

Copyright is owned by the Author of the thesis. Permission is given for a copy to be downloaded by an individual for the purpose of research and private study only. The thesis may not be reproduced elsewhere without the permission of the Author.

IMAGE REGISTRATION UNDER CONFORMAL DIFFEOMORPHISMS

A THESIS PRESENTED IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE
DEGREE OF
DOCTOR OF PHILOSOPHY
IN
MATHEMATICS
AT MASSEY UNIVERSITY, PALMERSTON NORTH,
NEW ZEALAND.

Muhammad Yousuf Tufail

2017

Contents

Declaration	xviii
Acknowledgements	xix
1 Introduction	1
1.1 Image registration	1
1.1.1 Brief Overview of the Literature	2
1.2 Definitions	5
1.2.1 Continuous and discrete images	5
1.2.2 Diffeomorphism	5
1.2.3 Conformal diffeomorphism	5
1.2.4 Action of diffeomorphisms on images	6
1.3 Applications of image registration	8
1.3.1 Morphometrics	8
1.3.2 Medical imaging	9
1.3.3 Computer vision	10
1.3.4 Remote sensing	10
1.3.5 Visual Cortex	11
1.4 Motivation	11
1.5 Brief overview of the thesis	18
2 Finite Dimensional Image Registration	20
2.1 Image Registration in Practice	20
2.1.1 Missing values in image registration	22
2.1.2 Image interpolation	23
2.2 Image registration using the rigid group	25
2.2.1 Using coarse search	29
2.2.2 Using gradient descent	33
2.3 Möbius registration	35
2.3.1 Image Registration with the Möbius Group	37

3	Method of Control Points	45
3.1	Introduction	45
3.1.1	Smoothing	46
3.1.2	Image Registration Using a Conformal Diffeomorphism	47
3.2	Control points method	48
3.3	The Cauchy–Riemann equations	53
3.4	First discrete form of the Cauchy–Riemann equations	54
3.4.1	First form with smooth images	56
3.4.2	First form with non-smooth images	60
3.5	Second and third discrete forms of the Cauchy–Riemann equations	68
3.5.1	Second/Third form with smooth images	72
3.5.2	Second/third form with non-smooth images	74
3.6	Fourth form of the penalty term	80
3.6.1	Fourth form with smooth images	81
3.6.2	Fourth form with non-smooth images	84
4	Gradient Flow	94
4.1	A Gradient Flow Algorithm for Conformal Diffeomorphisms	95
4.1.1	Implementation of the gradient flow method	98
4.2	Image registration with smooth images	100
4.3	Image registration with non-smooth images	108
5	Experiments and Comparison	120
5.1	Comparing the Algorithms	147
6	Conclusions and Future Work	149
6.1	Conclusion	149
6.2	Future work	150

List of Tables

1.1	The transformation groups relevant to this thesis.	2
1.2	This table suggests the connection between the transformations group and Thompson's images in [101]. This table is taken from [66].	13
1.3	<i>Top:</i> Four positions are marked on middle finger. <i>Bottom:</i> Numerical results of the computation of the cross-ratio of points on the middle finger in man at different ages, taken from [81].	15
1.4	<i>Top:</i> Numeric results of cross-ratios for postnatal stages are given. <i>Bottom:</i> Numeric results of cross-ratios for antenatal stages are provided. These numerical results are taken from [81]	16
1.5	Computation of cross-ratio at (marked) position of the middle finger (see the figure of a human hand in Table 1.3).	18
3.1	Registration errors and penalty terms for Example 3.7. Here the same registration error as in all other tables is used: $RE = \ (I_1 \circ \varphi^{-1})(x_{ij}) - I_2(x_{ij})\ $. However, we refer to it by name to save space. P_2 is the second (linear) penalty term given in Equation (3.15), and P_3 is the third (nonlinear) penalty term given in Equation (3.19).	73
3.2	Registration errors and penalty terms for Example 3.8.	74
3.3	Registration errors and penalty terms for Example 3.9. (The images are smoothed in the first 5 rows, and unsmoothed in the last row.)	76
3.4	Registration errors and penalty terms for Example 3.10. (The images are smoothed in the first 5 rows, and unsmoothed in the last row.)	77
3.5	Registration errors and penalty terms for Example 3.11. (The images are smoothed in the first 5 rows, and unsmoothed in the last row.)	78

5.1 The Taylor series (for Example 5.4) found by four different runs. Each run used an initial guess given by the results of the $n = 4$ registration plus a small random perturbation (independent normal random variates with mean 0 and standard deviation 0.04). While the algorithm has located a local minimum in each case, the minima are all different, and the coefficients a_n for $n \geq 4$ carry essentially no information about the image registrations. 131

List of Figures

1.1	A set of discrete images. <i>Left column:</i> 2-D and 3-D colour images. <i>Right:</i> greyscale images.	6
1.2	<i>Left:</i> A uniform grid (100×100) is given. <i>Right:</i> The image of the uniform grid under the conformal diffeomorphism $\varphi(z) = \frac{az+b}{cz+d}$, $z \in [-0.5, 0.5] \times [-0.5, 0.5]$, where $a = 0.4i, b = 0.2 + 0.4i, c = 0.2i, d = 1$).	7
1.3	Position $y \in \Omega$ in the pixel grid is mapped under φ to position $x = \varphi(y)$ in the transformed grid. In order to obtain the intensity of x it is necessary to invert φ and recover $y = \varphi^{-1}(x)$	7
1.4	Examples of Thompson's transformations between related species. These images are taken from [7] (p. no. 404).	12
1.5	These images of skulls are taken from Thompson's book [101]	12
1.6	<i>Top:</i> Images of foot bones of an ox, a sheep and a giraffe. Four positions are marked in each image. These images are taken from [101]. <i>Bottom:</i> Distances and cross-ratios between the foot bones of an ox, a sheep and a giraffe.	14
1.7	Scanned images of [81] (a) An adult, (b) a five year old and (c) a newborn.	15
1.8	<i>Left:</i> Antenatal stages (in lunar months). <i>Right:</i> Postnatal stages (in years). These images are taken from [81].	16
1.9	<i>Top:</i> Four positions are marked on each image showing the appearance of a human at different ages. <i>Below:</i> Numerical results for the cross-ratios are given (taken from [73]).	17
2.1	Set of images indicating a perfect registration in which a transformation φ is applied to the source and the transformed source is apparently identical to the destination image or target. This transformed source is subtracted from the target (using Equation (2.1)) and a mid-grey screen is obtained which is labelled as 'Difference'.	21

2.2	Set of images indicating a poor registration because the transformed source is not aligned with the target. In fact, the transformed source is rotated in the wrong direction. The poor registration is demonstrated by the fact that the difference image is not blank.	21
2.3	Missing values are shown in a pair of images that are obtained after two different rotations of the source image of Figure 2.1. These missing values have been set to 0 (black).	22
2.4	A uniform grid represents sixteen grid points in which four known values of a function f are marked with red circles corresponding to a set of grid points. A point ‘P’ at which f needs to be determined is marked with a blue circle. Two green circles represent two points P_1 and P_2 on a vertical line passing through a point P at which the function value is calculated (with the help of linear interpolation) before its computation at a point P.	24
2.5	The LOF caption	25
2.6	Square grid of 10,000 points in the range $[-0.5, 0.5] \times [-0.5, 0.5]$ with $N = 100$, i.e., 100 points per side.	28
2.7	Images for Example 2.2. In the first row, the source image, the Gaussian $\exp(-5x^2 - 3y^2)$, is shown on the left; the target image, a rigid transformation of the source, is shown on the right. The second row contains corresponding contour plots of these images.	30
2.8	The LOF caption	31
2.9	Coarse search results for Example 2.2. All of the parameter combinations tested are listed along the x -axis, with the registration errors plotted on the y -axis. The red circle indicates the smallest minimum that the coarse search optimisation obtained.	31
2.10	The non-smooth images used in Example 2.3. The target is a rigid transformation of the source.	32
2.11	<i>Top:</i> Results of rigid registration using coarse search for Example 2.3. A perfect registration is obtained. <i>Bottom:</i> All values of the objective function for coarse search optimisation for Example 2.3.	32
2.12	The non-smooth images used in Example 2.4. The target is a rigid transformation of the source.	33
2.13	<i>Top:</i> Results of rigid registration for Example 2.4. A satisfactory, but imperfect, registration is obtained from the coarse search. <i>Bottom:</i> All values of the objective function for coarse search optimisation for Example 2.4. Note that the y -axis only goes down to 20, not to 0 as in the previous Examples.	33

2.14	Results of least-squares rigid registration for Example 2.4 using <i>lsqnonlin</i> with initial guess the identity.	34
2.15	Results of least-squares rigid registration for Example 2.4 using <i>lsqnonlin</i> with initial guess the best result of coarse search. A perfect registration is obtained.	35
2.16	Data for Example 2.5. The source is a Gaussian and the target is a Möbius transformation of the source.	37
2.17	Results of Möbius registration for Example 2.5, with <i>lsqnonlin</i> and initial guess the identity. A perfect registration is obtained.	38
2.18	Data for Example 2.6, Möbius registration with synthetic non-smooth images. The mapping that generates the source from the target is also shown.	38
2.19	Results of first attempt at Möbius registration for Example 2.6. The desired mapping and the mapping computed by <i>lsqnonlin</i> are also shown; the results are poor.	39
2.20	Selection of corresponding landmarks on both the fish images.	40
2.21	Results of the second attempt at Möbius registration for Example 2.6, using an initial guess calculated from landmark matching. A perfect registration is obtained.	40
2.22	Data for Example 2.7 of Möbius registration, a cartoon version of Thompson’s fish.	41
2.23	Selection of the first set of landmarks on both the images.	41
2.24	Results of Möbius registration for Example 2.7: first set of landmarks.	42
2.25	Selection of a second set of landmarks on both the images.	42
2.26	Results of Möbius registration for Example 2.7: second set of landmarks.	42
2.27	Second dataset for Example 2.7: swapped source and target.	43
2.28	Selection of corresponding landmarks: first set.	43
2.29	Results of Möbius registration for Example 2.7: second dataset, first set of landmarks.	43
2.30	Selection of corresponding landmarks: second set.	44
2.31	Results of Möbius registration for Example 2.7: second dataset, second set of landmarks.	44
3.1	<i>Left:</i> Example smooth images. <i>Right:</i> Corresponding contour plots of the smooth images.	46
3.2	Examples of non-smooth images.	47
3.3	Smoothed images from Figure 3.2 using a 15×15 Gaussian filter with standard deviation of 5 pixels.	47
3.4	Grid of 10,000 points with $N = 100$	48

3.5	Blue circles represent selected control points from the discrete domain S .	49
3.6	Red circles represent the transformed control points under the action of some φ^{-1} .	50
3.7	Bilinear interpolation over the transformed control points, marked with red circles, was used to generate this grid.	50
3.8	In the first grid, the blue circles represent selected control points from the discrete domain. In the second grid, the images of the control points under φ^{-1} are shown with red circles. Bilinear interpolation is used to generate the rest of the grid points between these red circles.	51
3.9	The source and target images for Example 3.1 are shown as greyscale images in the first row and as contour plots in the second row. The source is $I_1(x, y) = \exp(-7x^2 - 2y^2)$, and the target is defined by $I_2 = I_1 \circ \varphi^{-1}$, where $\varphi^{-1}(z) = 0.1 + z + 0.3z^2$.	57
3.10	The results of conformal image registration using Algorithm 2 applied to Example 3.1. The source and target images are shown, together with the final transformed source, which nearly matches the target on the part of the domain on which it is defined. The difference $I_2 - I_1 \circ \varphi^{-1}$ is shown in greyscale, where uniform mid-grey indicates a perfect match. The grid of the transformed control points, i.e. $\varphi^{-1}(\hat{x}_{ij})$ (we call it the deformation grid) and the contour plots of target and transformed source are shown on the right. In the table, each row gives the registration error and the value of the penalty term for a step of Algorithm 2.	57
3.11	Four grids are shown that are corresponding to $\lambda = 20^4, 20^3, 20^2, 20$ respectively for sixteen control points for Example 3.1. The two grids in the first row are rigid because squares are mapped to squares under φ^{-1} ; the third and fourth grids, with smaller values of λ , are nonrigid. The deformation grid corresponding to $\lambda = 20^2$ shows asymmetry, which we are unable to explain.	58
3.12	The source and target images for Example 3.2. The source is $\exp(-5x^2 - 7y^2)$, the target is $\exp(-5x^2 - 2y^2)$. The corresponding contour plots can be seen in the second row.	59
3.13	The results of conformal registration with Algorithm 2 applied to Example 3.2, showing a near-perfect registration because some of the red contours are slightly mismatched with the corresponding green contours.	59
3.14	The images for Example 3.3. The source I_1 is a standard reference image; the target is $I_2(z) = I_1(0.1 + 1.2z + 0.3z^2)$ and is therefore known to be conformally related to the source.	60

3.15	The results of the first attempt at registration for Example 3.3. An extremely poor registration is obtained.	61
3.16	Left: the original set of control points; right: the same points with a small perturbation added to each point independently in x and y	61
3.17	The LOF caption	62
3.18	Two non-smooth independent images of ellipses are used as the source and target respectively.	62
3.19	The results of the first attempt at registration for Example 3.4. The optimiser remains stuck at the initial condition, the identity.	63
3.20	The results of the second attempt at registration for Example 3.4. The registration is still not successful even after the control points are perturbed.	64
3.21	Left: the image matching results for Example 3.4 initialised with smoothed versions of the images. The top row shows the source and the target, while the transformed source and the difference between the transformed source and the target can be seen in second row. Right: the deformation grid $\varphi^{-1}(\hat{x}_{ij}), \hat{x}_{ij} \in S$ with $\lambda = 20$. In the table below, each row gives the registration error and value of penalty term P_1 for each step of Algorithm 2. In the first 5 rows the source and target images are smoothed by a Gaussian filter with standard deviation 5 pixels; in the last row they are not smoothed.	64
3.22	Modified versions of the images of Thompson’s fish are shown. The first image serves as the source and the second image as the target for image registration.	65
3.23	The LOF caption	66
3.24	The LOF caption	66
3.25	The LOF caption	67
3.26	In both the figures, the three red circles represent three complex numbers. The blue circles represent the fourth complex number for the linear (left) and nonlinear (right) forms.	71
3.27	Results of conformal image registration for Example 3.7 using the linear discrete Cauchy–Riemann equations as a penalty term.	72
3.28	Results of conformal image registration for Example 3.7 using the nonlinear discrete Cauchy–Riemann equations as a penalty term.	73
3.29	The grids in Example 3.7 corresponding to $\lambda = 20^4, 20^3, 20^2, 20$ for (left) the linear form of the penalty term and (right) for the nonlinear form. . . .	74
3.30	Results of conformal image registration for Example 3.8 using the linear discrete Cauchy–Riemann equations as a penalty term.	75

3.31	Results of conformal image registration for Example 3.8 using the non-linear discrete Cauchy–Riemann equations as a penalty term.	75
3.32	Results of conformal image registration for Example 3.9 using the linear discrete Cauchy–Riemann equations as a penalty term.	75
3.33	Results of conformal image registration for Example 3.9 using the non-linear discrete Cauchy–Riemann equations as a penalty term.	76
3.34	Results of conformal image registration for Example 3.10 using the linear discrete Cauchy–Riemann equations as a penalty term (top) and the nonlinear form (bottom).	77
3.35	Results of conformal image registration for Example 3.11 using the linear discrete Cauchy–Riemann equations as a penalty term (top) and the nonlinear form (bottom). Note that the size of the top grid is $[-8, -1] \times [1, 4]$ which is far away from the domain $[-0.5, 0.5] \times [-0.5, 0.5]$. However, the majority portion of the bottom grid is within the domain.	78
3.36	<i>Left:</i> A uniform square grid of size sixteen is presented. <i>Right:</i> Its image under the conformal map $\varphi^{-1}(z) = 1.5 + 0.2z - 0.4z^2 + 0.1z^3 - 0.1z^4$	80
3.37	The results of conformal image registration with the fourth penalty term applied to the images in Example 3.12.	82
3.38	Four grids are shown that were calculated for Example 3.12 and correspond to $\lambda = 20^4, 20^3, 20^2, 20$ respectively for sixteen control points. The first three grids in the first row represent perfect rigid grids because squares are mapped to squares in each grid without showing any deformation. The fourth grid is a non-rigid conformal grid.	82
3.39	Results of conformal image registration applied to Example 3.37 with the fourth penalty term and $\lambda = 150$	83
3.40	The LOF caption	83
3.41	Results of conformal image registration for Example 3.14 with the fourth penalty term and $\lambda = 150$	84
3.42	Results of conformal image registration for Example 3.15 with the fourth penalty term with $\lambda = 150$ are given.	85
3.43	<i>Top:</i> A series of transformed sources as λ decreases geometrically. <i>Bottom:</i> The corresponding grids.	86
3.44	<i>Top:</i> Series of transformed sources are given when λ is increasing geometrically. <i>Bottom:</i> Series of grids are presented corresponding to transformed sources. A striking global change in the grid is observed between $\lambda = 200$ and $\lambda = 300$	87
3.45	L-curve for Example 3.15.	88

3.46	<i>Top:</i> Deformation grid of Example 3.4 (which used the first form of the penalty term); red circles indicate the non-smooth or non-conformal parts. <i>Bottom:</i> Deformation grid of Example 3.15 using the 4th form of the penalty term.	89
3.47	Results of conformal image registration for Example 3.15 with the fourth penalty term with $\lambda = 150$	90
3.48	<i>Top:</i> A series of transformed sources as λ decreases geometrically for Example 3.16. <i>Bottom:</i> The corresponding grids. A marked deterioration in the grid is observed at the smallest values of λ	91
3.49	<i>Top:</i> A series of transformed sources as λ increases geometrically for Example 3.16. <i>Bottom:</i> The corresponding grids.	92
3.50	Red lines shows the path of error while λ is decreasing. On the other hand, path of error function while λ is increasing is given with the green line. $\lambda = 150$ is a threshold point where both paths coincide.	93
4.1	<i>Top:</i> The source, the ellipse $I(x, y) = \exp(-7x^2 - 2y^2)$ is shown on the left, while the right shows the target, which is generated from the source using $\varphi^{-1} = \sum_{k=0}^2 a_k z^k$ using coefficients $a_0 = 0.3, a_1 = 0.8, a_2 = 0.4$. <i>Bottom:</i> The corresponding contour plots.	101
4.2	Registration with target generated from smooth source. <i>Left:</i> The top row shows the source and target, while the transformed source and the difference between the transformed source and the target can be seen in the second row. <i>Right:</i> A deformation grid $\varphi^{-1}(z_{ij}), \forall z_{ij} \in S$ is shown on top and the corresponding contour plots of the transformed source and the target are on the bottom.	102
4.3	The error graph shows the successful convergence of Algorithm 3 (with $K = 3$) for Example 4.1.	102
4.4	<i>Top:</i> Two ellipses for Example 4.2, the source $I_1 = \exp(-7x^2 - 2y^2)$, and the target $I_2 = \exp(-6.5x^2 - 1.5y^2)$. <i>Bottom:</i> The corresponding contour plots of the source and the target images.	103
4.5	Registration with independent smooth images for Example 4.2. <i>Left:</i> The top row shows the source and target images, while the bottom row shows the transformed source and the difference image. <i>Right:</i> The deformation grid $\varphi^{-1}(z_{ij}), \forall z_{ij} \in S$ on top and the corresponding contour plots on the bottom.	103
4.6	The algorithm successfully found a local minimum in Example 4.2. Note that the y -axis on this plot does not go down to 0, but stops at 1. Axis extrema are selected to make the plot as clear as possible for all of these error plots.	104

4.7	<i>Top:</i> Two ellipses, the source $I_1 = \exp(-7x^2 - 2y^2)$, and the target $I_2 = \exp(-4x^2 - y^2)$ respectively for Example 4.3. <i>Bottom:</i> The corresponding contour plots of the source and the target images for the same example.	104
4.8	Registration results for Example 4.3 with four terms. Successful convergence with good registration is obtained.	105
4.9	Registration results for Example 4.3 with eight terms. The final value of the objective function is rather lower than with four terms.	106
4.10	Registration results for Example 4.3 with twelve terms. The final value of the objective function is not much better than with eight terms.	107
4.11	Set of images for Example 4.4. <i>Left:</i> The source image. <i>Right:</i> The target image, which is generated from the source using $\varphi^{-1} = \sum_{k=0}^2 a_k z^k$ with coefficients $a_0 = 0.1i, a_1 = 1.5, a_2 = 0.2$.	108
4.12	Registration with dependent, non-smooth images for Example 4.4 with $K = 3$. <i>Left:</i> The top row shows the source and target images, while the bottom row shows the transformed source and the difference image. <i>Right:</i> The deformation grid $\varphi^{-1}(z_{ij}), \forall z_{ij} \in S$.	109
4.13	Plots of the error function for Example 4.4 with three terms of the Taylor series. <i>Left:</i> Smoothed images (first step). <i>Right:</i> Original images (second step).	109
4.14	Four sets of image registrations (along with their corresponding error graphs) for Example 4.4 are given. Algorithm 3 is run with four different stopping criteria (for $K = 3$ in each case) based on a common ratio ($r = 20\sqrt{3}$). Results (c) and (d) display the successful convergence of Algorithm 3 along with good registration, whereas set (a) has clearly stopped too early. Note that the runs lower down take many more steps, this is particularly clear for the registration of the smoothed image in (d).	110
4.15	<i>Top:</i> An ellipse with axes 0.8 units and 0.4 units (on the left) serves as the source, while another ellipse of axes 0.6 units, 0.4 units (on the right) serves as the target. <i>Bottom:</i> Corresponding contour plots of the source and the target respectively. Note that the background is not perfectly white (which has caused the artefacts in the contour plot) because of the resizing and resampling of the images.	111
4.16	<i>Top:</i> Registration with four terms using conformal gradient flow on independent images related by a conformal transformation (Example 4.5). <i>Bottom:</i> Error plots for smoothed (left) and original (right) images.	112

4.17	<i>Top:</i> Registration with eight terms using conformal gradient flow on independent images related by a conformal transformation (Example 4.5). <i>Bottom:</i> Error plots for smoothed (left) and original (right) images. The final error is little different to the run with four terms.	113
4.18	<i>Top:</i> Registration with twelve terms using conformal gradient flow on independent images related by a conformal transformation (Example 4.5). <i>Bottom:</i> Error plots for smoothed (left) and original (right) images. The final error is better than those with four or eight terms.	114
4.19	<i>Top:</i> A circle of radius 0.5 units serves as the source (on the left). The ellipse on the right has axes 0.5 units and 0.3 units and serves as the target. <i>Bottom:</i> Corresponding contour plots of the source and the target respectively.	115
4.20	<i>Top:</i> Registration using conformal gradient flow with four terms in the Taylor series of images of a circle and ellipse. <i>Bottom:</i> Error graphs for smoothed (left) and non-smoothed (right) images.	116
4.21	<i>Top:</i> Registration using conformal gradient flow with eight terms in the Taylor series of images of a circle and ellipse. <i>Bottom:</i> Error graphs for smoothed (left) and non-smoothed (right) images.	117
4.22	<i>Top:</i> Registration using conformal gradient flow with twelve terms in the Taylor series of images of a circle and ellipse. <i>Bottom:</i> Error graphs for smoothed (left) and non-smoothed (right) images.	118
5.1	<i>Top:</i> The source and target for Example 5.1. <i>Bottom:</i> Corresponding contour plots.	121
5.2	<i>Left:</i> Results of image registration for Example 5.1 using the control points method with the 4th penalty term and $\lambda = 150$. Images are smoothed for first two rows and non-smooth for the last row. Note that the algorithm has found a spurious rotation of the circle as well as the correct scaling. There is nothing to penalise this rotation in the registration.	121
5.3	Results of image registration with the gradient flow method for Example 5.1. <i>Top left:</i> The source and the target images, the transformed source, and the difference image. <i>Top right:</i> The final mapping $\varphi^{-1}(x_{ij})$ and the contour plots of the transformed source and the target. <i>Bottom:</i> The progress of the gradient descent algorithm for the smoothed (left) and original (right) images.	122

5.4	<i>Left:</i> The source and target images for Example 5.2 are shown in first row. The source is $I_1(x, y) = \exp(-7x^2 - 2y^2)$, and the target is defined by $I_2 = I_1 \circ \varphi^{-1}$, where φ^{-1} , where $\varphi^{-1}(z) = 0.1 + z + 0.2z^2 + 0.5iz^2$. Corresponding contour plots are given in the second row. <i>Right:</i> The non-conformal mapping that generates the target.	123
5.5	<i>Top:</i> Results of image matching for Example 5.2 with the control points method. <i>Bottom:</i> Numerical results of the registration for $\lambda = 150$	123
5.6	Results of conformal registration using the gradient flow method (with four terms) for Example 5.2.	123
5.7	Progress of the gradient descent algorithm towards finding a minimum for Example 5.2.	124
5.8	<i>Left:</i> The source and target images for Example 5.3 are given. The target is defined by $I_2 = I_1 \circ \varphi^{-1}$, where $\varphi^{-1}(z) = 0.1 + z + 0.2z^2 + 0.1\bar{z}^3$. <i>Right:</i> The non-conformal map that generates the target.	124
5.9	Results of conformal registration using the control points method for Example 5.3.	125
5.10	<i>Top:</i> Results of conformal registration using the gradient flow method (with four terms) for Example 5.3. <i>Bottom:</i> Progress of the gradient descent algorithm towards finding a minimum for Example 5.3.	125
5.11	The source and target images for Example 5.4, the cartoon versions of Thompson’s fish that he believed were ‘isogonally’ related.	126
5.12	Results of conformal registration using the control points method for Example 5.4. Although the shapes are very different, an extremely good registration is obtained, with an invertible mapping, at $\lambda = 150$	126
5.13	Results of the continuation method for Example 5.4 for $25 \leq \lambda \leq 150$	127
5.14	Results of the continuation method for Example 5.4 for $150 \leq \lambda \leq 800$. The basic shape of the mapping remains unchanged over a wide range of λ values.	127
5.15	The L-curve for Example 5.4.	127
5.16	Registration results for Example 5.4 with the gradient flow method with four terms.	128
5.17	Registration results for Example 5.4 with the gradient flow method with eight terms.	129
5.18	Registration results for Example 5.4 with the gradient flow method with twelve terms.	129

5.19	The three grids found by the gradient flow method for Example 5.4 are plotted on the same graph. It can be seen that all three grids (with 4 terms in green, 8 terms in blue and 12 terms in black) display slightly different deformations.	130
5.20	Convergence of the gradient flow for Example 5.4; note that the y axis is different for the smooth case. <i>Top left:</i> Smoothed images, 4 terms; <i>top right:</i> original images, 4 terms; <i>bottom left:</i> original images, 8 terms; <i>bottom right:</i> original images, 12 terms. In each case the gradient flow ran successfully to a local minimum.	130
5.21	Registration results for Example 5.4 with the gradient flow method with six terms.	132
5.22	Registration results for Example 5.4 with the gradient flow method with nine terms.	132
5.23	Registration results for Example 5.4 with the gradient flow method with twelve terms.	133
5.24	The three grids found with six, nine and twelve terms for Example 5.4; compare to the somewhat different registrations found with 4, 8, and 12 terms shown in Figure 5.19.	133
5.25	Convergence of the gradient flow for Example 5.4. <i>Top left:</i> Smoothed images, 6 terms; <i>top right:</i> original images, 6 terms; <i>bottom left:</i> original images, 9 terms; <i>bottom right:</i> original images, 12 terms. The two graphs on the left do not show convergence to a local minimum because the gradient flow was stopped due to non-invertibility of the mapping. . . .	134
5.26	A pair of black and white images. <i>Left:</i> Image of a cat, which serves as the source. <i>Right:</i> Image of a chicken, which is used as the target. . . .	135
5.27	Results of conformal registration using the control points method for Example 5.5. A good registration is obtained at the cost of an irregular grid. Note that the first two values in the table are with the smoothed images, while the last row is not.	135
5.28	Results of the continuation method for Example 5.5 for $25 \leq \lambda \leq 150$. . .	136
5.29	Results of the continuation method for Example 5.5 for $150le\lambda \leq 800$. When λ is sufficiently large (greater than about 240) the mapping is invertible.	136
5.30	The L-curve for Example 5.5.	137
5.31	Registration results for Example 5.5 with the gradient flow method with four terms. The deformation is conformal, but the match is poor.	137

5.32	Registration results for Example 5.5 with the gradient flow method with eight terms. The registration is not markedly different to that with four terms.	138
5.33	Registration results for Example 5.5 with the gradient flow method with twelve terms. The registration is not markedly different to that with four or eight terms.	138
5.34	Three grids for Example 5.5 (with four, eight and twelve terms) are shown. Each grid shows a slightly different deformation.	139
5.35	Convergence of the gradient flow for Example 5.5. <i>Top left</i> : Smoothed images, 4 terms; <i>top right</i> : original images, 4 terms; <i>bottom left</i> : original images, 8 terms; <i>bottom right</i> : original images, 12 terms. For 12 terms the gradient flow stopped at $t = 0.0115$ (before locating a local minimum), due to the emergence of a non-invertible grid.	139
5.36	Five images of immature human skulls are presented at different ontogenetic times. These images of skulls are taken from [63].	140
5.37	Results of conformal registration using the control points for Example 5.6 at $\lambda = 150$, ‘Skull1’ (source) registered to ‘Skull4’ (target).	141
5.38	Results of the continuation method for Example 5.6 for $25 \leq \lambda \leq 150$	141
5.39	Results of the continuation method for Example 5.6 for $150 \leq \lambda \leq 800$	142
5.40	L-curve for Example 5.6, ‘Skull1’ registered to ‘Skull4’.	142
5.41	Results of conformal registration using the gradient flow method, ‘Skull1’ registered to ‘Skull4’.	143
5.42	Results of registration (using gradient flow) when ‘Skull1’ is registered with ‘Skull2’.	144
5.43	Results of registration (using gradient flow) when ‘Skull1’ is registered with ‘Skull3’.	145
5.44	Results of registration (using gradient flow) when ‘Skull1’ is registered with ‘Skull5’.	146
5.45	This graph indicates a smooth conformal pattern of skull growth in Example 5.6, with shape data relating the skull shapes (from the set of conformal mapping) projected to two dimensions. The data points are the absolute values of two Taylor coefficients in the conformal registration of ‘Skull1’ to each of the 5 skulls.	147
6.1	One of Thompson’s isogonally related fish examples. These scanned images are taken from [101] (p 1064).	150

Declaration

It is hereby declare that this dissertation is my own work. It is being submitted for the degree of Doctor of Philosophy in Mathematics at the Massey University, Palmerston North. It has not been submitted before for any degree or examination at this or any other institution.

Muhammad Yousuf Tufail
(Candidate's Signature)

Acknowledgements

Glory and Praise to Allah, the most gracious and merciful. *Indeed, the worst of living creatures in the sight of Allah are the deaf and dumb who do not use reason.*

- Al-Quran (8: 22)

Thousands and hundreds thousands salam to Prophet Muhammad Sall Allāhu ‘alayhi wa-sallam. *No two things have been combined better than knowledge and patience.*

- Sahih Bukhari (Book 11: Hadees 644)

I would like to express my deepest gratitude to my supervisors, Prof. Dr. Stephen Richard Marsland and Dist. Prof. Dr. Robert Ian McLachlan. This journey was so long that I thought it would never end, but they have been supportive and encouraging throughout my research. Their reassuring comments and strong belief in me were sometimes all that kept me going. I am truly indebted to them for their selfless time and unforgettable help during the writing process of this thesis. Their rational ideas, honest criticism, valuable suggestions and patience made this thesis possible. I am also very grateful to my supervisors for organizing many useful workshops where I got the chance to know amazing people. These workshops were a great source of learning through sharing and exchanging ideas with well-known researchers.

I gratefully acknowledge the Marsden Fund scholarship that I received towards my PhD at Massey University. This PhD study would not have been possible without their financial support. I would like to thank all technical and non-technical staff at SEAT for helping me through various stages.

I also appreciate the support and trust of my employer NED University of Engineering and Technology.

My heartfelt thanks to Abu, Ami and my siblings for always believing in me and encouraging me to follow my dreams. I am also thankful to my close friends who kept my sense of humour alive.

And finally to my sweetheart Hamza who wonders what I do all day and without whose never ending love, this thesis would have been finished in due time. And to my lovely wife who has been a source of encouragement and forbearance through thick and thin. She also provided motivation, often by asking: when are you going to submit this thesis?

Abstract

Image registration is the process of finding an alignment between two or more images so that their appearance matches. It has been widely studied and applied to several fields, including medical imaging and biology (where it is related to morphometrics). In biology, one motivation for image registration comes from the work of Sir D’Arcy Thompson. In his book *On Growth and Form* he presented several examples where a grid superimposed onto a two-dimensional image of one species was smoothly deformed to suggest a transformation to an image of another species. His examples include relationships between species of fish and comparison of human skulls with higher apes.

One of Thompson’s points was that these deformations should be as ‘simple’ as possible. In several of his examples, he uses what he calls an *isogonal* transformation, which would now be called conformal, i.e., angle-preserving. His claims of conformally-related change between species were investigated further by Petukhov, who used Thompson’s grid method as well as computing the cross-ratio (which is an invariant of the Möbius group, a finite-dimensional subgroup of the group of conformal diffeomorphisms) to check whether sets of points in the images could be related by a Möbius transformation. His results suggest that there are examples of growth and evolution where a Möbius transformation cannot be ruled out. In this thesis, we investigate whether or not this is true by using image registration, rather than a point-based invariant: we develop algorithms to construct conformal transformations between images, and use them to register images by minimising the sum-of-squares distance between the pixel intensities. In this way we can see how close to conformal the image relationships are.

We develop and present two algorithms for constructing the conformal transformation, one based on constrained optimisation of a set of control points, and one based on gradient flow. For the first method we consider a set of different penalty terms that aim to enforce conformality, based either on discretisations of the Cauchy-Riemann equations, or geometric principles, while in the second the conformal transformation is represented as a discrete Taylor series. The algorithms are tested on a variety of datasets, including synthetic data (i.e., the target is generated from the source using a known conformal transformation; the easiest possible case), and real images, including some that are not actually conformally related. The two methods are compared on a set of images that include Thompson’s fish example, and a small dataset demonstrating the growth of a human skull. The conformal growth model does appear to be validated for the skulls, but interestingly, not for Thompson’s fish.