maximisation.

$$U(A|F) = \frac{H_A + H_F - H_{AF}}{H_A} \quad (\text{B.44})$$

$$\frac{dU(A|F)}{dx} = \frac{1}{H_A} \frac{dMI}{dx} \quad (\text{B.45})$$

$$U(F|A) = \frac{H_A + H_F - H_{AF}}{H_F} \quad (\text{B.46})$$

$$\frac{dU(F|A)}{dx} = \frac{1}{H_F^2} \left[H_F \frac{dMI}{dx} - MI \frac{dH_F}{dx} \right] \quad (\text{B.47})$$

Appendix C

Formation of a Fluid-based Image Registration Algorithm

The results of the two previous appendices may now be combined. From Appendix B, for a given cost-function we have expressions that give a force map that seeks to maximise (or minimise) the similarity measure. From Appendix A we have an expression for the viscous fluid equation from which is solved for the velocity of the registration correction \mathbf{v} from the similarity measure force \mathbf{F} (Equation C.1).

$$\mu \nabla^2 \mathbf{v} + (\mu + \lambda) \nabla (\nabla \cdot \mathbf{v}) = \mathbf{F} \quad (\text{C.1})$$

The solution for the velocity given the force \mathbf{F} at a given iteration is found using the method developed by Cahill (Cahill et al., 2007b) for a zero-boundary condition. This is analogous to the solution for the elastic deformation used in Chapter 2 and the method is identical making use of the Fast Sine Transform coded in Numerical Recipes (Press et al., 2007). The algorithm is presented in Table C.1, where we update using the full derivative for $d\mathbf{u}$ since we are considering the velocity of elements of fluid, rather than the velocities at fixed points in the lattice.

Table C.1: Algorithm for fluid registration (see text for discussion)

Given two images:

Initial Displacement $\mathbf{u}_{old} = 0$.

Start Loop

- 1) Find force gradients \mathbf{F} using Appendix B.
- 2) If images are too similar or force gradients too small, break loop.
- 3) Solve Equation C.1 for velocity field given result of 1.
- 4) Update displacement field using full derivative: $\mathbf{u}_{new} = \mathbf{u}_{old} + \mathbf{v}dt - dt \sum_{i=1}^3 \mathbf{v}_{xi} \frac{d\mathbf{u}}{dx_i}$
- 5) Transform float image by current displacement.

End Loop

Output Result

Appendix D

Formation of a KLAMP capable Fluid-based Image Registration Algorithm

Here we present the internal implementation of the KLAMP algorithm as discussed in Chapter 8. Instead of the masking and patching of the float image as a preprocessing step, we calculate the image mask and patch at each iteration but use the masked and patched image to calculate the joint entropy force gradient component prior to the solution of the velocity field using the fluid equation in Equation D.1. The gradient formulation outlined in Appendix B is adjusted so that the marginal probability distributions used in H_A and H_F are calculated using the unmodified float and anchor, but the joint probability distribution used to form $H_{A,F}$ uses the masked and patch float image $F' = M \cdot F + P$ to give $MI = H_A + H_F - H_{A,M \cdot F + P}$. The resulting algorithm is presented in Table D.1.

$$\mu \nabla^2 \mathbf{v} + (\mu + \lambda) \nabla (\nabla \cdot \mathbf{v}) = \mathbf{F} \quad (\text{D.1})$$

Table D.1: Algorithm for KLAMPed fluid registration (see text for discussion)

Given three images, \mathbf{A}_1 , \mathbf{A}_2 and a training image, \mathbf{A}_3 :

Initial Displacement $\mathbf{u}_{old} = 0$.

Start Loop

- 1) Calculate the float image mask, M and patch, P using KLAMP.
- 2) Find force gradients \mathbf{F} combining $H_{\mathbf{A}_1}$, $H_{\mathbf{A}_2}$ and $H_{\mathbf{A}_1, \mathbf{M} \cdot \mathbf{A}_2 + P}$ using Appendix B.
- 3) If images are too similar or force gradients too small, break loop.
- 4) Solve Equation D.1 for velocity field given result of 2.
- 5) Update displacement field using full derivative: $\mathbf{u}_{new} = \mathbf{u}_{old} + \mathbf{v}dt - dt \sum_{i=1}^3 \mathbf{v}_{xi} \frac{d\mathbf{u}}{dx_i}$
- 6) Transform float image \mathbf{A}_2 by current displacement.

End Loop

Output Result

Appendix E

List of Movies included on supplementary CD

This appendix contains descriptions of the movie files included on the attached CD as supplementary material.

Chapter 2

- *Movie-2-01.avi* Illustration of cost-function values with horizontal displacement of two (identical) brain images. Note that all cost-function values are normalised to fall in the range 0-1.

Chapter 3

- *Movie-3-01.avi* Illustration of motion model (coronal).
- *Movie-3-02.avi* Illustration of enhancement model (coronal).
- *Movie-3-03.avi* Illustration of combined motion and enhancement model (sagittal).

Chapter 6

All movie files have the same format from left to right. The left-most movie is the original, unregistered DCE-MRI dataset. The second is registration by direct fluid registration using cross-correlation to the first (unenhanced) image in the dataset. The third is the result of registration using the PPCR algorithm. If a fourth movie exists, this is the result of a scanner-based image registration algorithm. Since the data from Table 1.2 is 3D, only the central slice is shown.

- *Movie-6-01.avi* DCE-MRI movie for Patient 2 from Table 1.1.
- *Movie-6-02.avi* DCE-MRI movie for Patient 3 from Table 1.1.
- *Movie-6-03.avi* DCE-MRI movie for Patient 5 from Table 1.1.
- *Movie-6-04.avi* DCE-MRI movie for Patient 1 from Table 1.2.
- *Movie-6-05.avi* DCE-MRI movie for Patient 5 from Table 1.2.
- *Movie-6-06.avi* DCE-MRI movie for Patient 6 from Table 1.2.

Bibliography

- Acton, P. D., Pilowsky, L. S., Suckling, J., Brammer, M. J. & Ell, P. J. (1997), "Registration of dynamic dopamine D2 receptor images using principal component analysis.", *Eur J Nucl Med* **24**(11), 1405–1412.
- Adluru, G., DiBella, E. V. R. & Schabel, M. C. (2006), "Model-based registration for dynamic cardiac perfusion MRI.", *J Magn Reson Imaging* **24**(5), 1062–1070.
- Alexander, D. C., Pierpaoli, C., Basser, P. J. & Gee, J. C. (2001), "Spatial transformations of diffusion tensor magnetic resonance images.", *IEEE Trans Med Imaging* **20**(11), 1131–1139.
- Alexander, D., Gee, J. & Bajcsy, R. (1999), Elastic Matching of Diffusion Tensor MRIs, in "IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'99)", Vol. 1, p.1244.
- Alvarez (2000), "Reliable Estimation of Dense Optical Flow Fields with Large Displacements", *International Journal of Computer Vision* **39**(1), 41–56.
- Andersen, C., Taagehj, J. F., Mhler, A. & Rehling, M. (1996), "Approximation of arterial input curve data in MRI estimation of cerebral blood-tumor-barrier leakage: comparison between Gd-DTPA and 99mTc-DTPA input curves.", *Magn Reson Imaging* **14**(3), 235–241.
- Armitage, P., Behrenbruch, C., Brady, M. & Moore, N. (2005), "Extracting and visualizing physiological parameters using dynamic contrast-enhanced magnetic resonance imaging of the breast.", *Med Image Anal* **9**(4), 315–329.
- Atkinson, D., Hill, D. L., Stoye, P. N., Summers, P. E. & Keevil, S. F. (1997), "Automatic correction of motion artifacts in magnetic resonance images using an entropy focus criterion.", *IEEE Trans Med Imaging* **16**(6), 903–910.
- Bai, Y. & Alexander, D. C. (2008), Model-based registration to correct for motion between acquisitions in diffusion MR imaging, in "ISBI", pp.947–950.
- Bajcsy, R. & Kovacic, S. (1989), "Multiresolution elastic matching", *Computer Vision, Graphics and Image Processing* **46**, 1–21.
- Bammer, R., Acar, B. & Moseley, M. E. (2003), "In vivo MR tractography using diffusion imaging.", *Eur J Radiol* **45**(3), 223–234.

- Barber, D. C. & Hose, D. R. (2005), "Automatic segmentation of medical images using image registration: diagnostic and simulation applications.", *J Med Eng Technol* **29**(2), 53–63.
- Basser, P. J. & Pierpaoli, C. (1996), "Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI.", *J Magn Reson B* **111**(3), 209–219.
- Basser, P. J., Mattiello, J. & LeBihan, D. (1994), "Estimation of the effective self-diffusion tensor from the NMR spin echo.", *J Magn Reson B* **103**(3), 247–254.
- Batchelor, P. G., Atkinson, D., Hill, D. L. G., Calamante, F. & Connelly, A. (2003), "Anisotropic noise propagation in diffusion tensor MRI sampling schemes.", *Magn Reson Med* **49**(6), 1143–1151.
- Batchelor, P. G., Atkinson, D., Irarrazaval, P., Hill, D. L. G., Hajnal, J. & Larkman, D. (2005), "Matrix description of general motion correction applied to multishot images.", *Magn Reson Med* **54**(5), 1273–1280.
- Baudelet, C. & Gallez, B. (2002), "How does blood oxygen level-dependent (BOLD) contrast correlate with oxygen partial pressure (pO₂) inside tumors?", *Magn Reson Med* **48**(6), 980–986.
- Behrens, T. E. J., Woolrich, M. W., Jenkinson, M., Johansen-Berg, H., Nunes, R. G., Clare, S., Matthews, P. M., Brady, J. M. & Smith, S. M. (2003), "Characterization and propagation of uncertainty in diffusion-weighted MR imaging.", *Magn Reson Med* **50**(5), 1077–1088.
- Blackall, J. M., Ahmad, S., Miquel, M. E., McClelland, J. R., Landau, D. B. & Hawkes, D. J. (2006), "MRI-based measurements of respiratory motion variability and assessment of imaging strategies for radiotherapy planning.", *Phys Med Biol* **51**(17), 4147–4169.
- Bookstein, F. L. (1989), "Principal Warps: Thin-Plate Splines and the Decomposition of Deformations", *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**, 567–585.
- Brix, G., Semmler, W., Port, R., Schad, L., Layer, G. & Lorenz, W. (1990), Parameterization of the MRI signal enhancement during and after iv. infusion of Gd-DTPA: A Pharmacokinetic Model, in "9th Annual Meeting of the Society of Magnetic Resonance in Medicine", number 758 in "2".
- Buckley, D. L. (2002), "Uncertainty in the analysis of tracer kinetics using dynamic contrast-enhanced T1-weighted MRI.", *Magn Reson Med* **47**(3), 601–606.
- Buckley, D. L., Bui, J. D., Phillips, M. I. & Blackband, S. J. (1999), "MRI measurement of cell volume fraction in the perfused rat hippocampal slice.", *Magn Reson Med* **42**(3), 603–607.
- Buckley, D. L., Kerslake, R. W., Blackband, S. J. & Horsman, A. (1994), "Quantitative analysis of multi-slice Gd-DTPA enhanced dynamic MR images using an automated simplex minimization procedure.", *Magn Reson Med* **32**(5), 646–651.
- Buonaccorsi, G. A., Roberts, C., Cheung, S., Watson, Y., Davies, K., Jackson, A., Jayson, G. C. & Parker, G. J. M. (2005), "Tracer kinetic model-driven registration for dynamic contrast enhanced

- MRI time series.”, *Med Image Comput Comput Assist Interv Int Conf Med Image Comput Comput Assist Interv* **8**(Pt 1), 91–98.
- Buonaccorsi, G. A., Roberts, C., Cheung, S., Watson, Y., O’Connor, J. P. B., Davies, K., Jackson, A., Jayson, G. C. & Parker, G. J. M. (2006), “Comparison of the performance of tracer kinetic model-driven registration for dynamic contrast enhanced MRI using different models of contrast enhancement.”, *Acad Radiol* **13**(9), 1112–1123.
- Cahill, N. D., Noble, J. A. & Hawkes, D. J. (2007a), Fourier Methods for Nonparametric Image Registration, in “Computer Vision and Pattern Recognition, 2007”.
- Cahill, N. D., Noble, J. A., Hawkes, D. J. & Ray, L. A. (2007b), Fast Fluid Registration with Dirichlet Boundary Conditions, in “International Symposium on Biomedical Imaging”, pp.712–715.
- Chan, C., Barratt, D., Edwards, P., Penney, G., Slomczykowski, M., Carter, T. & Hawkes, D. (2004), Cadaver Validation of the Use of Ultrasound for 3D Model Instantiation of Bony Anatomy in Image Guided Orthopaedic Surgery, in “Medical Image Computing and Computer-Assisted Intervention MICCAI 2004”, Vol. 3217, pp.397–404.
- Chang, L.-C., Jones, D. K. & Pierpaoli, C. (2005), “RESTORE: robust estimation of tensors by outlier rejection.”, *Magn Reson Med* **53**(5), 1088–1095.
- Chiang, M., Leow, A. D., Klunder, A. D., Dutton, R. A., Barysheva, M., Rose, S., McMahon, K. L., de Zubizaray, G. I., Toga, A. W. & Thompson, P. M. (2008), “Fluid Registration of Diffusion Tensor Images Using Information Theory”, *IEEE-TMI* .
- Choyke, P. L., Dwyer, A. J. & Knopp, M. V. (2003), “Functional tumor imaging with dynamic contrast-enhanced magnetic resonance imaging.”, *J Magn Reson Imaging* **17**(5), 509–520.
- Christensen, G. E., Joshi, S. C. & Miller, M. I. (1997), “Volumetric transformation of brain anatomy.”, *IEEE Trans Med Imaging* **16**(6), 864–877.
- Christensen, G., Rabbitt, R. & Miller, M. (1996), “Deformable templates using large deformation kinematics”, *IEEE Transactions on Image Processing* **5**, 1435–1447.
- Commowick, O., Arsigny, V., Isambert, A., Costa, J., Dhermain, F., Bidault, F., Bondiau, P.-Y., Ayache, N. & Malandain, G. (2008), “An efficient locally affine framework for the smooth registration of anatomical structures.”, *Med Image Anal* .
- Cootes, T., Twining, C., Babalola, K. & Taylor, C. (2008), “Diffeomorphic statistical shape models”, *Image and Vision Computing* **26**, 326–332.
- Crum, W. R., Hill, D. L. G. & Hawkes, D. J. (2003), “Information theoretic similarity measures in non-rigid registration.”, *Inf Process Med Imaging* **18**, 378–387.

- Crum, W. R., Scahill, R. I. & Fox, N. C. (2001), "Automated hippocampal segmentation by regional fluid registration of serial MRI: validation and application in Alzheimer's disease.", *Neuroimage* **13**(5), 847–855.
- Crum, W. R., Tanner, C. & Hawkes, D. J. (2005), "Anisotropic multi-scale fluid registration: evaluation in magnetic resonance breast imaging.", *Phys Med Biol* **50**(21), 5153–5174.
- Cuenod, C. A., Fournier, L., Balvay, D. & Guinebretire, J.-M. (2006), "Tumor angiogenesis: pathophysiology and implications for contrast-enhanced MRI and CT assessment.", *Abdom Imaging* **31**(2), 188–193.
- D'Agostino, E., Maes, F., Vandermeulen, D. & Suetens, P. (2003), "A viscous fluid model for multimodal non-rigid image registration using mutual information.", *Med Image Anal* **7**(4), 565–575.
- d'Arcy, J. A., Collins, D. J., Padhani, A. R., Walker-Samuel, S., Suckling, J. & Leach, M. O. (2006), "Informatics in Radiology (infoRAD): Magnetic Resonance Imaging Workbench: analysis and visualization of dynamic contrast-enhanced MR imaging data.", *Radiographics* **26**(2), 621–632.
- de Bazelaire, C. M. J., Duhamel, G. D., Rofsky, N. M. & Alsop, D. C. (2004), "MR imaging relaxation times of abdominal and pelvic tissues measured in vivo at 3.0 T: preliminary results.", *Radiology* **230**(3), 652–659.
- Donahue, K. M., Burstein, D., Manning, W. J. & Gray, M. L. (1994), "Studies of Gd-DTPA relaxivity and proton exchange rates in tissue.", *Magn Reson Med* **32**(1), 66–76.
- Duhamel, G., Schlaug, G. & Alsop, D. C. (2006), "Measurement of arterial input functions for dynamic susceptibility contrast magnetic resonance imaging using echoplanar images: comparison of physical simulations with in vivo results.", *Magn Reson Med* **55**(3), 514–523.
- Feynman, R., Leighton, R. B. & Sands, M. (1998), *Lectures on Physics*, Addison Wesley.
- Fletcher, P. T., Lu, C., Pizer, S. M. & Joshi, S. (2004), "Principal geodesic analysis for the study of nonlinear statistics of shape.", *IEEE Trans Med Imaging* **23**(8), 995–1005.
- Frank, L. R., Lu, K. & Wong, E. C. (2008), "Perfusion Tensor Imaging.", *Magn Reson Med* **60**(6), 1284–1291.
- Gennert & Negahdaripour (1987), "Relaxing the Brightness Constancy Assumption in Computing Optical Flow", *Massachusetts Institute of Technology* .
- George, R., Vedam, S. S., Chung, T. D., Ramakrishnan, V. & Keall, P. J. (2005), "The application of the sinusoidal model to lung cancer patient respiratory motion.", *Med Phys* **32**(9), 2850–2861.
- Gray, H. . (1918), *Gray's Anatomy*, Philadelphia: Lea & Febiger.

- Guimond, A., Guttman, C., Warfield, S. & Westin, C. (2002), Deformable registration of DT-MRI data based on transformation invariant tensor characteristic, *in* "IEEE International Symposium on Biomedical Imaging", pp.761–764.
- Harrer, J. U., Parker, G. J. M., Haroon, H. A., Buckley, D. L., Embelton, K., Roberts, C., Balriaux, D. & Jackson, A. (2004), "Comparative study of methods for determining vascular permeability and blood volume in human gliomas.", *J Magn Reson Imaging* **20**(5), 748–757.
- Hayton, P., Brady, M., Tarassenko, L. & Moore, N. (1997), "Analysis of dynamic MR breast images using a model of contrast enhancement.", *Med Image Anal* **1**(3), 207–224.
- Hecke, W. V., Leemans, A., DAgostino, E., Backer, S. D., Vandervliet, E., Parizel, P. M. & Sijbers, J. (2007), "Nonrigid Coregistration of Diffusion Tensor Images Using a Viscous Fluid Model and Mutual Information", *IEEE TRANSACTIONS ON MEDICAL IMAGING* **26**, 1598–1612.
- Hermosillo, G. (2002), Variational Methods for Mutlimodal Image Matching, PhD thesis, Universite de Nice - Sophia Antipolis.
- Hill, D. L., Batchelor, P. G., Holden, M. & Hawkes, D. J. (2001), "Medical image registration.", *Phys Med Biol* **46**(3), R1–45.
- Hill, D. L., Studholme, C. & Hawkes, D. J. (1994), Voxel similarity measures for automated image registration, *in* R. A. Robb (ed.), "Proc. SPIE Visualization in Biomedical Computing", Vol. 2359, pp.205–216.
- Horn, B. K. & Schunck, B. G. (1981), "Determining Optical Flow", *ARTIFICIAL INTELLIGENCE* **17**, 185–203.
- Jiang, L., Zhao, D., Constantinescu, A. & Mason, R. P. (2004), "Comparison of BOLD contrast and Gd-DTPA dynamic contrast-enhanced imaging in rat prostate tumor.", *Magn Reson Med* **51**(5), 953–960.
- Kingsley, P. B. (2006), "Introduction to diffusion tensor imaging mathematics: Part III. Tensor calculation, noise, simulations, and optimization", *Concepts in Magnetic Resonance Part A* **28A**, 155–179.
- Knopp, M. V., Weiss, E., Sinn, H. P., Mattern, J., Junkermann, H., Radeleff, J., Magener, A., Brix, G., Delorme, S., Zuna, I. & van Kaick, G. (1999), "Pathophysiologic basis of contrast enhancement in breast tumors.", *J Magn Reson Imaging* **10**(3), 260–266.
- Koh, D.-M. & Collins, D. J. (2007), "Diffusion-weighted MRI in the body: applications and challenges in oncology.", *AJR Am J Roentgenol* **188**(6), 1622–1635.
- Kwee, T. C., Takahara, T., Koh, D.-M., Niveststein, R. A. J. & Luijten, P. R. (2008), "Comparison and reproducibility of ADC measurements in breathhold, respiratory triggered, and free-breathing diffusion-weighted MR imaging of the liver.", *J Magn Reson Imaging* **28**(5), 1141–1148.

- Larsson, H. B., Stubgaard, M., Frederiksen, J. L., Jensen, M., Henriksen, O. & Paulson, O. B. (1990), "Quantitation of blood-brain barrier defect by magnetic resonance imaging and gadolinium-DTPA in patients with multiple sclerosis and brain tumors.", *Magn Reson Med* **16**(1), 117–131.
- Leach, M. O., Brindle, K. M., Evelhoch, J. L., Griffiths, J. R., Horsman, M. R., Jackson, A., Jayson, G. C., Judson, I. R., Knopp, M. V., Maxwell, R. J., McIntyre, D., Padhani, A. R., Price, P., Rathbone, R., Rustin, G. J., Tofts, P. S., Tozer, G. M., Vennart, W., Waterton, J. C., Williams, S. R., Workman, P. & Pharmacodynamic/Pharmacokinetic Technologies Advisory Committee, D. D. O. R. U. K. (2005), "The assessment of antiangiogenic and antivasular therapies in early-stage clinical trials using magnetic resonance imaging: issues and recommendations.", *Br J Cancer* **92**(9), 1599–1610.
- Leemans, A., Sijbers, J., Backer, S. D., Vandervliet, E. & Parizel, P. M. (2005), "Affine Coregistration of Diffusion Tensor Magnetic Resonance Images Using Mutual Information", *Lecture Notes in Computer Science* **3708**, 523–530.
- Lester, H., Arridge, S. R., Jansons, K. M., Lemieux, L., Hajnal, J. V. & Oatridge, A. (1998), Non-linear Registration with the Variable Viscosity Fluid Algorithm, in "Lecture Notes in Computer Science: Information Processing in Medical Imaging", Vol. 1613, pp.238–251.
- Little, J., Hill, D. & Hawkes, D. (1997), "Deformation incorporating rigid structures.", *Computer vision and image understanding* **66**, 223–32.
- Liu, Z. & Bilston, L. (2000), "On the viscoelastic character of liver tissue: experiments and modelling of the linear behaviour.", *Biorheology* **37**(3), 191–201.
- Martel, A. L., Froh, M. S., Brock, K. K., Plewes, D. B. & Barber, D. C. (2007), "Evaluating an optical-flow-based registration algorithm for contrast-enhanced magnetic resonance imaging of the breast", *Phys. Med. Biol* **52**, 3803–3816.
- McClelland, J. R., Blackall, J. M., Tarte, S., Chandler, A. C., Hughes, S., Ahmad, S., Landau, D. B. & Hawkes, D. J. (2006), "A continuous 4D motion model from multiple respiratory cycles for use in lung radiotherapy.", *Med Phys* **33**(9), 3348–3358.
- McLeish, K., Hill, D. L. G., Atkinson, D., Blackall, J. M. & Razavi, R. (2002), "A Study of the Motion and Deformation of the Heart Due to Respiration", *IEEE TRANSACTIONS ON MEDICAL IMAGING* **9**, 1142–1150.
- McRobbie, D. W., Moore, E. A., Graves, M. J. & Prince, M. R. (2006), *MRI from Picture to Proton (2nd Edition)*, Cambridge.
- Melbourne, A., Atkinson, D., White, M. J., Collins, D. J., Leach, M. O. & Hawkes, D. J. (2007a), Using Registration to Quantify the Consistency of Whole Liver Position during Patient Breath-hold in Dynamic Contrast-Enhanced MRI 2007, 3709., in "ISMRM", number 3709 in "1".

- Melbourne, A., Atkinson, D., White, M. J., Collins, D., Leach, M. & Hawkes, D. (2007b), "Registration of dynamic contrast-enhanced MRI using a progressive principal component registration (PPCR).", *Phys Med Biol* **52**(17), 5147–5156.
- Melbourne, A., Hawkes, D. & Atkinson, D. (2008a), "Influence of Organ Motion and Contrast Enhancement on Image Registration", *MICCAI 2008 LNCS* **1**, 1.
- Melbourne, A., Hawkes, D. & Atkinson, D. (2008b), Non-Rigid Registration of Diffusion Weighted MRI Using Progressive Principal Component Registration (PPCR), in "ISMRM", number 3097 in "1".
- Melhem, E. R., Mori, S., Mukundan, G., Kraut, M. A., Pomper, M. G. & van Zijl, P. C. M. (2002), "Diffusion tensor MR imaging of the brain and white matter tractography.", *AJR Am J Roentgenol* **178**(1), 3–16.
- Milles, J., van der Geest, R. J., Jerosch-Herold, M., Reiber, J. H. C. & Lelieveldt, B. P. F. (2008), "Fully automated motion correction in first-pass myocardial perfusion MR image sequences.", *IEEE Trans Med Imaging* **27**(11), 1611–1621.
- Narayanan, R., Fessler, J. A., Park, H. & Meyerl, C. R. (2005), "Diffeomorphic nonlinear transformations: a local parametric approach for image registration.", *Inf Process Med Imaging* **19**, 174–185.
- Orton, M. R., Collins, D. J., Walker-Samuel, S., d'Arcy, J. A., Hawkes, D. J., Atkinson, D. & Leach, M. O. (2007), "Bayesian estimation of pharmacokinetic parameters for DCE-MRI with a robust treatment of enhancement onset time.", *Phys Med Biol* **52**(9), 2393–2408.
- Orton, M. R., d'Arcy, J. A., Walker-Samuel, S., Hawkes, D. J., Atkinson, D., Collins, D. J. & Leach, M. O. (2008), "Computationally efficient vascular input function models for quantitative kinetic modelling using DCE-MRI.", *Phys Med Biol* **53**(5), 1225–1239.
- Parker, G. J. M. & Alexander, D. C. (2003), "Probabilistic Monte Carlo based mapping of cerebral connections utilising whole-brain crossing fibre information.", *Inf Process Med Imaging* **18**, 684–695.
- Parker, G. J. M., Roberts, C., Macdonald, A., Buonaccorsi, G. A., Cheung, S., Buckley, D. L., Jackson, A., Watson, Y., Davies, K. & Jayson, G. C. (2006), "Experimentally-derived functional form for a population-averaged high-temporal-resolution arterial input function for dynamic contrast-enhanced MRI.", *Magn Reson Med* **56**(5), 993–1000.
- Parker, G. J., Suckling, J., Tanner, S. F., Padhani, A. R., Husband, J. E. & Leach, M. O. (1998), "MRIW: parametric analysis software for contrast-enhanced dynamic MR imaging in cancer.", *Radiographics* **18**(2), 497–506.
- Passe, T. J., Bluemke, D. A. & Siegelman, S. S. (1997), "Tumor angiogenesis: tutorial on implications for imaging.", *Radiology* **203**(3), 593–600.

- Patankar, T. F., Haroon, H. A., Mills, S. J., Balriaux, D., Buckley, D. L., Parker, G. J. M. & Jackson, A. (2005), "Is volume transfer coefficient (K(trans)) related to histologic grade in human gliomas?", *AJNR Am J Neuroradiol* **26**(10), 2455–2465.
- Pennec, X., Cachier, P. & Ayache, N. (1999), Understanding the Demons Algorithm: 3D Non-rigid Registration by Gradient Descent, in "MICCAI 1999, Lecture Notes in Computer Science 1679", p.597606.
- Pluim, J. P., Maintz, J. B. & Viergever, M. A. (2000), "Image registration by maximization of combined mutual information and gradient information.", *IEEE Trans Med Imaging* **19**(8), 809–814.
- Pluim, J. P. W., Maintz, J. B. A. & Viergever, M. A. (2003), "Mutual-information-based registration of medical images: a survey.", *IEEE Trans Med Imaging* **22**(8), 986–1004.
- Port, R. E., Knopp, M. V. & Brix, G. (2001), "Dynamic Contrast-Enhanced MRI Using Gd-DTPA: Interindividual Variability of the Arterial Input Function and Consequences for the Assessment of Kinetics in Tumors", *Magnetic Resonance in Medicine* **45**, 1030–1038.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. (2007), *Numerical Recipes: Third Edition*, Cambridge.
- Provenzale, J. M., Mukundan, S. & Barboriak, D. P. (2006), "Diffusion-weighted and perfusion MR imaging for brain tumor characterization and assessment of treatment response.", *Radiology* **239**(3), 632–649.
- Rinck, P. A. & Muller, R. N. (1999), "Field strength and dose dependence of contrast enhancement by gadolinium-based MR contrast agents.", *Eur Radiol* **9**(5), 998–1004.
- Roberts, C., Buckley, D. L. & Parker, G. J. M. (2006a), "Comparison of errors associated with single- and multi-bolus injection protocols in low-temporal-resolution dynamic contrast-enhanced tracer kinetic analysis.", *Magn Reson Med* **56**(3), 611–619.
- Roberts, C., Issa, B., Stone, A., Jackson, A., Waterton, J. C. & Parker, G. J. M. (2006b), "Comparative study into the robustness of compartmental modeling and model-free analysis in DCE-MRI studies.", *J Magn Reson Imaging* **23**(4), 554–563.
- Rohlfing, T., Maurer, C. R., Bluemke, D. A. & Jacobs, M. A. (2003), "Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint.", *IEEE Trans Med Imaging* **22**(6), 730–741.
- Rohlfing, T., Maurer, C. R., O'Dell, W. G. & Zhong, J. (2004), "Modeling liver motion and deformation during the respiratory cycle using intensity-based nonrigid registration of gated MR images.", *Med Phys* **31**(3), 427–432.

- Rueckert, D., Sonoda, L. I., Hayes, C., Hill, D. L., Leach, M. O. & Hawkes, D. J. (1999), "Nonrigid registration using free-form deformations: application to breast MR images.", *IEEE Trans Med Imaging* **18**(8), 712–721.
- Sarvazyan, A., Goukassian, D., Maevsky, E. & Oranskaja, G. (1994), Elastic Imaging as a new Modality of Medical Imaging for Cancer Detection, in "International Workshop on Interaction of Ultrasound with Biological Media", p.6981.
- Schnabel, J. A., Tanner, C., Castellano-Smith, A. D., Degenhard, A., Leach, M. O., Hose, D. R., Hill, D. L. G. & Hawkes, D. J. (2003), "Validation of nonrigid image registration using finite-element methods: application to breast MR images.", *IEEE Trans Med Imaging* **22**(2), 238–247.
- Stefanescu, R., Pennec, X. & Ayache, N. (2004), "Grid powered nonlinear image registration with locally adaptive regularization.", *Med Image Anal* **8**(3), 325–342.
- Stegmann, M. B., Olafsdttir, H. & Larsson, H. B. W. (2005), "Unsupervised motion-compensation of multi-slice cardiac perfusion MRI.", *Med Image Anal* **9**(4), 394–410.
- Stejskal & Tanner (1965), "Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient", *The Journal of Chemical Physics* **1**, 288–298.
- Studholme, C., Hill, D. & Hawkes, D. (1999), "An overlap invariant entropy measure of 3D medical image alignment", *Pattern Recognition* **32**(1), 71–86.
- Sychra, J. J., Pavel, D. G., Chen, Y. & Jani, A. (1994), "The accuracy of SPECT brain activation images: propagation of registration errors.", *Med Phys* **21**(12), 1927–1932.
- Tanner, C., Schnabel, J., Chung, D., Clarkson, M., Rueckert, D., Hill, D. & Hawkes, D. (2000), "Volume and Shape Preservation of Enhancing Lesions when Applying Non-rigid Registration to a Time Series of Contrast Enhancing MR Breast Images", *Lect. Notes Comput. Sc.* **1935**, 327–337.
- Tanner, C., Schnabel, J., Degenhard, A., Castellano-Smith, A., Hayes, C., Leach, M., Hose, D., Hill, D. & Hawkes, D. (2002), "Validation of Volume-Preserving Non-rigid Registration: Application to Contrast-Enhanced MR-Mammography", *Lect. Notes Comput. Sc.* **2488**, 307–314.
- Tao, X. & Miller, J. V. (2006), "A method for registering diffusion weighted magnetic resonance images.", *LNCS MICCAI 2006* **9**(Pt 2), 594–602.
- Taylor, N. J., Baddeley, H., Goodchild, K. A., Powell, M. E., Thoumine, M., Culver, L. A., Stirling, J. J., Saunders, M. I., Hoskin, P. J., Phillips, H., Padhani, A. R. & Griffiths, J. R. (2001), "BOLD MRI of human tumor oxygenation during carbogen breathing.", *J Magn Reson Imaging* **14**(2), 156–163.
- Thirion, J. P. (1998), "Image matching as a diffusion process: an analogy with Maxwell's demons.", *Med Image Anal* **2**(3), 243–260.

- Tofts, P. S. (1997), "Modeling tracer kinetics in dynamic Gd-DTPA MR imaging.", *J Magn Reson Imaging* **7**(1), 91–101.
- Tofts, P. S. & Kermode, A. G. (1991), "Measurement of the blood-brain barrier permeability and leakage space using dynamic MR imaging. 1. Fundamental concepts.", *Magn Reson Med* **17**(2), 357–367.
- Tofts, P. S., Brix, G., Buckley, D. L., Evelhoch, J. L., Henderson, E., Knopp, M. V., Larsson, H. B., Lee, T. Y., Mayr, N. A., Parker, G. J., Port, R. E., Taylor, J. & Weisskoff, R. M. (1999), "Estimating kinetic parameters from dynamic contrast-enhanced T(1)-weighted MRI of a diffusable tracer: standardized quantities and symbols.", *J Magn Reson Imaging* **10**(3), 223–232.
- Totman, J. J., O'gorman, R. L., Kane, P. A. & Karani, J. B. (2005), "Comparison of the hepatic perfusion index measured with gadolinium-enhanced volumetric MRI in controls and in patients with colorectal cancer.", *Br J Radiol* **78**(926), 105–109.
- Vandecaveye, V., Keyzer, F. D. & Dymarkowski, S. (2007), "Perfusion-and diffusion-weighted imaging of hepatocellular carcinoma.", *JBR-BTR* **90**(6), 492–496.
- Vercauteren, T., Pennec, X., Malis, E., Perchant, A. & Ayache, N. (2007), "Insight into efficient image registration techniques and the demons algorithm.", *Inf Process Med Imaging* **20**, 495–506.
- Viola & Wells (1997), "Alignment by Maximization of Mutual Information", *International Journal of Computer Vision* **24**, 137154.
- Weinmann, H. J., Brasch, R. C., Press, W. R. & Wesbey, G. E. (1984), "Characteristics of gadolinium-DTPA complex: a potential NMR contrast agent.", *AJR Am J Roentgenol* **142**(3), 619–624.
- Weisenfeld, N. & Warfield, S. (2004), Normalization of joint image-intensity statistics in MRI using the Kullback-Leibler divergence, in "ISBI 2004", Vol. 1, pp.101–104.
- White, M. J., Atkinson, D. & Hawkes, D. (2008), Image Deformation Recovery Using Overlapping Partial Samples (IDROPS): Model-Based Respiratory Artefact Correction in Free-Breathing Liver MRI, in "International Society for Magnetic Resonance in Medicine", Vol. 16.
- Woolrich, M. W., Jenkinson, M., Brady, J. M. & Smith, S. M. (2004), "Fully Bayesian spatio-temporal modeling of fMRI data.", *IEEE Trans Med Imaging* **23**(2), 213–231.
- Xiaohua, C., Brady, M., Lo, J. L.-C. & Moore, N. (2005), Simultaneous Segmentation and Registration of Contrast-Enhanced Breast MRI, in "IPMI LNCS", Vol. 3565, pp.126–137.
- Zuo, C. S., Jiang, A., Buff, B. L., Mahon, T. G. & Wong, T. Z. (1996), "Automatic motion correction for breast MR imaging.", *Radiology* **198**(3), 903–906.